

# Are “Market Neutral” Hedge Funds Really Market Neutral?

Andrew J. Patton\*

*London School of Economics*

First version: 11 March, 2004. This version: 28 October, 2004.

## Abstract

One can consider the concept of market neutrality for hedge funds as having breadth and depth: “breadth” reflects the number of market risks to which a fund is neutral, while “depth” reflects the “completeness” of the neutrality of the fund to market risks. We focus on market neutrality depth, and propose five different neutrality concepts. “Mean neutrality” nests the standard correlation-based definition of neutrality. “Variance neutrality”, “Value-at-Risk neutrality” and “tail neutrality” all relate to the neutrality of the risk of the hedge fund to market risks. Finally, “complete neutrality” corresponds to independence of the fund to market risks. We suggest statistical tests for each neutrality concept, and apply the tests to a combined database of monthly “market neutral” hedge fund returns from the HFR and TASS hedge fund databases. We find that around one-quarter of these funds exhibit some significant exposure to market risk.

Keywords: diversification, alternative investments, correlation, portfolio decisions.

J.E.L. Codes: G23, G11.

---

\*Helpful comments and suggestions were received from Vikas Agarwal, Carol Alexander, Ron Anderson, Andrew Ang, Giovanni Belossi, Greg Connor, Oliver Linton, Andrew Lo, Sophocles Mavroidis, Narayan Naik, Burt Porter, Joshua Rosenberg, Chester Spatt, Bas Werker, Elliot Williams, and seminar participants at the 2004 European Finance Association meeting, Inquire UK, Aarhus School of Business, University of Amsterdam, Bank of England, LSE, NYU Stern School of Business and Tilberg University. The author would like to thank IAM for providing the data used in this study, and for financial support through the IAM Programme in Hedge Fund Research at LSE. Contact address: London School of Economics, Houghton Street, London WC2A 2AE, United Kingdom. Email: a.patton@lse.ac.uk. This paper is available from <http://fmg.lse.ac.uk/~patton/research.html>.

# 1 Introduction

The hedge fund industry is one of the fastest-growing sectors of the economy. The value of assets under the management of hedge funds has grown from \$50 billion in 1990 to around \$1 trillion in 2004<sup>1</sup>. In addition to the impressive returns observed on some hedge funds recently, the low correlation between hedge fund returns and market returns is an oft-cited favorable characteristic of hedge funds generally, see Brown, *et al.* (1999), Fung and Hsieh (2001) and Agarwal and Naik (2002). Indeed, the term ‘hedge fund’ was coined with reference to the goal of the first such funds, which was to invest in under-valued securities using the proceeds from short-sales of related securities, thereby creating a “market neutral” strategy, see Caldwell (1995).

Hedge funds are often classified according to their self-described investment strategies or styles, and the “equity market neutral” strategy is one of the largest such categories, representing about 20% of funds under the management of hedge funds according to Fung and Hsieh (1999). But despite their size, what exactly is meant by the moniker “market neutral” can be hard to pin down. Most definitions<sup>2</sup> of an equity market neutral strategy include phrases like “neutralize market risk(s) by combining long and short positions in related securities”, with limited detail on how neutrality should be measured and what risks should be considered “market” risks. Clarity and precision in the use of these terms would be beneficial. Indeed, the case of *Weyerhaeuser vs. Geewax Terker & Company*, which featured several prominent finance academics as expert witnesses, centered on whether Geewax Terker had truly followed a “market neutral” strategy<sup>3</sup>.

The most commonly used measure of neutrality is based on correlation or *beta*: a fund may be said to be “market neutral” if it generates returns that are uncorrelated with the returns on some market index, or a collection of market risk factors. Several studies, see Fung and Hsieh (2001), Mitchell and Pulvino (2001) and Agarwal and Naik (2002), have observed the nonlinear relation between hedge fund returns and market returns and proposed more sophisticated methods for modeling the expected returns on hedge funds: Fung and Hsieh (2001) suggest using payoffs from “lookback straddle” options on the market to approximate the pay-off structure of hedge funds.

---

<sup>1</sup>Source: The Economist, June 10, 2004.

<sup>2</sup>See Hedge Fund Research’s Strategy Definitions, Nicholas (2000) and Beliossi (2002) for example.

<sup>3</sup>Sources: The Wall Street Journal, March 25, 1992, and Pension & Investments, March 30, 1992. Geewax Terker eventually settled the lawsuit by agreeing to pay Weyerhaeuser \$8 million, almost as much as Weyerhaeuser had originally given Geewax Terker to manage.

Mitchell and Pulvino (2001) and Agarwal and Naik (2002) similarly suggest using piece-wise linear models for the hedge fund returns as a function of the market return.

We consider the concept of neutrality more generally than that implied by the use of correlations or betas. We suggest that one can consider the concept of market neutrality as having “breadth” and “depth”. The “breadth” of the neutrality of a hedge fund refers to the number of sources of “market” risk, such as equity market index risk, exchange rate risks and interest rate risks, etc., to which the returns on the hedge fund are neutral. The “depth” of the neutrality of a hedge fund refers to the “completeness” of the neutrality of the fund to market risks. We focus on market neutrality “depth” and propose five different neutrality concepts: “mean neutrality”, which nests the standard correlation- or beta-based definition of neutrality. “Variance neutrality”, “Value-at-Risk neutrality” and “tail neutrality” all relate to the neutrality of the risk of the hedge fund returns to market returns. The final concept, “complete neutrality”, corresponds to statistical independence of the fund and the market returns. We suggest statistical tests for each neutrality concept, and apply the tests to a combined database of monthly “market neutral” hedge fund returns from the HFR and TASS hedge fund databases, using the S&P 500 or the MSCI World indices to represent the market. Focusing solely on a single equity market index may be interpreted as testing a necessary condition for neutrality to a wider set of market variables.

By presenting a battery of neutrality concepts and tests we hope to aid investors’ evaluation of these funds, in a similar way to the use of the “Greeks” to evaluate the exposure of an option position, see Hull (2003). The concepts and tests proposed in this paper may be used as methods of analyzing the non-neutrality of hedge funds, rather than solely as strict tests of their neutrality. As the neutrality of a “market neutral” fund is one of its selling points, we conjecture that when comparing a collection of such funds, the risk, reward *and* the nature of the dependence between each fund and the market is of interest to investors. Of course, if the investor’s utility function was known, then funds could be directly ranked by expected utility, however such a case is not common in practise.

The “market neutral” class of hedge funds has not received a great deal of attention in the academic literature, though it currently represents a significant fraction of the hedge fund industry and has been growing at a rapid rate: from 2% of the hedge fund market in the early 1990’s to around 20% of the market in the late 1990’s, see Fung and Hsieh (1999) and Nicholas (2000). Mitchell and Pulvino (2001) focus on “risk arbitrage” hedge funds, Fung and Hsieh (2001) on “trend

following” hedge funds, and Agarwal and Naik (2002) on “event arbitrage”, “restructuring”, “event driven”, “relative value arbitrage”, “convertible arbitrage” and “equity hedge<sup>4</sup>” funds.

This paper makes two main contributions. First, we propose a number of different neutrality concepts relevant for hedge fund return and risk analysis, and present statistical tests of each neutrality concept. We pay particular attention to the types of non-neutral alternatives considered in each test: we consider either a general non-neutral alternative, or only those non-neutral alternatives that are disliked by risk averse investors with some existing exposure to market risk. For example, a risk averse investor prefers zero correlation to positive correlation, but prefers negative correlation to zero correlation. Thus we may test zero correlation against non-zero correlation, or only against positive correlation.

The second contribution of the paper is a detailed study of the neutrality of a combined database of “market neutral” hedge funds from the HFR and TASS hedge fund databases over the period April 1993 to April 2003. We use monthly data on 194 live and 23 dead “market neutral” hedge funds to evaluate their neutrality against a market index, the S&P 500. We find that approximately one-quarter of these funds exhibit some significant non-neutrality, at the 0.05 level. These proportions are lower than those found for other categories of hedge funds. Thus our findings suggest that many “market neutral” hedge funds are in fact *not* market neutral, but overall they are, at least, more market neutral than other categories of hedge funds. In a series of robustness checks we verify that our results are not overly affected by our choice of market index, our use of U.S. dollar returns, or by the last few, or first few, observations on fund returns.

The remainder of the paper is structured as follows. In Section 2 we describe the data used in this study. In Section 3 we present definitions of different types of neutrality, tests for each definition, and the results of these tests when applied to our collection of hedge funds. In Section 4 we present robustness checks of our results. Section 5 contains some discussion of extensions of the neutrality definitions presented in Section 3, and Section 6 concludes. An Appendix contains details on the bootstrap methods used in the paper.

---

<sup>4</sup>Although similar-sounding in name, the “equity hedge” index in the HFR database is distinct from the “equity market neutral” index. See [https://www.hedgefundresearch.com/pdf/HFR\\_Strategy\\_Definitions.pdf](https://www.hedgefundresearch.com/pdf/HFR_Strategy_Definitions.pdf). Nicholas (2000) includes some of these categories in the broad category of “market neutral”, and categorises the funds in our analysis as a subset of “equity market neutral” funds.

## 2 Description of the data and results using correlation

Our data set consists of those funds that categorize themselves as being “market neutral” and an equity market index, the S&P 500. We will focus on the S&P 500 as the market index for most of this paper, and show in Section 4 that our results do not change greatly if we instead use other equity market indices. The fund returns are monthly, net of management fees. Summary statistics on the funds are presented in Table 1, and summary statistics on the number of observations available on each of the funds are presented in Table 2. The latter of these two tables shows that we have between 59 and 213 “market neutral” funds available for analysis, depending on the data requirements of the test being considered.

[ INSERT TABLES 1 AND 2 ABOUT HERE ]

When computing measures of dependence between the market and a fund we do so using all data from the period when the fund was in the data base. The database includes both live and dead funds<sup>5</sup> and one may question whether the behavior of some funds in the period leading up to their dropping out of the database distorts our results. We show in Section 4 that this is not the case.

