

Are “Market Neutral” Hedge Funds Really Market Neutral?

Andrew J. Patton*

London School of Economics

PRELIMINARY. COMMENTS WELCOME

11 March, 2004

Abstract

One can consider the concept of market neutrality as having “breadth” and “depth”: “breadth” reflects the number of market risks to which the hedge 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 for hedge funds. “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 between one-quarter and one-third of these funds exhibit some significant exposure to market risk.

Keywords: hedge funds, market neutrality, dependence, correlation, risk, portfolio decisions, copulas.

J.E.L. Codes: G10, G11, G19, G23.

*Helpful comments and suggestions were received from Ron Anderson, Greg Connor, Oliver Linton, Chester Spatt, Elliot Williams, and seminar participants at LSE. 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. Assets under the management of hedge funds has grown from \$50 billion in 1990 to around \$650 billion in 2003¹. The low correlation between hedge fund returns and market returns is an oft-cited favourable characteristic of hedge funds, 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-described² investment strategies or styles, and there are roughly seven categories. These categories and their size as a proportion of the total hedge fund industry are reported by Fung and Hsieh (1999) as being “event driven” (9.6%), “global” (34.3%), “global/macro” (33.1%), “market neutral” (20.0%), “sectors” (2.0%), “short sellers” (0.6%) and “long only” (0.4%). “Market neutral” funds thus make up a large fraction of the total hedge fund industry.

If the moniker “market neutral” could be taken at face value, it would offer investors a very valuable piece of information about the fund. Most hedge funds have relatively short histories of returns available, making it hard to determine the risk/return characteristics of the fund, and similarly making it hard to determine the dependence of the fund on the market, which is required when determining the optimal amount to invest in a hedge fund. Knowing that a fund is (completely) market neutral would eliminate the need to study its relationship with the market, though of course the risk/return characteristics would still remain to be estimated.

We study in detail the market neutrality of funds with the “market neutral” label. The most commonly used measure of “market neutrality” is based on correlation or *beta*: a fund is “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 studying neutrality: Fung and Hsieh (2001) suggest using payoffs from “lookback straddle” options on the market to

¹Source: The Economist, September 20, 2003.

²Brown and Goetzmann (1997) and Fung and Hsieh (1997) have instead used actual returns on funds to classify them.

approximate the pay-off structure of hedge funds. Mitchell and Pulvino (2001) and Agarwal and Naik (2002) suggest using piece-wise linear models for the hedge fund returns as a function of the market return, which is obviously related to the method proposed by Fung and Hsieh (2001).

We consider the concept of neutrality more generally than that implied by the use of correlations or betas. One can consider the concept of market neutrality as having “breadth” and “depth”. The “breadth” of the market 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 the market neutrality “depth” and propose five different neutrality concepts for hedge funds: “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&P500 or the MSCI World indices as representing the market. Our focus solely on a single equity market index may be interpreted as testing a necessary condition for neutrality to a wider set of market variables. The methods presented in this paper extend naturally when considering a collection of market variables rather than a single market variable.

Our focus on the class of “market neutral”, or more specifically “equity market neutral”, hedge funds is motivated by two factors: the first is that “market neutral” hedge funds actually define themselves by their relation (or lack thereof) with the market, making them a prime target for an investigation of different neutrality concepts. Further, 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. 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). Of course, if the each investor’s utility function was known, then funds should be ranked by expected utility, however such a case is not common in practise.

The second motivation is that this is class of hedge funds has not received a great deal of

attention in the academic literature, though it 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"³ funds.

The paper makes two main contributions. First, we propose a number of different neutrality concepts, and present statistical tests of each neutrality concept against either a general non-neutral alternative, or against only those non-neutral alternatives that are disliked by risk averse investors. 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&P500. We find that between one-quarter and one-third of the 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 are 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 observations on funds that died during the sample period. We also check for a relation between the degree of dependence between a fund and the market and the number of observations available, and find in general that older funds tend to be more dependent on the market in levels of returns, but less dependent on the market in terms of risk.

The remainder of the paper is structured as follows. In Section 2 we describe the data used

³Although 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. 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.

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 concludes.

2 Description of the data and results using correlation

Our data set consists of those funds that categorise themselves as being “market neutral” and an equity market index, the S&P500. We will focus on the S&P500 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 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 funds⁴ and one may question whether the behaviour 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, let us firstly analyse 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 5th and 95th 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.