Before moving on to consider refinements of the definition of market neutrality, we will first analyze the relationship between the funds and the market index using standard linear correlation. The average correlation between the 171 hedge funds with 18 or more observations and the market index was 0.016, and the 5<sup>th</sup> and 95<sup>th</sup> sample quantiles of the cross-sectional distribution of correlation coefficients was [-0.64, 0.39] indicating substantial cross-sectional dispersion in the degree of correlation with the market portfolio.

Using a bootstrap procedure described in detail in the Appendix, which is designed to yield tests that are robust to serial correlation, volatility clustering and return non-normality, we find that 29.2% of the funds in our sample exhibit significant correlation with the market portfolio at the 0.05 level. This is statistic is surprisingly high: these are funds are (self-) described as “market neutral”, possibly to more factors than our single market index, and yet almost one-third of them have significant correlations with the market. If we instead focus our test only on deviations from

---

<sup>5</sup>As Agarwal, *et al.* (2003) point out, these funds are misnomered, since funds may drop out of the database for numerous reasons: liquidation (death), mergers, or simply a withdrawal from reporting to the database while continuing to operate.

zero correlation to *positive* correlation with the market, which is the sign of correlation a risk averse investor seeks to avoid, then we find 28.0% of funds with significant positive correlation. That is, over a quarter of the “market neutral” funds in our sample exhibit significant positive correlation with the market.

Under a joint null hypothesis that *all* funds in our sample are correlation neutral we would expect 5% of funds to be rejected. To determine whether the proportion of rejections we observe are significantly different from what we expect under the null hypothesis we consider two methods of obtaining a critical value. The first is based on the unrealistic assumption that the test statistics for each fund are independent. In this case the critical value can be obtained from the Binomial distribution, and it is 7.74% in this case. A bootstrap method for obtaining a critical value that allows for dependence between the test statistics is discussed in the Appendix, and yields a critical value of 18.75%, much higher than that obtained assuming independence, but still less than the observed proportion of rejections. Thus we conclude that we have significant evidence against “correlation neutrality” for this collection of funds as a whole.

Correlation neutrality is just one of many types of neutrality that may be of interest to a risk averse investor. An investor with quadratic utility, or one facing returns that are multivariate normally distributed, will only require linear correlation as the measure of dependence, and so this standard concept of market neutrality would suffice. However neither quadratic utility nor multivariate normality is an empirically reasonable assumption, particularly for hedge fund returns, and so we now consider alternative types of market neutrality.

### **3 Definitions and tests of versions of ‘market neutrality’**

In this section we consider refinements of the concept of market neutrality, using the preferences of a generic risk averse investor to motivate each concept and to determine the alternative hypotheses to consider. Depending on the preferences of the investor, one or more of the following definitions, or perhaps some “neutrality index” formed by some combination of these, may be of interest. We will start with the simplest generalization of correlation neutrality, and proceed through to the strictest form of neutrality: that of independence between the fund return and the market return.

### 3.1 Mean neutrality

The simplest neutrality concept, and the one that nests the standard “correlation neutral” concept, is that of “mean neutrality”. This is defined as the expected return on the fund being independent of the return on the market:

$$[it | mt] = [it] \forall mt \quad (1)$$

$$\text{or } [it | \mathcal{F}_{t-1}, mt] = [it | \mathcal{F}_{t-1}] \forall mt \quad (2)$$

and corresponds to the statement that the market return does not Granger-cause the fund return in mean. The second of the above conditions allows us to consider mean neutrality conditional on some other information,  $\mathcal{F}_{t-1}$ . We will focus primarily on the first of the above two conditions due to the data limitations we face.

Equation (1) is a stricter statement than correlation neutrality, as it rules out any function, not just a linear function, of the market return being useful for explaining fund returns. To test mean neutrality we could employ a number of methods. The most general would employ nonparametric regression to estimate  $\mu_i(mt) = [it | mt]$ :

$$it = \mu_i(mt) + \epsilon_{it}$$

and then test that  $\mu_i$  is equal to a constant. A simple alternative would be to employ a Taylor series approximation<sup>6,7</sup> to the conditional mean function:

$$it = \mu_0 + \mu_1 mt + \mu_2 \frac{mt^2}{2} + \dots + \epsilon_{it}$$

To capture the possibility that hedge funds exhibit different dependence on the market during market downturns than market upturns, as in Longin and Solnik (2001) or Ang and Chen (2002) for example, we consider a polynomial approximation that differs for positive versus negative market

---

<sup>6</sup>Numerous authors have, in various contexts, proposed using a polynomial in the market return to explain asset returns, see Bansal, *et al.* (1993), Chapman (1997), Harvey and Siddique (2000) and Dittmar (2002), amongst others.

<sup>7</sup>Mitchell and Pulvino (2001) and Agarwal and Naik (2002) both use piece-wise linear regressions rather than polynomials in their approximation of the conditional mean function. Under certain conditions on how the models expand as the sample size increases both methods can be considered nonparametric models for the conditional mean, see Andrews (1991) or Chen and Shen (1998). The use of a piece-wise linear specification with estimated kink points may lead to parameter identification problems when testing, however.

returns:

$$i_t = \alpha_0 + \alpha_1 m_t + \alpha_2 \frac{m_t^2}{|m_t|} + \dots \quad (3)$$

$$\text{where } i_t = \begin{cases} \alpha_1 m_t + \alpha_2 \frac{m_t^2}{|m_t|} + \dots & m_t \leq 0 \\ \alpha_0 + \alpha_1 m_t + \alpha_2 \frac{m_t^2}{|m_t|} + \dots & m_t > 0 \end{cases}$$

and then test

$$H_0 : \alpha_1 = \alpha_2 = 0 \text{ for all } m_t \quad (4)$$

$$\text{vs. } H_a : \alpha_1 \neq 0 \text{ or } \alpha_2 \neq 0 \text{ for at least one } m_t$$

via a standard  $F$  test. We estimated a simple second-order polynomial version of the model in equation (3) on the 150 funds with at least 24 observations, and found that for 23.3% of funds we could reject the null hypothesis of mean neutrality at the 0.05 level.

The definition above, however, ignores the fact that there are certain types of relations between the expected return on a fund and the market return that a risk averse investor would desire, and others that he/she would dislike. For example, a risk averse investor would prefer a negative relation between the fund and the market when the market return is negative, and a positive relation when the market return is positive, to zero correlation in both states. Thus it may not be mean neutrality that investors truly seek, or that “market neutral” hedge funds truly seek to provide, but rather a restricted type of dependence between the fund and the market. Below we derive a test of mean neutrality which tests only for violations of mean neutrality that are disliked by risk averse investors.

Consider the following refinement of mean neutrality, which we will call “mean neutrality on the downside”. This form of neutrality imposes that the expected return on the fund is neutral or negatively related to the market the market return is negative. That is:

$$\frac{\partial i(m_t)}{\partial m_t} \leq 0 \text{ for all } m_t \leq 0 \quad (5)$$

where  $i(m_t) \equiv E[i_t | m_t]$ . This version of neutrality ignores the relation between the fund and the market when the market return is positive, focusing solely on the ability of the fund to provide diversification benefits when the market return is negative. If we use the second-order polynomial version of equation (3) to approximate the conditional mean function, then

$$\frac{\partial i(m_t)}{\partial m_t} = \alpha_1 + 2\alpha_2 m_t \text{ for } m_t \leq 0$$



The first derivative of  $\beta_i$  is negative for all values of  $\beta_{mt} \leq 0$  if and only if  $\beta_1 \leq 0$  and  $\beta_2 \geq 0$ . Thus a test of mean neutrality on the downside may be obtained by testing the following hypothesis:

$$\begin{aligned} H_0 &: \beta_1 \leq 0 \cap \beta_2 \geq 0 \\ \text{vs. } H_a &: \beta_1 > 0 \cup \beta_2 < 0 \end{aligned} \tag{6}$$

As a test statistic we use  $\max[\beta_1 - \beta_2]$ , and we compare the observed value of this statistic with the 95<sup>th</sup> quantile of the bootstrap distribution of this test statistic. We were able to reject “mean neutrality on the downside” for only 0.7% of funds at the 0.05 level. Thus while we observe significant evidence against mean neutrality for these funds, we see no such evidence when restricting attention to alternatives that risk averse investors dislike<sup>8</sup>.

One concern about the above regressions and the subsequent tests on functions of the estimated parameters regards omitted variables bias. Consider an example where a single factor, different from the market portfolio, drives a given fund’s returns. If this factor is positively correlated with the market, then with sufficient data the above method will reject the null hypothesis that the fund is market neutral even though the factor driving returns is not the market portfolio. However, if the factor is correlated with the market then part of its risk is market risk and part is non-market risk, and so exposure to this factor does indeed involve some exposure to market risk. Thus the conclusion that this hypothetical fund is not market neutral seems reasonable. If one wanted to test market neutrality controlling for fund exposure to some other sources of risk, then this could be done by simply including the returns on the other factors in the above regression as control variables.

### 3.2 Variance neutrality

Another form of neutrality that one might expect from a “market neutral” fund is that the risk of the fund is neutral to market risk. In particular, we might expect that the risk of the fund, while not constant, does not increase at the same time as the risk of the market index. In this section

---

<sup>8</sup>In a previous version of this paper we used a second-order polynomial that did not allow the parameters to differ according to the sign of the market return. The proportion of funds that failed the general test of mean neutrality using this model was qualitatively similar to the figure reported here, but the proportion of funds that failed the test of mean neutrality on the downside was much higher: around 25%. This indicates the importance of allowing the dependence between funds and the market to depend on the sign of the market return.

we consider risk as measured by variance, and in the next section we consider risk as measured by Value-at-Risk. To our knowledge this paper is the first to consider the market neutrality of the risk of a hedge fund.

In the tests below we will control for mean non-neutrality before testing variance, VaR, or tail neutrality. This fact, combined with the limited data available, makes it unsurprising that we have only limited power in these tests. Nevertheless, these tests provide us with an alternative view on the relationship between a fund and the market and so may offer information not available in more standard mean-based tests.

Risk averse investors can be shown to have preferences over the dependence between the variance of the fund and the market return. Non-increasing absolute risk aversion, a property suggested by Arrow (1971) as being desirable in a utility function, leads to a preference for positive skewness in the distribution of portfolio returns. Kimball (1993) suggested further that reasonable utility functions should exhibit decreasing *absolute prudence*, which can be shown to imply an aversion to kurtosis in the distribution of portfolio returns<sup>9</sup>. Together these imply that risk averse investors prefer

$$\begin{aligned} \left[ (it - i)^2 \quad mt - m \right] &\geq 0, \text{ and} \\ \left[ (it - i)^2 \quad (mt - m)^2 \right] &\leq 0 \end{aligned}$$

so that the skewness of a portfolio of the fund and the market is larger and the kurtosis of the portfolio is smaller. With this motivation, we define “variance neutrality”, controlling for mean non-neutrality, as

$$[it - i(mt) | mt] = [it - i(mt)] \quad (7)$$

$$\text{or} \quad [it - i(mt) | \mathcal{F}_{t-1}, mt] = [it - i(mt) | \mathcal{F}_{t-1}] \quad (8)$$

In a similar manner to the previous section, we can obtain a test by approximating the true conditional variance function,  $\frac{\partial^2}{\partial i^2}(mt)$ , by a Taylor series polynomial:

---

<sup>9</sup>Related papers on investor preferences over higher-order moments include Kraus and Litzenberger (1976), Harvey and Siddique (2000) and Dittmar (2002), amongst many others.

$$i_t = i(mt) + i_t \quad (9)$$

$$i_t = i(mt) + i_t \sim (0, 1) \quad (10)$$

$$i(mt) = 0 + 1 \cdot mt + 2 \cdot \frac{1}{2} mt^2 + 1 \cdot mt(1-t) + 2 \cdot \frac{1}{2} mt(1-t) + i_t \quad (11)$$

$$\frac{\partial}{\partial i} i(mt) = 0 + 1 \cdot mt + 2 \cdot \frac{1}{2} mt + 3(1-t) \cdot mt + 4 \cdot \frac{1}{2} mt(1-t), \text{ or} \quad (12)$$

$$\frac{\partial}{\partial i} i(mt - i_{t-1}) = 0 + 1 \cdot mt + 2 \cdot \frac{1}{2} mt + 3(1-t) \cdot mt + 4 \cdot \frac{1}{2} mt(1-t) + 5 \cdot \frac{1}{2} i_{t-1} \quad (13)$$

$$\text{where } i_t = \begin{cases} 1 & mt \leq 0 \\ 0 & mt > 0 \end{cases} \quad (14)$$

where the latter conditional variance specification is designed to control for an ARCH(1) effect in the fund return. To test variance neutrality we would then test

$$H_0 : i = 0 \text{ for } i = 1, 2, 3, 4 \quad (15)$$

$$\text{vs. } H_a : i \neq 0 \text{ for some } i$$

We conducted this test on the 150 funds with more than 24 observations, with the ARCH(1) term as a control, and were able to reject the null at the 0.05 level for only 6.0% of funds. Thus most of these funds appear to be variance neutral to the market portfolio.

We can also consider “variance neutrality on the downside”, where we use the preferences of a risk averse investor to determine the desired sign of the first derivative of the conditional variance function when the market return is negative. The above preferences of a risk averse investor imply a preference for:

$$\frac{\partial}{\partial i} i(mt) \geq 0 \text{ for all } mt \leq 0 \quad (16)$$

Using the above specification for the conditional variance, we then obtain

$$\frac{\partial}{\partial i} i(mt) = 1 + 2 \cdot \frac{1}{2} mt \geq 0 \text{ for } mt \leq 0$$

and so the relevant hypotheses are:

$$H_0 : 1 \geq 0 \cap 2 \leq 0 \quad (17)$$

$$\text{vs. } H_a : 1 < 0 \cup 2 > 0$$

We use  $\max[-1, 2]$  as the test statistic, and compare the observed value of this statistic with the 95<sup>th</sup> quantile of its bootstrap distribution. We found significant violations of variance neutrality on the downside for 4.0% of funds when including an ARCH(1) term in the variance specification.

Overall, after controlling for violations of mean neutrality, we find no statistical evidence of violations of variance neutrality. This is in contrast with the widely-known fact that many hedge funds take “volatility bets”, that is, they take positions that pay off when market volatility is high, regardless of the direction of the movement. Thus it may be that our failure to find evidence against variance neutrality is due to a lack of data, and thus limited test power.

### 3.3 Value-at-Risk neutrality

The second risk-related neutrality concept we propose is that of “Value-at-Risk neutrality”, or “VaR neutrality”. Given that the VaR of an asset is simply a quantile of its distribution of returns<sup>10</sup>, this could also be called “quantile neutrality”. If we set the quantile to be 0.5, we would have a test of “median neutrality”, though we do not pursue that here. Quantile neutrality is a special case of complete neutrality, discussed below, which implies that *all* quantiles of the fund are neutral to the market, but differs from the previous two neutrality concepts in that it focuses on quantiles rather than moments. A VaR neutral portfolio is one with a VaR that is unaffected by the market portfolio return. That is:

$$P(r_{it} | mt) = P(r_{it}) \quad (18)$$

$$\text{or } P(r_{it} | mt, \mathcal{F}_{t-1}) = P(r_{it} | \mathcal{F}_{t-1}) \quad (19)$$

Violations of mean neutrality or variance neutrality will generally lead to violations of VaR neutrality, which leads us to consider “conditional VaR neutrality<sup>11</sup>”:

$$\left( \frac{r_{it} - VaR(r_{it})}{VaR(r_{it})} \middle| mt \right) = \left( \frac{r_{it} - VaR(r_{it})}{VaR(r_{it})} \right) \quad (20)$$

where we consider only the VaR of the standardized returns and not the returns themselves. Gupta and Liang (2003) have used VaR to examine the risk in hedge funds from a regulatory perspective.

---

<sup>10</sup>Formally, the Value-at-Risk of an asset is obtained from the following equality:  $\Pr[r_t \leq VaR_t | \mathcal{F}_{t-1}] = \alpha$ , where  $\alpha$  usually equals 0.10, 0.05 or 0.01.

<sup>11</sup>Note that we use the term “conditional VaR” to refer to the quantile of a conditional distribution. Other authors have used this term to describe the expected return conditioning on the VaR being breached, that is,  $E[r_{it} | r_{it} \leq VaR(r_{it})]$ , a quantity otherwise known as “expected shortfall” or “tail conditional loss”.

Though hedge fund returns are generally not normally or elliptically distributed, see Gupta and Liang (2003) for example, it is interesting to note that if the market and fund returns were jointly elliptically distributed then the portfolio VaR would be an affine function of the portfolio variance, and VaR neutrality would then follow directly from mean and variance neutrality. Under normality, conditional VaR neutrality would always hold, even if mean and variance neutrality did not, but for other elliptical distributions this need not be the case. Embrechts, *et al.* (2001) provide further discussion on VaR for portfolios, and see Artzner, *et al.* (1999) for a criticism of VaR as a measure of risk.

There are a number of ways that one might test the null hypothesis

$$\begin{aligned} 0 & : & (it | mt) &= & (it) \forall mt \\ \text{vs. } a & : & (it | mt) &\neq & (it) \text{ for some } mt \end{aligned} \quad (21)$$

or

$$\begin{aligned} 0 & : & \left( \frac{it - i(mt)}{i(mt)} \middle| mt \right) &= & \left( \frac{it - i(mt)}{i(mt)} \right) \forall mt \\ \text{vs. } a & : & \left( \frac{it - i(mt)}{i(mt)} \middle| mt \right) &\neq & \left( \frac{it - i(mt)}{i(mt)} \right) \text{ for some } mt \end{aligned} \quad (22)$$

With sufficient data one could use quantile regression, see Koenker and Bassett (1978), to test for the influence of the market return on a quantile of the fund return distribution in a similar way to our tests for mean and variance neutrality. However hedge fund return histories are notoriously short and the quantiles of interest are in the tail, so it is likely that data shortages will be a problem.

A simple alternative way of testing a necessary condition for VaR neutrality is via a test of Christoffersen (1998). This test examines whether the probability of one variable exceeding its VaR is affected by another variable exceeding or not exceeding its VaR. Specifically, we test:

$$\begin{aligned} 0 & : & \Pr \left[ it \leq \widehat{(\cdot)}(it) \middle| mt \leq \widehat{(\cdot)}(mt) \right] &= & \Pr \left[ it \leq \widehat{(\cdot)}(it) \middle| mt \leq \widehat{(\cdot)}(mt) \right] \\ \text{vs. } a & : & \Pr \left[ it \leq \widehat{(\cdot)}(it) \middle| mt \leq \widehat{(\cdot)}(mt) \right] &\neq & \Pr \left[ it \leq \widehat{(\cdot)}(it) \middle| mt \leq \widehat{(\cdot)}(mt) \right] \end{aligned} \quad (23)$$

where  $it \equiv (it - i(mt)) / i(mt)$  and  $mt \equiv (mt - m_t) / m_t$ . For the fund we again use a second-order polynomial for the conditional mean, and a second-order polynomial for the conditional variance with an ARCH(1) term. For the market we use a simple AR(1)-ARCH(1) model.  $\widehat{(\cdot)}(it)$  and  $\widehat{(\cdot)}(mt)$  are estimated by the empirical quantiles of  $it$  and  $mt$

Due to the data-intensive nature of studies of VaR, we only considered funds that had at least 66 months of observations available, which left us with 59 funds, and we tested the 10% VaR rather

than the more common 1% or 5% VaR. We conducted the conditional VaR neutrality test on these funds and found evidence against VaR neutrality for none of the funds at the 0.05 level. Thus we have no evidence against VaR neutrality for these funds, having controlled for mean and variance non-neutrality.