⁴As Agarwal, *et al.* (2003) point out, these funds are misnomered, as 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.

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 statistic is surprisingly high: these funds are (self-) described as “market neutral”, possibly to more factors than our single market index, and yet close to one-third of them have significant correlations with the market. As we show in Section 4 this proportion does not change greatly when we consider other market indices.

If we instead focus our test only on deviations from zero correlation to *positive* correlation with the market, then we find 28.0% of funds with significant positive correlation, ie, over a quarter of the “market neutral” funds in our sample exhibit significant positive correlation with the market.

Correlation is just one measure of dependence, and thus this type of neutrality is but one of the many types 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 a palatable assumption 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, with an emphasis on the preferences of a risk averse investor. We present a variety of different neutrality measures, with the hope of assisting in investors’ evaluation of “market neutral” hedge funds. 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 generalisation 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 neutral

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:

$$E[r_{it}|r_{mt}] = E[r_{it}] \quad \forall r_{mt} \quad (1)$$

and corresponds to the statement that the market return does not Granger-cause the fund return in mean. 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 $\mu_i(r_{mt}) = E[r_{it}|r_{mt}]$:

$$r_{it} = \mu_i(r_{mt}) + e_{it}$$

and then test that μ_i is equal to some constant. A simple alternative would be to employ a Taylor series approximation^{5,6} to the conditional mean function:

$$\begin{aligned} r_{it} &= \beta_0 + \beta_1 r_{mt} + \beta_2 r_{mt}^2 + e_{it}, \text{ or} \\ r_{it} &= \beta_0 + \beta_1 r_{mt} + \beta_2 r_{mt}^2 + \beta_3 r_{mt}^3 + \beta_4 r_{mt}^4 + e_{it} \end{aligned} \quad (2)$$

and then test

$$\begin{aligned} H_0 &: \beta_i = 0 \text{ for all } i > 0, \text{ vs} \\ H_a &: \beta_i \neq 0 \text{ for at least one } i > 0 \end{aligned} \quad (3)$$

via a standard χ^2 test. We estimated the simple second-order polynomial⁷ model on the 150 funds with at least 24 observations, and found that for 22.0% of funds we could reject the null hypothesis of mean neutrality at the 0.05 level.

⁵Numerous authors have, in various contexts, proposed using a polynomial in the market return to explain individual asset returns, see Bansal, *et al.* (1993), Chapman (1997), Harvey and Siddique (2000) and Dittmar (2002), amongst others.

⁶Mitchell 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.

⁷Not surprisingly for financial returns with a limited sample size, the third- and fourth-order terms in the fourth-order polynomial model were not significant for most of the funds, and so we instead focussed on the second-order polynomial.

This strict 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 a 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 return when the market return is negative. That is:

$$\frac{\partial \mu_i(r_{mt})}{\partial r_{mt}} \leq 0 \text{ for } r_{mt} \leq 0$$

where $\mu_i(r_{mt}) \equiv E[r_{it}|r_{mt}]$. This version of neutrality ignores the relation between the fund and the market when the market return is positive, focussing solely on the ability of the fund to provide diversification benefits when the market return is negative.

If we use the second-order polynomial in equation (2) to approximate the conditional mean function, then

$$\frac{\partial \mu_i(r_{mt})}{\partial r_{mt}} = \beta_1 + 2\beta_2 r_{mt}$$

A point-wise confidence interval for the first derivative of the conditional mean function is simple to construct using the covariance matrix of the estimated parameters from equation (2), and we can then determine whether the first derivative of the conditional mean function is significantly greater than zero, for some $r_{mt} \leq 0$. However, using point-wise confidence intervals and searching across all values of r_{mt} leads to a size distortion in the test, and so we instead conduct a test on the average value of this derivative across values of $r_{mt} \leq 0$. We test the following hypothesis:

$$\begin{aligned} H_0 & : E \left[\left. \frac{\partial \mu_i(r_{mt})}{\partial r_{mt}} \right| r_{mt} \leq 0 \right] \leq 0, \text{ vs} \\ H_a & : E \left[\left. \frac{\partial \mu_i(r_{mt})}{\partial r_{mt}} \right| r_{mt} \leq 0 \right] > 0 \end{aligned} \tag{4}$$