We can also consider a “downside” version of this test, which focusses specifically on testing whether the probability of the fund breaching its VaR is greater given that the market return has breached its VaR.

$$\begin{aligned} 0 & : \Pr \left[ it \leq \widehat{\phantom{it}} ( it ) \mid mt \leq \widehat{\phantom{mt}} ( mt ) \right] \leq \Pr \left[ it \leq \widehat{\phantom{it}} ( it ) \right] \\ \text{vs. } a & : \Pr \left[ it \leq \widehat{\phantom{it}} ( it ) \mid mt \leq \widehat{\phantom{mt}} ( mt ) \right] > \Pr \left[ it \leq \widehat{\phantom{it}} ( it ) \right] \end{aligned} \quad (24)$$

This version of VaR neutrality uses the fact that a risk averse investor would be averse to a fund that has a higher probability of a VaR exceedance given that the market has exceeded its VaR, and would have a preference for the opposite. Conducting this test on the funds we again find that none of these funds are rejected at the 0.05 level, and so conclude that we have no evidence against the downside VaR neutrality of these funds.

### 3.4 Tail neutrality

Now we consider the concept of neutrality during extreme events, or ‘tail neutrality’. Intuitively this can be thought of as an extension of VaR neutrality to the extreme tail: a market neutral fund should have a probability of extreme events that is unaffected by the market return. The formal definition of tail neutrality that we will use is:

$$L \equiv \lim_{q \rightarrow 0} \Pr [ i ( i ) \mid m ( m ) ] = \lim_{q \rightarrow 0} \Pr [ i ( i ) ] = 0 \quad (25)$$

where  $i|\Omega_{t-1} \sim i$  and  $m|\Omega_{t-1} \sim m$ . In words, our definition imposes that the probability of an extremely low return on the fund is not affected by conditioning on the fact that an extremely low return on the market is observed. The variable  $L$  is known as the coefficient of lower tail dependence, see Joe (1997) for example. If the fund return and the market return have zero lower tail dependence then the probability of an extreme negative return on the fund is unaffected by an extreme negative return on the market portfolio, and limits to zero as we consider more and more extreme returns. The alternative to tail neutrality is tail dependence, when  $L > 0$ . If the tail dependence coefficient is positive then there is a non-zero chance that both the fund and the

market will simultaneously experience an extremely low return. It is intuitively clear that risk averse investors would prefer tail neutrality to positive tail dependence: a higher probability of a joint crash increases the probability of a large negative return on a portfolio of these two assets. That is, positive lower tail dependence will generally lead to a fatter left tail.

A number of recent studies have proposed methods for detecting dependence in the tails of joint distributions. Longin and Solnik (2001) propose specifying a specific copula for the joint tails, Clayton’s copula in Nelsen (1999), and then testing that the parameter of this copula is such that no tail dependence is present. Bae, *et al.* (2003) model the probability of the joint occurrence of large returns across assets using parametric multinomial logistic regression. We employ the method of Quintos (2003), who proposes a nonparametric approach using extreme value theory, to derive a statistic to test for tail dependence<sup>12</sup>. Due to the heavy data requirements of tail analyses we restricted our sample to the 28 funds in our sample with at least 100 observations. Of these, 9 had enough observations in the joint tail to complete the test, and only one of these 9 funds rejected the null of no tail dependence at the 0.05 level<sup>13</sup>. Thus we conclude that no evidence of violations of tail neutrality is present for the funds in our database. This conclusion, however, may be overturned in the future when more data becomes available and our estimates of tail behavior become more precise.

It should be noted that the heavy data requirements of the VaR neutrality and tail neutrality tests introduce the possibility that survivorship bias affects our results. It may be that the funds that survive for a minimum of 66 or 100 months are those that live up to the name “market neutral”. This may be because surviving funds are those that have maintained a “good” return regardless of the market (which is a definition of market neutrality) or because investors desire “market neutral” funds that are truly market neutral and so these funds remain alive. In either of these scenarios, surviving funds would be more likely to pass VaR and tail neutrality tests, and thus the low proportion of rejections of VaR neutrality and tail neutrality would not be representative of the VaR and tail neutrality of “market neutral” funds with a shorter histories. We investigate fund longevity and market neutrality further in Section 4.

---

<sup>12</sup>We also implemented the method of Longin and Solnik (2001) on our hedge funds. While this method is quite different in implementation from the method of Quintos (2003), we drew similar conclusions using both methods.

<sup>13</sup>We used the asymptotic theory provided by Quintos (2003) rather than the bootstrap for this test.

### 3.5 Complete neutrality

“Complete neutrality” is the strictest form of neutrality, and requires that the distribution of fund returns is completely independent of the market return. The formal definition is:

$$f_{i|m} = f_i \quad (26)$$

where “ $=^d$ ” indicates equality in distribution. If we let  $f_{it} \sim f_i$  and  $f_{mt} \sim f_m$ , and  $(f_{it}, f_{mt}) \sim f_{i,m}$ , this implies that

$$f_{(it, mt)} = f_i(f_{it}) \cdot f_m(f_{mt}) \quad (27)$$

whereas in general the joint distribution of the fund return and the market return is written as

$$f_{(it, mt)} = f_i(f_{it}) \cdot f_m(f_{mt}) \cdot c(f_i(f_{it}), f_m(f_{mt})) \quad (28)$$

where  $c$  is the “copula density”<sup>14</sup> or “dependence function” of the fund and the market returns. Under complete neutrality the assets’ copula is the “independence copula”, denoted  $C_I$ , which takes the value 1 everywhere. We can use the preferences of a risk averse investor to derive a ranking of copulas between the fund and the market using a result of Epstein and Tanny (1980).

A general alternative to complete neutrality is a dependence function,  $C^*$ , that differs from  $C_I$  by a “correlation-increasing transformation” (CIT) of Epstein and Tanny (1980). A CIT is the dependence equivalent of the better-known “mean-preserving spread” of Rothschild and Stiglitz (1970). A CIT shifts some probability mass towards realizations where both variables are “large” or “small” and away from realizations where one is “large” and the other is “small” in such a way that the marginal distributions of the variables are preserved. From Epstein and Tanny (1980) we know that:

$$C_I(x, y) \leq C^*(x, y) \quad \forall (x, y) \in [0, 1] \times [0, 1] \quad (29)$$

and we say that  $C^*$  is “more concordant”<sup>15</sup> than  $C_I$ , or simply that  $C_I \leq C^*$ . This ordering is a multivariate first-order stochastic dominance ordering.

<sup>14</sup>The copula *cdf* is denoted with an upper case  $C$  while the copula density is denoted with a lower case  $c$ . See Nelsen (1999) for an introduction to copulas.

<sup>15</sup>Epstein and Tanny (1980) interpret the condition in equation (29) as saying that  $C^*$  exhibits “greater correlation” than  $C_I$  but we will refrain from using the term “correlation” unless referring directly to Pearson’s linear correlation or Spearman’s rank correlation.



To relate the above statistical ordering of dependence functions to some economic ordering, Epstein and Tanny (1980) introduce the concept of “correlation aversion”. A utility function involving two random variables is “correlation averse” if expected utility is reduced by a CIT. This can be checked directly for utility functions that are twice differentiable by checking whether

$$\frac{\partial^2 U(x, y)}{\partial x \partial y} \leq 0 \quad (30)$$

In this paper we consider the case that:

$$\begin{aligned} U(x, y) &= \mathcal{U}(w_1 x + w_2 y), \text{ where } \mathcal{U} \text{ is some utility function} \\ \frac{\partial^2 U(x, y)}{\partial x \partial y} &= \mathcal{U}''(w_1 x + w_2 y) w_1 w_2 \\ &\leq 0 \text{ if } w_1 w_2 \geq 0 \end{aligned}$$

for any concave utility function  $\mathcal{U}$ . We rule out short selling either the hedge fund or the market portfolio, both reasonable restrictions for most investors, and so  $w_1 w_2 \geq 0$ . The weak inequality above holds strictly if  $w_1 w_2 > 0$ ; of course if either portfolio weight is zero then the investor is “correlation neutral”. Thus any risk averse investor subject to short selling constraints will be (weakly) “correlation averse”.

The concept of correlation aversion can be used to derive an economic ranking of dependence functions for risk averse investors: if  $F_1$  and  $F_2$  are two possible joint distribution functions for  $(x, y)$  with common marginal distributions, and if

$$F_1[(x, y)] \leq F_2[(x, y)] \quad (31)$$

for all correlation averse utility functions  $U$ , then Epstein and Tanny (1980) write  $F_2 \preceq_u F_1$  and say that  $F_1$  exhibits greater correlation than  $F_2$ . The main theorem in Epstein and Tanny (1980) shows that the ranking obtained from the expected utility of risk averse investors is equivalent to the purely statistical concordance ordering discussed above. That is,

$$F_2 \preceq_u F_1 \Leftrightarrow F_2 \leq F_1 \quad (32)$$

In terms of dependence functions and neutrality, this implies that

$$I \preceq_u I^* \Leftrightarrow I \leq I^* \quad (33)$$

and so any dependence function that is a CIT away from complete neutrality will be less preferred by risk averse investors. Epstein and Tanny (1980) thus show theoretically that general risk averse

investors care about the dependence (not just the correlation) between hedge fund returns and market returns.

We could use the above results to motivate tests for a concordance ordering of hedge funds, using tests for multivariate first-order stochastic dominance. Instead we propose the more modest task of examining the ordering of a scalar measure of dependence, namely Spearman’s rank correlation. Nelsen (1999) shows that Spearman’s rank correlation will reflect the concordance ordering of two dependence functions. That is,

$$F \leq G \Rightarrow S(F) \leq S(G) \text{ for any copulas } F, G \quad (34)$$

and so we have

$$F \preceq_u G \iff F \leq G \iff S(F) \leq 0 \leq S(G) \quad (35)$$

Thus we may obtain an approximate ordering of the funds for a general risk averse investor by categorizing the funds as having significant negative rank correlation, non-significant rank correlation or significant positive rank correlation with the market index. Rank correlation can detect monotonic nonlinear relationships, in addition to the linear relationships that the usual correlation coefficient may be used to detect.

Average rank correlation across the 171 funds with at least 18 observations was 0.016; a similar figure to that obtained using linear correlation. From tests for non-zero rank correlation we found 26.3% of funds had significant rank correlation at the 0.05 level, and 25.1% of funds had significantly positive rank correlation at the 0.05 level. Of course, complete neutrality implies neutrality of any other type, and so all other tests in this paper may also be thought of as tests of necessary conditions for complete neutrality.

### 3.6 Summary: are ‘market neutral’ hedge funds really market neutral?

In this section we combine the results of the tests introduced above to draw an overall conclusion about the market neutrality of funds with the label “market neutral”. Given that so few of the funds in our sample had sufficient data for the test of tail neutrality to be applied, we will not consider this test when drawing overall conclusions.

Declaring a fund to be non-neutral if it fails at least one test for market neutrality leads to a size distortion. For example, the probability that a truly market neutral fund fails at least one

of five independent tests of market neutrality, each with size 0.05, is 0.23. Further, we must take into account the fact that the test statistics for each of the five tests considered here (correlation neutrality, mean neutrality, variance neutrality, VaR neutrality and complete neutrality) are not likely to be independent. We deal with these two problems by looking at the number of tests failed by a set of bootstrapped data series, generated imposing the independence of the fund and the market returns. If the actual number of tests failed is greater than the 95<sup>th</sup> percentile of the number of tests failed by the bootstrapped data sets, then we conclude that the fund fails a joint test of market neutrality. Further details are in the Appendix. The results obtained for “market neutral” hedge funds are collected in the first column of Table 4.

At the 0.05 level, we found that 28.1% of funds failed a joint test of market neutrality against general non-neutral alternatives, while 21.6% of funds failed a joint test of market neutrality against alternatives that are disliked by risk averse investors. The 95% critical value for the proportion of funds failing a test of neutrality is 18.75%, and so both of these proportions represent significant evidence against market neutrality for these funds as a whole.

A natural question to ask is how the neutral and non-neutral funds differ along various dimensions. In Table 3 we compare some simple summary statistics on two portfolios of these funds: the first (second) portfolio is constructed as an equally-weighted average of the neutral (non-neutral) funds for each month in our sample. Table 3 reveals that the non-neutral portfolio yielded a significantly greater average return than the neutral portfolio, and had significantly greater standard deviation. The correlation with the market was significantly greater for the non-neutral portfolio (a predictable outcome, given the way the two portfolios were constructed) and the average age of non-neutral funds was significantly greater than neutral funds. The fact that non-neutral funds are, on average, almost twice as old as neutral funds may either reflect the fact that more observations lead to greater power to reject the null hypothesis of neutrality if it is false, or it may genuinely reflect the fact that older funds tend to stray from market neutrality more than younger funds. There was no significant difference in the average size of neutral and non-neutral funds.

[ INSERT TABLE 3 ABOUT HERE ]

In summary, we conclude that approximately one quarter of “market neutral” funds in our sample exhibit significant deviations from market neutrality, of a type that is specifically disliked by risk averse investors. Our sample sizes are not extremely large (the median sample size is just

42 observations) which means that the power of the tests employed may be low, suggesting that the true proportion of non-neutral funds may be even higher. Our findings suggest that careful analysis of fund returns is required to reap the widely-cited diversification benefits of hedge funds.

### **3.7 Are ‘market neutral’ hedge funds more market neutral than other funds?**

In this section we apply the tests introduced above to collections of hedge funds with other styles. We look at four other hedge fund styles from the HFR data base: equity hedge, equity non-hedge, event driven, and funds of funds. “Equity hedge” funds hold some exposure to the market, with the degree of exposure ranging from near zero to over 100%, along with some hedge, either through short sales of stocks or through stock options. “Equity non-hedge” funds are otherwise known as “stock pickers”. These funds may also hedge their exposures, though generally not consistently. “Event driven” funds seek returns from mergers, takeovers, bankruptcies, etc. These funds may or may not hedge their exposures to the market. Funds of hedge funds invest in multiple funds, which may or may not be in the same category.

Table 4 reports the results of tests for different versions of neutrality on these four categories of hedge funds, as well as on the market neutral hedge funds discussed above. This table shows that a far higher proportion of funds in the equity hedge, equity non-hedge, event driven, and funds of hedge funds categories exhibit significant exposure to equity market risk. Over 80% of equity non-hedge funds, for example, exhibit some significant violation of market neutrality at the 0.05 level, and over 88% of these funds have a significantly positive correlation coefficient with the market. The average correlation coefficient across the 77 funds in this category with at least 18 observations is 0.51. The funds of funds category is the most market neutral of these four categories, with about half of these funds exhibiting some significant violation of market neutrality. The average correlation coefficient across the 457 funds in this category with at least 18 observations was 0.25.

Recalling that only 25% of “market neutral” funds exhibited significant violations of market neutrality, and that the average correlation coefficient across funds was 0.016, we draw the conclusion that while not all “market neutral” funds are truly market neutral, they are, as a category, substantially more market neutral than other hedge fund categories.

[ INSERT TABLE 4 ABOUT HERE ]

## 4 Robustness checks

In this section we conduct robustness checks of the results reported above. Firstly, we consider an alternative index for the “market” portfolio, the MSCI World index. We then consider changing the outlook of our hypothetical investor from one who cares about U.S. dollar returns to one who cares about British pound returns. We then analyze whether our results change when we drop the last six months, or the first 12 months, of available data on firms, and finally we look at the relation between the number of observations available on a fund and its dependence characteristics.

**Choice of market portfolio.** Obviously the choice of market index is an important input to tests of market neutrality. In the paper we considered using the S&P 500 index as the market index, and a summary of results for this case are presented in the first column of Table 5. Corresponding results when the MSCI World index is instead used are presented in the second column of Table 5. Comparing these two columns shows that our results are robust to this choice. We also obtained results (not reported) when the MSCI Europe index was employed and again no substantial differences were found.

[ INSERT TABLE 5 ABOUT HERE ]

**Choice of currency.** To consider the impact of our choice to examine the neutrality of these hedge funds from the perspective of a U.S. investor, we re-computed all results from the perspective of a U.K. investor. The results for a U.K. investor using the MSCI World index are presented in the third column of Table 5. The differences in the results are small: the proportion of rejected funds falls from roughly one-quarter to one-fifth, but remains significant.

**End-game behavior.** In the months leading up to a fund dropping out of the HFR or TASS databases it is conceivable that the behavior of a fund’s returns changes. If a fund is doing poorly and is about to be liquidated then the investment decisions of the hedge fund manager may place greater emphasis on objectives other than maintaining the market neutrality of the fund. For this reason, we re-computed all the results for the U.S. based investor using the S&P 500 index as the market index, dropping the last six observations on each fund. The results are presented in the fourth column of Table 5, and are not substantially different from the results for the benchmark case.

**Backfill bias.** Hedge funds usually enter databases with a history of returns, usually ranging from 6 to 18 months in length. It might be reasonable to think that the emphasis placed on maintaining market neutrality during the first year of a fund’s life, relative to simply achieving positive returns, is lower than for the rest of a “market neutral” fund’s life. Thus a rejection of market neutrality may come from the fund’s first year of life, and not be representative of its neutrality following that first year. To allow for this we re-computed all the results for the U.S. based investor using the S&P 500 index as the market index, dropping the first twelve observations on each fund. The results are presented in the last column of Table 5 and are not substantially different from the results for the benchmark case.

**Age of the fund and its market neutrality.** Above we reported proportions of rejections of market neutrality concepts, averaging across all funds with sufficient data to conduct the test. But an interesting, and possibly important, question is whether the older funds have different market dependence properties to newer funds. As an example, in Figure 1 we plot the linear correlation between a fund and the S&P 500 market index against the number of observations available on that fund. This plot indicates a significant positive relation between the correlation coefficient and the age of the fund<sup>16</sup>. The robust t-statistics associated with each of these correlation coefficients also have a positive relation with the number of observations available. Further, a probit regression (not reported) of the probability of a t-statistic being greater than 1.96 revealed a positive and significant dependence on the number of observations available. A similar picture emerged when comparing the average age of funds that passed and failed the joint test of market neutrality on the downside, presented in Table 3.

These findings suggest that market neutral hedge funds that survive for a relatively long time are more positively dependent on the market return than younger funds. This may be related to the fact that many hedge fund prospectuses commit the fund to a “market neutral” strategy only for a fixed period of time (three years seems to be a common choice), if at all. “Material” changes of investment strategy within that time often require the approval of shareholders, but changes after that time are not discussed. It is thus possible that hedge funds exhibit “style drift”, and that funds that were once correctly classified as “market neutral” may not follow such strategies

---

<sup>16</sup>The reader may notice two observations in the upper right-hand corner of this plot, representing two funds that have the maximum number of observations (121) and correlation coefficients of over 0.9. We re-did the regression without these observations and the relation was still significantly positive.

any more. This observation suggests that a classification method based on relating fund returns to various risk factors, as in Brown and Goetzmann (1997) and Fung and Hsieh (1997) for example, may be more reliable than one based on the self-reported styles of funds.

[ INSERT FIGURE 1 ABOUT HERE ]

## 5 Extensions

In this section we discuss the issues surrounding two of the more important extensions of the definitions and tests presented in Section 3. A full treatment of these and other extensions is being pursued in separate work.

### 5.1 Time-varying market neutrality

That hedge fund managers employ dynamic trading strategies is widely known, and thus one may ask whether conducting tests on unconditional correlation coefficients or market betas, for example, is less interesting than studying the conditional correlations or betas between the return on a fund and the return on the market index. In this section we propose extending (unconditional) beta (or correlation) neutrality to “conditional beta neutrality”. We leave the consideration of time variation in the other forms of neutrality proposed in this paper for future work.

Consider the following simple time-varying “beta” model, in the spirit of Harvey (1989) and Shanken (1990):

$$\begin{aligned}
 r_{it} &= \alpha_i + \beta_{it} r_{mt} + \epsilon_{it} \\
 \beta_{it} &= \beta_0 + \beta_1 \bar{r}_{it,6} + \beta_2 \bar{r}_{mt,6} + \beta_3 \bar{\sigma}_{mt,6}^2
 \end{aligned} \tag{36}$$

where  $\bar{r}_{it,6} \equiv \sum_{j=1}^6 r_{i,t-j}$ ,  $\bar{r}_{mt,6} \equiv \sum_{j=1}^6 r_{m,t-j}$ ,  $\bar{\sigma}_{mt,6}^2 \equiv \sum_{j=1}^6 \sigma_{m,t-j}^2$ . This model is designed to capture the possibility that “market neutral” hedge fund managers may dynamically adjust their exposure to the market according to the past performance of their fund, the past performance of the market, and the recent volatility of the market. We did not conduct an extensive search of the possible determinants of the conditional betas of these hedge funds; these three variables seemed reasonable variables to use in this initial investigation. Correlation neutrality, as defined in Section 2, is tested by estimating  $\beta_0$ , under the restriction that  $\beta_1 = \beta_2 = \beta_3 = 0$ , and then testing

that  $\beta_0 = 0$ . But it is conceivable that a fund may have a zero beta on average, but a non-zero



sufficient condition for market neutrality to the wider array of market risks. Greater power may be obtained by constructing tests of neutrality to multiple sources of market risk.