Using the second-order polynomial model, and plugging in the parameter estimates we can see that

$$E \left[\frac{\partial \hat{\mu}_i(r_{mt})}{\partial r_{mt}} \Big| r_{mt} \leq 0 \right] = \hat{\beta}_1 + 2\hat{\beta}_2 E[r_{mt} | r_{mt} \leq 0]$$

and so testing that $E[\partial \mu_i(r_{mt}) / \partial r_{mt} | r_{mt} \leq 0] \leq 0$ reduces to checking the point-wise confidence interval on the first derivative of the conditional mean function at the point $\hat{E}[r_{mt} | r_{mt} \leq 0]$. We did this, again using the bootstrap distribution of the test statistic rather than the asymptotic distribution, and found that we were able to reject “mean neutrality on the downside” for 25.3% of funds at the 0.05 level. Thus over one quarter of all “market neutral” funds exhibit conditional mean dependence on the market of a form disliked by risk averse investors.

One concern about the above regressions of the fund return on powers of the market return and the subsequent tests on functions of the estimated parameters is about omitted variables bias. It is quite likely that other factors affect r_{it} , and if any of those are also correlated with powers of r_{mt} then the parameters will be biased and the test results misleading. A worrying case would be where r_{it} follows an AR(1) and r_{it} Granger-causes $r_{m,t+1}$. In this case we may find evidence against mean neutrality even though the fund return is truly mean neutral. This is quite an unlikely scenario. If one were truly concerned about this possibility then including r_{it-1} in the regression in equation (2) would solve the problem. We did this and the results did not change significantly. A similar remedy could be applied using any other variable that one thought might be correlated with both the fund return and the market return.

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 we consider risk as measured by variance, and in the next section we consider risk as measured by Value-at-Risk. While “mean neutrality” has been tested and studied previously in various ways, to our knowledge this paper is the first to consider the market neutrality of the risk of a hedge fund.

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 returns⁸. Together these imply that risk averse investors prefer

$$\begin{aligned} \text{Corr} \left[(r_{it} - \mu_i)^2, r_{mt} - \mu_m \right] &\geq 0, \text{ and} \\ \text{Corr} \left[(r_{it} - \mu_i)^2, (r_{mt} - \mu_m)^2 \right] &\leq 0 \end{aligned}$$

so that the skewness of a portfolio of the fund and the market is larger and the kurtosis of the portfolio is smaller.

In a similar manner to the previous section, we could consider approximating the true conditional variance function, $\sigma_i^2(r_{mt})$, by a Taylor series polynomial

$$r_{it} = \mu_i(r_{mt}) + e_{it} \quad (5)$$

$$e_{it} = \sigma_i(r_{mt}) \varepsilon_{it}, \quad \varepsilon_{it} \sim (0, 1) \quad (6)$$

$$\mu_i(r_{mt}) = \beta_0 + \beta_1 r_{mt} + \beta_2 r_{mt}^2$$

$$\sigma_i^2(r_{mt}) = \alpha_0 + \alpha_1 r_{mt} + \alpha_2 r_{mt}^2, \text{ or} \quad (7)$$

$$\sigma_i^2(r_{mt}, e_{it-1}) = \alpha_0 + \alpha_1 r_{mt} + \alpha_2 r_{mt}^2 + \alpha_3 e_{it-1}^2 \quad (8)$$

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

$$H_0 : \alpha_1 = \alpha_2 = 0, \text{ vs} \quad (9)$$

$$H_a : \alpha_1 \neq 0 \cup \alpha_2 \neq 0$$

We did this for 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 9.3% 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 that the desired sign of the first derivative of the conditional variance function is positive when the market return is negative. Again, rather than search over

⁸Related papers on investor preferences over higher-order moments include Kraus and Litzenberger (1976), Harvey and Siddique (2000) and Dittmar (2002), amongst many others.

all values of $r_{mt} \leq 0$, we instead focus on the average value of the first derivative. So we seek to test the hypothesis:

$$\begin{aligned} H_0 & : E \left[\frac{\partial \sigma_i^2(r_{mt})}{\partial r_{mt}} \middle| r_{mt} \leq 0 \right] \geq 0, \text{ vs} \\ H_a & : E \left[\frac{\partial \sigma_i^2(r_{mt})}{\partial r_{mt}} \middle| r_{mt} \leq 0 \right] < 0 \end{aligned} \quad (10)$$

Following the same method as for the test of mean neutrality, we can again quite easily obtain a confidence interval for this first derivative and check that it is non-negative at $\hat{E}[r_{mt}|r_{mt} \leq 0]$. We did this for the funds in our database and found significant violations of variance neutrality on the downside for 9.3% when including an ARCH(1) term in the variance specification.