All of the tests proposed in this paper generalize quite straightforwardly to consider multiple sources of market risk. However the limited amount of hedge fund returns data available may mean that the most obvious extensions, where the additional market risks are simply added as additional explanatory variables, will actually lead to a loss of power rather than a gain; increased estimation error may dominate the increased information from the additional risk sources. One alternative is to combine the sources of risk into a single index or portfolio, and then test for neutrality against this portfolio of market risks. For example, to allow both contemporaneous and lagged market returns to affect the hedge fund return, we could use the following simple index:

$$\tilde{m}_t \equiv 0.4 m_t + 0.3 m_{t-1} + 0.2 m_{t-2} + 0.1 m_{t-3} \quad (38)$$

Using  $\tilde{m}_t$  rather than  $m_t$  as the market index allows lagged returns to affect hedge fund risk, as suggested by Asness, *et al.* (2001) and Getmansky, *et al.* (2003), without increasing the number of parameters that need to be estimated. The above choice of weights is arbitrary, but reflects the spirit of these two papers, and is roughly in accordance with the median values, across funds, of the regression coefficients obtained from regressing the fund return on a constant, and lags zero through three of the market return<sup>17</sup>. When we applied the tests of the previous sections using this weighted average of lagged market returns as the market index we found results quite similar to those obtained just using the contemporaneous value of the market return, see Table 6. This suggests that lagged market returns *do* impact these hedge fund returns: if lagged market returns were unrelated to hedge fund returns then the  $\tilde{m}_t$  variable would simply be a noisier measure of market risk, and the proportion of rejections would fall. It may be possible to construct alternative combinations of current and lagged market returns that lead to a greater proportion of rejections, keeping in mind the strict trade-off between flexibility and parsimony, but we leave this for future work.

[ INSERT TABLE 6 ABOUT HERE ]

---

<sup>17</sup>The median values of these four coefficients, across funds, were 0.51, 0.20, 0.14, 0.15, when the coefficients were normalised to sum to one.

## 6 Concluding remarks

“Market neutral” hedge funds manage about 20% of the \$1 trillion currently under the management of hedge funds. One of the attractions of a “market neutral” fund is a low degree of dependence between the fund and the market. We considered generalizing the concept of “market neutrality” to reflect both “breadth” and “depth”. The “breadth” of neutrality of a fund reflects the *number* of market-type risks to which the fund is neutral. The “depth” of neutrality of a fund reflects the *completeness* of the fund’s neutrality to market risks. We proposed five new neutrality concepts for hedge funds: “mean neutrality”, which nests the standard correlation-based definition of neutrality; “variance neutrality”, “Value-at-Risk neutrality”, and “tail neutrality”, which examine the neutrality of the risk of a fund to market risk; and “complete neutrality” which corresponds to independence of the fund and the market returns.

We proposed statistical tests of each of these neutrality concepts. These tests take neutrality as the null hypothesis and compare it with either a general non-neutral alternative hypothesis, or

proposed in this paper may help investors to learn more about the relationship between a given hedge fund and the market index.

The work in this paper leaves unanswered many interesting questions. It is well-known that hedge fund managers employ dynamic trading strategies and so it would be interesting to extend the tests in this paper to allow for time-varying neutrality. Further, allowing for multiple sources of “market” risk, rather than just a single source as considered in this paper, is an important direction for future work. We discussed some of the issues surrounding these two extensions but we have by no means solved those problems completely. Another open problem relates to forming portfolios of market neutral hedge funds that attain some desired degree of neutrality to the market. Funds of hedge funds are a fast-growing sector of the hedge fund industry and currently account for about 30% of invested funds. Clearly the problem of combining “market neutral” funds to form a market neutral (no quotation marks) fund would be of interest to this sector.

## 7 Appendix: Details of the bootstrap tests

Concerns about short sample sizes, serial correlation, volatility clustering, and non-normality of asset returns prompted the use of a block bootstrap to obtain critical values for the various tests proposed in this paper. We used the stationary bootstrap of Politis and Romano (1994), and the algorithm of Politis and White (2004) to determine the optimal average block size for each asset. Specifically, for the fund and the market, we applied the Politis and White algorithm to the return, squared return, and the product of the fund return and the market return. We then selected the largest of these three lengths to use as the block length for that asset. The block lengths selected ranged from 1 to 10, and averaged 2.4. We used 1000 bootstrap replications.

To obtain the distribution of each test statistic under the null hypothesis of neutrality we re-sampled the fund and market returns separately, rather than re-sampling the vector of fund and market returns. By using the stationary bootstrap and imposing independence between the bootstrapped fund and market data we ensure that the null hypothesis in each of the tests is satisfied, while not changing the univariate distributions of the fund and market returns, at least asymptotically. For all but the test of complete neutrality, independence is a sufficient but not necessary condition for the null hypothesis to hold.

**Obtaining a joint test.** Using five individual tests (correlation neutrality, mean neutrality, variance neutrality, VaR neutrality and complete neutrality) and obtaining a “joint test” by simply checking whether at least one test was failed clearly leads to a size distortion. Instead we employed a method related to that of Westfall and Young (1993): on each bootstrap sample we conduct the five tests, we then count the number of tests that lead to a rejection of a null hypothesis. If the test statistics were independent then we would expect  $0.05 \times 5 = 0.25$  tests to be failed for each sample, however these test statistics are almost certainly not independent, and by using this bootstrap procedure we capture this dependence. We then compute the 95<sup>th</sup> percentile of the distribution of the number of tests failed. Across all funds this quantile ranged from 0 to 2, with a median of 1 and a mean of 0.8, though for each fund the quantile is of course an integer between 0 and 5. (If the tests were independent then the 95<sup>th</sup> percentile would be 1: the *cdf* of a Binomial(5 0.05) evaluated at 1 is 0.98 ) We conclude that a fund failed the *joint* test of market neutrality if it failed more tests than the 95<sup>th</sup> percentile of the distribution of the number of tests failed by the bootstrap data. This test controls the size of the joint test, and enables us to draw an overall conclusion about

the neutrality of a fund.

**Obtaining critical values on the proportion of funds failing a test.** Under the null hypothesis that all funds are neutral the expected proportion of funds to fail a given test is 0.05, but in order to conclude that the proportion of funds that fail a given test is significantly more than we would expect under the null we must obtain the 95<sup>th</sup> percentile of the distribution of the proportion of funds that fail a given test under the null hypothesis. We consider two approaches. The first approach relies on the assumption that the test statistics are independent across funds, and so the number of funds that fail a given test is a  $\text{Binomial}(n, 0.05)$  random variable, where  $n$  is the number of funds that are tested. The 95<sup>th</sup> percentile of this distribution, divided by  $n$ , ranged from 0.077 to 0.102 depending on the test (which affected the value for  $n$ ).

The assumption that the test statistics are independent across funds is not realistic, and we considered a second approach again using the bootstrap to estimate the joint distribution of these test statistics under the null. Ideally, we would employ a vector bootstrap using all  $n$  funds to obtain this distribution. However because some of the funds have missing observations, either because they entered the sample late or left the sample early, we cannot use this method: it would generate bootstrap samples that have missing observations in the middle of the sample, which would mean that we could not employ all of our tests (the estimation of a GARCH model, for example, requires a contiguous sample). We employed an alternative approach, using a vector bootstrap only for the  $n^*$  funds that are present for the full sample<sup>18</sup>. For each bootstrap sample we compute the proportion of these funds that failed each test, and then use the 95<sup>th</sup> percentile of the bootstrap distribution of these proportions as the critical value. This method relies on the assumption that the proportion of the  $n^*$  funds that fail each test has the same distribution as the proportion of all  $n$  that fail each test, which may or may not be a reasonable assumption, but is likely a more reasonable assumption than that of independence between all test statistics. For “market neutral” hedge funds this method yielded a critical proportion of 0.125 for the VaR-neutrality test, and 0.188 for the remaining tests.

---

<sup>18</sup>The number of funds that were present for the full sample was 16, 51, 20, 14 and 49 for the five hedge fund styles considered here: market neutral, equity hedge, equity non-hedge, event driven, and funds of funds.

**Table 1**  
**Summary statistics of fund and market index returns**

This table presents descriptive statistics on the monthly fund and market index returns over the sample period, April 1993 to April 2003. The column headed “median fund” presents the median of the statistic in the row across the 213 funds with more than 6 observations. The “Jarque-Bera statistic” refers to the Jarque-Bera (1980) test of normality, the p-value for this test is also reported. †The mean and standard deviation statistics have been annualized to ease interpretation. The annualized standard deviation was computed using a Newey-West estimate of the variance, to take into account serial correlation in returns.

	Median fund	S&P 500	MSCI World
Mean <sup>†</sup>	8.11	10.16	6.63
Standard deviation <sup>†</sup>	8.01	15.22	14.73
Skewness	0.02	-0.54	-0.5425
Kurtosis	3.65	3.26	3.28
Minimum	-4.54	-14.44	-13.35
Maximum	5.87	9.78	9.02
Jarque-Bera statistic	2.56	5.99	6.04
Jarque-Bera p-value	0.28	0.05	0.05
Number of obs	42	121	121

**Table 2**  
**Summary statistics on the number of observations available**

This table presents descriptive statistics on the number of monthly return observations available on each of the five categories of hedge funds used in this paper. The sample period is April 1993 to April 2003 and so the maximum number of observations is 121. The bottom four rows present the number of funds available for analysis in the tests of the various types of market neutrality.

	Market Neutral	Equity Hedge	Equity Non-hedge	Event Driven	Funds of Funds
Minimum	1	1	1	6	1
0.25 quantile	19	23	40	28	18
Median	42	44	79	58	40
Mean	49.5	52.9	74.1	61.6	48.6
0.75 quantile	69	77	114	90	73
Maximum	121	121	121	121	121
Number of funds with $\geq 6$ obs	213	543	84	110	569
Number of funds with $\geq 18$ obs	171	466	77	90	463
Number of funds with $\geq 24$ obs	150	422	77	86	414
Number of funds with $\geq 66$ obs	59	182	58	47	169

**Table 3**  
**Comparison of neutral and non-neutral hedge funds**

This table presents descriptive statistics on two portfolios constructed using the results of the joint test of neutrality on the downside, reported in Section 3.6. The first (second) portfolio is constructed as an equally-weighted average of the neutral (non-neutral) funds for each month in our sample period, April 1993 to April 2003. 21.6% of the 171 funds studied failed the joint test of neutrality on the downside, and so the “neutral” portfolio contains up to 134 funds, while the “non-neutral” portfolio contains up to 37 funds. The fact that some funds dropped out of the database and others joined sometime after April 1993 implies that the two portfolios contained varying numbers of funds. The bottom two rows report the median age and size (in millions of U.S. dollars, as at the last available observation for each fund) of the funds that comprise the two portfolios. The third column reports the difference between the “neutral” and “non-neutral” portfolios, and an asterisk denotes a difference that is significantly different from zero at the 0.05 level. †The mean and standard deviation statistics have been annualized to ease interpretation. The annualized standard deviation was computed using a Newey-West estimate of the variance, to take into account serial correlation in returns.

	Neutral fund portfolio	Non-neutral fund portfolio	Difference
Mean <sup>†</sup>	3.15	6.49	-3.34*
Standard deviation <sup>†</sup>	0.79	1.48	-0.69*
Skewness	0.20	-0.33	0.53
Kurtosis	3.19	3.14	0.04
Correlation with market	-0.06	0.78	-0.85*
Median size	\$23m	\$15m	\$8m
Median age	46 months	88 months	-42 months*

**Table 4**  
**Proportion of funds that fail tests of neutrality, by hedge fund style**

This table presents the results of tests of various types of market neutrality presented in the paper, applied to five different hedge fund styles. These styles are described in Section 3.7. The tests are conducted on monthly hedge fund returns over the period April 1993 to April 2003. The S&P 500 is used as the market index. Panel I reports the proportion of funds for which the null hypothesis of market neutrality could be rejected at the 0.05 level in favour of a general alternative. Panel II reports the proportion of funds for which the null hypothesis of market neutrality could be rejected at the 0.05 level in favour of a dependence structure of a type that is disliked by U.S.-based risk averse investors with a pre-existing exposure to the S&P 500. An asterisk indicates that the proportion of rejected funds is significantly different from zero (using the bootstrap procedure described in the appendix) and thus that collection of funds taken as a whole are significantly non-neutral.

	Hedge fund style				
	Market neutral	Equity hedge	Equity non-hedge	Event driven	Funds of funds
Panel I: Proportion rejected null of neutrality					
Correlation neutrality	0.298*	0.562*	0.870*	0.667*	0.497*
Mean neutrality	0.233	0.481*	0.766*	0.640*	0.531*
Variance neutrality	0.060	0.024	0.000	0.012	0.031
VaR neutrality	0.000	0.033	0.052	0.064	0.024
Complete neutrality	0.263*	0.584*	0.883*	0.589*	0.492*
Joint test	0.281*	0.554*	0.870*	0.644*	0.503*
Panel II: Proportion rejected null of neutrality on the downside					
Correlation neutrality	0.269*	0.573*	0.883*	0.744*	0.534*
Mean neutrality	0.007	0.012	0.065	0.012	0.007
Variance neutrality	0.040	0.000	0.000	0.035	0.000
VaR neutrality	0.000	0.033	0.052	0.064	0.024
Tail neutrality	0.039	0.005	0.417*	0.077	0.029
Complete neutrality	0.251*	0.577*	0.883*	0.711*	0.523*
Joint test	0.216*	0.537*	0.805*	0.644*	0.486*



**Table 5**  
**Results from robustness checks**

This table presents the results of tests of various types of market neutrality presented in the paper, applied to variations on the original data set to check the robustness of the conclusions in various dimensions. The tests are conducted on monthly hedge fund returns over the period April 1993 to April 2003. The S&P 500 is used as the market index. Panel I reports the proportion of funds for which the null hypothesis of market neutrality could be rejected at the 0.05 level in favour of a general alternative. Panel II reports the proportion of funds for which the null hypothesis of market neutrality could be rejected at the 0.05 level in favour of a dependence structure of a type that is specifically disliked by U.S.-based risk averse investors with a pre-existing exposure to the S&P 500. An asterisk indicates that the proportion of rejected funds is significantly different from zero (using the bootstrap procedure described in the appendix) and thus that the collection of funds taken as a whole are significantly non-neutral. The first column of results are for the benchmark case, and are the same as those reported in the first column of Table 4. The second column considers the impact of switching the market index from the S&P 500 to the MSCI World index. The third column considers the MSCI World index as the market index, from the perspective of a U.K. based investor, and so all US dollar returns are converted to British pound returns. The fourth and fifth columns check the robustness of the results to “end-game” behavior and to “backfill bias”, by using all but the last 6 months of data, or all but the first 12 months of data.

Home country	U.S.A.	U.S.A.	U.K.	U.S.A.	U.S.A.
Market index	S&P 500	MSCI World	MSCI World	S&P 500	S&P 500
Dropped obs	None	None	None	Last 6	First 12
Panel I: Proportion rejected null of neutrality					
Correlation neutrality	0.298*	0.269*	0.211*	0.252*	0.280*
Mean neutrality	0.233*	0.180	0.267*	0.167	0.154
Variance neutrality	0.060	0.027	0.060	0.047	0.037
VaR neutrality	0.000	0.017	0.017	0.000	0.000
Complete neutrality	0.263*	0.292*	0.193*	0.216*	0.280*
Joint test	0.281*	0.240*	0.199*	0.199*	0.253*
Panel II: Proportion rejected null of neutrality on the downside					
Correlation neutrality	0.269*	0.251*	0.252*	0.246*	0.260*
Mean neutrality	0.007	0.013	0.007	0.013	0.007
Variance neutrality	0.040	0.033	0.033	0.033	0.037
VaR neutrality	0.000	0.017	0.017	0.000	0.000
Tail neutrality	0.039	0.039	0.000	0.039	0.000
Complete neutrality	0.251*	0.257*	0.246*	0.222*	0.280*
Joint test	0.216*	0.240*	0.222*	0.205*	0.247*

**Table 6**  
**Results using a weighted average of lagged market returns**

This table presents the results of tests of various types of market neutrality presented in the paper. The tests are conducted on monthly “market neutral” hedge fund returns over the period April 1993 to April 2003. We present results for the benchmark case when the market index is the contemporaneous return on the S&P 500,  $r_{mt}$ , and for the case when the market index is a weighted average of current and lagged returns on the S&P 500 index:  $\tilde{r}_{mt} = \sum_{j=0}^3 \alpha_j r_{mt-j} = 0.4 r_{mt} + 0.3 r_{mt-1} + 0.2 r_{mt-2} + 0.1 r_{mt-3}$ . The first two columns report the proportion of funds for which the null hypothesis of market neutrality could be rejected at the 0.05 level in favour of a general alternative. The third and fourth columns report the proportion of funds for which the null hypothesis of market neutrality could be rejected at the 0.05 level in favour of a dependence structure of a type that is specifically disliked by U.S.-based risk averse investors with a pre-existing exposure to the S&P 500. An asterisk indicates that the proportion of rejected funds is significantly different from zero (using the bootstrap procedure described in the appendix) and thus that the collection of funds taken as a whole are significantly non-neutral.

Market index	Proportion rejected null of neutrality		Proportion rejected null of neutrality on downside	
	$r_{mt}$	$\sum_{j=0}^3 \alpha_j r_{mt-j}$	$r_{mt}$	$\sum_{j=0}^3 \alpha_j r_{mt-j}$
Correlation neutrality	0.298*	0.246*	0.269*	0.216*
Mean neutrality	0.233*	0.153	0.007	0.013
Variance neutrality	0.060	0.033	0.040	0.033
VaR neutrality	0.000	0.017	0.000	0.017
Tail neutrality	-	-	0.039	0.000
Complete neutrality	0.263*	0.275*	0.251*	0.304*
Joint test	0.281*	0.222*	0.216*	0.211*