Overall, controlling for violations of mean neutrality, only about one-tenth of the hedge funds in our sample exhibited some significant conditional variance dependence on the market return.

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, usually the 0.10, 0.05 or 0.01 percentile, this could also be called “quantile neutrality”. This 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 extreme quantiles rather than moments. A VaR neutral portfolio is one with a VaR that is unaffected by the market portfolio return. That is:

$$VaR(r_{it}|r_{mt}) = VaR(r_{it}) \quad (11)$$

Violations of mean neutrality or variance neutrality will generally lead to violations of VaR neutrality, which leads us to consider “conditional VaR neutrality⁹”:

$$VaR \left(\frac{r_{it} - \mu_i(r_{mt})}{\sigma_i(r_{mt})} \middle| r_{mt} \right) = VaR \left(\frac{r_{it} - \mu_i(r_{mt})}{\sigma_i(r_{mt})} \right) \quad (12)$$

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

⁹Note 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.”

If the market and fund returns are jointly elliptically distributed then the portfolio VaR is an affine function of the portfolio variance, and VaR neutrality will then follow directly from mean and variance neutrality. Under normality, conditional VaR neutrality will always hold, even if mean and variance neutrality do 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

$$H_0 : VaR(r_{it}|r_{mt}) = VaR(r_{it}) \forall r_{mt}, \text{ vs} \quad (13)$$

$$H_a : VaR(r_{it}|r_{mt}) \neq VaR(r_{it}) \text{ for some } r_{mt}$$

or

$$H_0 : VaR\left(\frac{r_{it} - \mu_i(r_{mt})}{\sigma_i(r_{mt})} \middle| r_{mt}\right) = VaR\left(\frac{r_{it} - \mu_i(r_{mt})}{\sigma_i(r_{mt})}\right) \forall r_{mt} \quad (14)$$

$$H_a : VaR\left(\frac{r_{it} - \mu_i(r_{mt})}{\sigma_i(r_{mt})} \middle| r_{mt}\right) \neq VaR\left(\frac{r_{it} - \mu_i(r_{mt})}{\sigma_i(r_{mt})}\right) \text{ for some } r_{mt}$$

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:

$$\begin{aligned} H_0 : & \Pr\left[\varepsilon_{it} \leq \widehat{VaR}(\varepsilon_{it}) \middle| \varepsilon_{mt} \leq \widehat{VaR}(\varepsilon_{mt})\right] = \Pr\left[\varepsilon_{it} \leq \widehat{VaR}(\varepsilon_{it}) \middle| \varepsilon_{mt} > \widehat{VaR}(\varepsilon_{mt})\right] \cap \\ & \Pr\left[\varepsilon_{it} \leq \widehat{VaR}(\varepsilon_{it}) \middle| \varepsilon_{mt} \leq \widehat{VaR}(\varepsilon_{mt})\right] = \Pr\left[\varepsilon_{it} \leq \widehat{VaR}(\varepsilon_{it})\right] \quad \text{vs} \quad (15) \\ H_a : & \Pr\left[\varepsilon_{it} \leq \widehat{VaR}(\varepsilon_{it}) \middle| \varepsilon_{mt} \leq \widehat{VaR}(\varepsilon_{mt})\right] \neq \Pr\left[\varepsilon_{it} \leq \widehat{VaR}(\varepsilon_{it})\right] \cup \\ & \Pr\left[\varepsilon_{it} \leq \widehat{VaR}(\varepsilon_{it}) \middle| \varepsilon_{mt} > \widehat{VaR}(\varepsilon_{mt})\right] \neq \Pr\left[\varepsilon_{it} \leq \widehat{VaR}(\varepsilon_{it})\right] \end{aligned}$$

where $\varepsilon_{it} \equiv (r_{it} - \mu_i(r_{mt}))/\sigma_i(r_{mt})$ and $\varepsilon_{mt} \equiv (r_{mt} - \mu_{mt})/\sigma_{mt}$. For the fund we again use a second-order polynomial for the conditional mean, along with an AR(1) term, 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. $\widehat{VaR}(\varepsilon_{it})$ and $\widehat{VaR}(\varepsilon_{mt})$ are estimated by the empirical quantiles of ε_{it}

and ε_{mt} . We ignore the estimation error in the mean and variance models, and in the empirical quantile estimates.