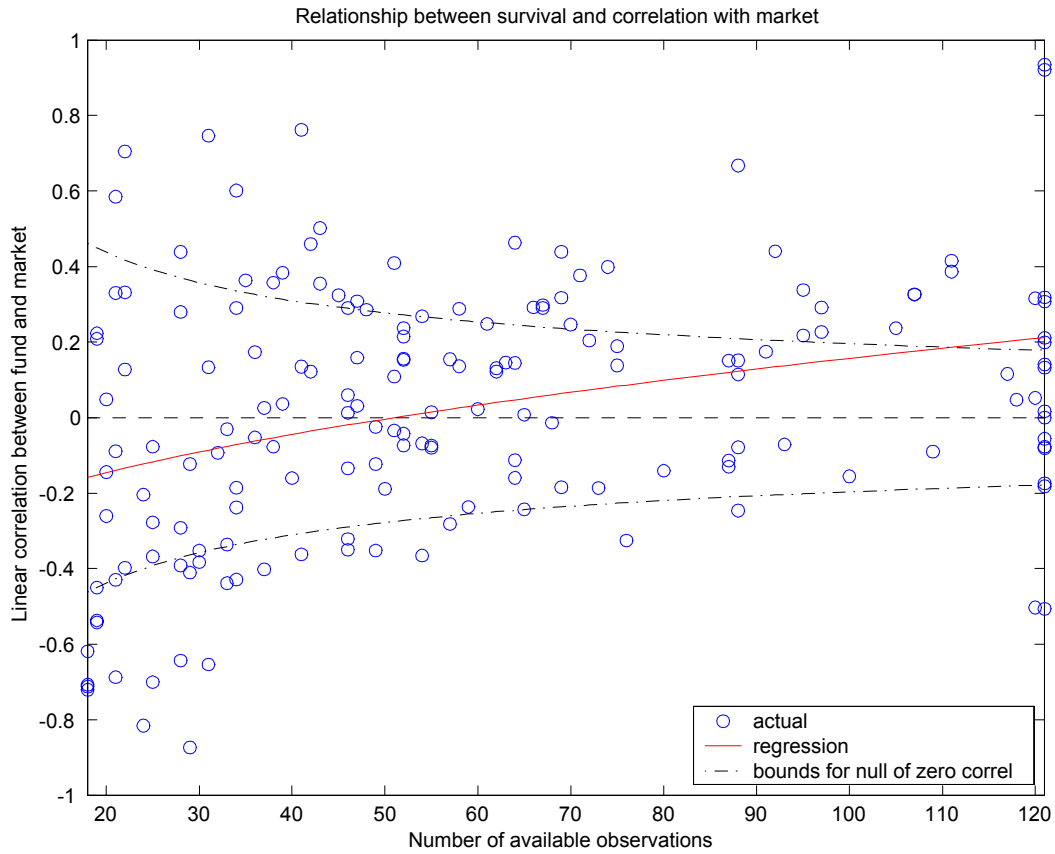


Figure 1: *The relation between linear correlation and the number of available observations on a hedge fund. The circles represent the number of observations/linear correlation with the market return observations for each fund in our sample with at least 18 observations. The solid line is the result of a linear regression of correlation on a constant and the number of observations available. The dashed lines are simple Bartlett 95% confidence bounds for testing a null of zero correlation, and equal  $\pm 1.96 \sqrt{\quad}$  (For formal tests we use serial correlation and heteroscedasticity robust confidence intervals based on a block bootstrap.)*

## References

- [1] Ackermann, Carl, McEnally, Richard and Ravenscraft, David, 1999, The Performance of Hedge Funds: Risk, Return and Incentives, *Journal of Finance*, 54(3), 833-874.
- [2] Andrews, Donald W. K., 1991, Asymptotic Normality of Series Estimators for Nonparametric and Semi-parametric Regression Models, *Econometrica*, 59(2), 307-346.
- [3] Ang, Andrew, and Bekaert, Geert, 2002, International Asset Allocation with Regime Shifts, *Review of Financial Studies*, 15(4), 1137-1187.
- [4] Ang, Andrew, and Chen, Joe, 2002, Asymmetric Correlations of Equity Portfolios, *Journal of Financial Economics*, 63(3), 443-494.
- [5] Argarwal, Vikas, and Naik, Narayan Y., 2002, Risks and Portfolio Decisions Involving Hedge Funds, forthcoming in the *Review of Financial Studies*.
- [6] Argarwal, Vikas, Daniel, Naveen D., and Naik, Narayan Y., 2003, Flows, Performance, and Managerial Incentives in the Hedge Fund Industry, working paper, London Business School.
- [7] Arrow, Kenneth J., 1971, *Essays in the Theory of Risk Bearing*, Markham Publishing Co., Chicago.
- [8] Artzner, P., Delbaen, F., Eber, J.-M. and Heath, D., 1999, Coherent Measures of Risk, *Mathematical Finance*, 9(3), 203-228.
- [9] Asness, Clifford, Krail, Robert, and Liew, John, 2001, Do Hedge Funds Hedge?, *Journal of Portfolio Management*, 28(1), 6-19.
- [10] Bae, Kee-Hong, Karolyi, G. Andrew, and Stulz, René M., 2003, A New Approach to Measuring Financial Contagion, *Review of Financial Studies*, 16(3), 717-763.
- [11] Bansal, Ravi, Hsieh, David A., and Viswanathan, S., 1993, A New Approach to International Arbitrage Pricing, *Journal of Finance*, 48(5), 1719-1747.
- [12] Bauwens, Luc, Laurent, Sébastien, and Rombouts, Jeroen, 2003, Multivariate GARCH Models: A Survey, *Journal of Applied Econometrics*, forthcoming.
- [13] Beliossi, Giovanni, 2002, Market Neutral Strategies, *Journal of Alternative Investments*, 5(2), 1-4.
- [14] Bollerslev, Tim, Engle, Robert F., and Wooldridge, Jeffrey M., 1988, A Capital Asset Pricing Model with Time-Varying Covariances, *Journal of Political Economy*, 96(1), 116-131.
- [15] Brown, Stephen J., and Goetzmann, William N., 1997, Mutual Fund Styles, *Journal of Financial Economics*, 43(3), 373-399.
- [16] Brown, Stephen J., Goetzmann, William N., and Ibbotson, Roger G., 1999, Offshore Hedge Funds: Survival and Performance 1989-1995, *Journal of Business*, 72(1), 91-117.
- [17] Caldwell, Ted, 1995, Introduction: The Model for Superior Performance, in J. Lederman and R. A. Klein eds., *Hedge Funds*, Irwin, New York.

- [18] Chapman, David, 1997, Approximating the Asset Pricing Kernel, *Journal of Finance*, 52(4), 1383-1410.
- [19] Chen, Xiaohong, and Shen Xiaotong, 1998, Sieve Extremum Estimates for Weakly Dependent Data, *Econometrica*, 66(2), 289-314.
- [20] Christoffersen, Peter F., 1998, Evaluating Interval Forecasts, *International Economic Review*, 39(4), 841-862.
- [21] Dittmar, Robert F., 2002, Nonlinear Pricing Kernels, Kurtosis Preference, and Evidence from the Cross Section of Equity Returns, *Journal of Finance*, 57(1), 369-403.
- [22] Embrechts, Paul, McNeil, Alexander and Straumann, Daniel, 2001, Correlation and Dependence Properties in Risk Management: Properties and Pitfalls, in M. Dempster, ed., *Risk Management: Value at Risk and Beyond*, Cambridge University Press.
- [23] Epstein, Larry G., and Tanny, Stephen M., 1980, Increasing Generalized Correlation: A Definition and some Economic Consequences, *Canadian Journal of Economics*, 13, 16-34.
- [24] Fung, William, and Hsieh, David A., 1997, Empirical Characteristics of Dynamic Trading Strategies: The Case of Hedge Funds, *The Review of Financial Studies*, 10(2), 275-302.
- [25] Fung, William, and Hsieh, David A., 1999, A Primer on Hedge Funds, *Journal of Empirical Finance*, 6, 309-331.
- [26] Fung, William, and Hsieh, David A., 2001, The Risk in Hedge Fund Strategies: Theory and Evidence from Trend Followers, *The Review of Financial Studies*, 14(2), 313-341.
- [27] Getmansky, Mila, Lo, Andrew W., and Makarov, Igor, 2003, An Econometric Model of Serial Correlation and Illiquidity in Hedge Fund Returns, *Journal of Financial Economics*, forthcoming.
- [28] Gupta, Anurag and Liang, Bing, 2003, Do Hedge Funds Have Enough Capital? A Value-at-Risk Approach, *Journal of Financial Economics*, forthcoming.
- [29] Harvey, Campbell R., 1989, Time-Varying Conditional Covariances in Tests of Asset Pricing Models, *Journal of Financial Economics*, 24, 289-317.
- [30] Harvey, Campbell R., and Siddique, Akhtar, 2000, Conditional Skewness in Asset Pricing Tests, *Journal of Finance*, 55(3), 1263-1295.
- [31] Hull, John C., 2003, *Options, Futures, and Other Derivatives*, Prentice Hall, New Jersey.
- [32] Jarque, C. M., and Bera, Anil K., 1980, Efficient Tests for Normality, Heteroskedasticity, and Serial Independence of Regression Residuals, *Economics Letters*, 6, 255-259.
- [33] Joe, Harry, 1997, *Multivariate Models and Dependence Concepts*, Monographs on Statistics and Applied Probability 73, Chapman and Hall, London, U.K.
- [34] Kimball, Miles, 1993, Standard Risk Aversion, *Econometrica*, 61, 589-611.
- [35] Koenker, R., and Bassett, G., 1978, Regression Quantiles, *Econometrica*, 46, 33-50.

- [36] Kraus, Alan, and Litzenberger, Robert H., 1976, Skewness Preference and the Valuation of Risk Assets, *Journal of Finance*, 31(4) 1085-1100.
- [37] Longin, François, and Solnik, Bruno, 2001, Extreme Correlation of International Equity Markets, *Journal of Finance*, 56(2), 649-676.
- [38] Nelsen, Roger B., 1999, *An Introduction to Copulas*, Springer-Verlag, New York.
- [39] Newey, Whitney K., and West, Kenneth D., 1987, A Simple, Positive Semidefinite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix, *Econometrica*, 55(3), 703-708.
- [40] Nicholas, Joseph G., 2000, *Market-Neutral Investing*, Bloomberg Press, New Jersey.
- [41] Mitchell, Mark and Pulvino, Todd, 2001, Characteristics of Risk and Return in Risk Arbitrage, *Journal of Finance*, 56(6), 2135-2175.
- [42] Politis, Dimitris N., and Romano, Joseph P., 1994, The Stationary Bootstrap, *Journal of the American Statistical Association*, 89(428), 1303-1313.
- [43] Politis, Dimitris N., and White, Halbert, 2004, Automatic block-length selection for the dependent bootstrap, *Econometric Reviews*, 23(1), 53-70.
- [44] Quintos, Carmela, 2003, Extremal Correlation for GARCH Data, working paper, Simon Graduate School of Business, University of Rochester.
- [45] Rothschild, Michael, and Stiglitz, Joseph E., 1970, Increasing Risk: I. A Definition, *Journal of Economic Theory*, 2, 225-243.
- [46] Shanken, Jay, 1990, Intertemporal Asset Pricing: An Empirical Investigation, *Journal of Econometrics*, 45, 99-120.
- [47] Westfall, Peter H., and Young, S. Stanley, 1993, *Resampling-Based Multiple Testing: Examples and Methods for p-Value Adjustment*, John Wiley & Sons, USA.