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 (which controls for mean and variance non-neutrality) on these funds and found evidence against VaR neutrality for only 3.4% of funds at the 0.05 level. Thus we conclude that we have no evidence against VaR neutrality for these funds, having controlled for mean and variance non-neutrality.

We also consider a “downside” version of this test, which focusses specifically on

$$\begin{aligned} H_0 & : \Pr \left[\varepsilon_{it} \leq \widehat{VaR}(\varepsilon_{it}) \mid \varepsilon_{mt} \leq \widehat{VaR}(\varepsilon_{mt}) \right] \leq \Pr \left[\varepsilon_{it} \leq \widehat{VaR}(\varepsilon_{it}) \right] \text{ vs} \\ H_a & : \Pr \left[\varepsilon_{it} \leq \widehat{VaR}(\varepsilon_{it}) \mid \varepsilon_{mt} \leq \widehat{VaR}(\varepsilon_{mt}) \right] > \Pr \left[\varepsilon_{it} \leq \widehat{VaR}(\varepsilon_{it}) \right] \end{aligned} \quad (16)$$

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 exceedence 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 only 3.4% of funds are rejected at the 0.05 level, and so again conclude that we have no evidence against the VaR neutrality of these funds.

3.4 Tail neutrality

Finally, we consider the concept of neutrality during extreme events, or ‘tail neutrality’. Intuitively this can be thought of as an extension of VaR neutrality: 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:

$$\tau^L \equiv \lim_{\varepsilon \rightarrow 0} \Pr [F_i(r_i) < \varepsilon \mid F_m(r_m) < \varepsilon] = \lim_{\varepsilon \rightarrow 0} \Pr [F_i(r_i) < \varepsilon] = 0 \quad (17)$$

where $r_i | \Omega_{t-1} \sim F_i$ and $r_m | \Omega_{t-1} \sim F_m$. 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 τ^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 $\tau^L > 0$. 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 dependence¹⁰. 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. 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 behaviour 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.

¹⁰We 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. This is equivalent to the statement that there is no Granger-causality (in distribution) from the market return to the fund return. The formal definition is:

$$r_i | r_m =^d r_i \quad (18)$$

where “ $=^d$ ” indicates equality in distribution. If we let $r_{it} \sim F_i$ and $r_{mt} \sim F_m$, and $(r_{it}, r_{mt}) \sim F$, this implies that

$$f(r_{it}, r_{mt}) = f_i(r_{it}) \cdot f_m(r_{mt}) \quad (19)$$

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

$$f(r_{it}, r_{mt}) = f_i(r_{it}) \cdot f_m(r_{mt}) \cdot c(F_i(r_{it}), F_m(r_{mt})) \quad (20)$$

where c is the “copula density” or “dependence function” of the fund and the market returns. Under complete neutrality the copula¹¹ of the fund and the market 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).

In considering alternatives to complete neutrality we should be more general than simply allowing for non-zero correlation. A general alternative to complete neutrality is a dependence function, C^* , that differs from C_I solely 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 realisations where both variables are “large” or “small” and away from realisations 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:

$$C_I(u, v) \leq C^*(u, v) \quad \forall (u, v) \in [0, 1] \times [0, 1] \quad (21)$$

¹¹The copula *cdf* is denoted with an upper case C while the copula density is denoted with a lower case c . See Nelsen (1998) for an introduction to copulas.

and we say that C^* is “more concordant¹²” than C_I , or simply that $C_I \leq C^*$. This ordering has obvious similarities to first-order stochastic dominance.

Epstein and Tanny (1980) define a utility function involving two random variables to be “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

$$\frac{\partial^2 u(x, y)}{\partial x \partial y} < 0 \quad (22)$$

In this paper we consider the case that:

$$\begin{aligned} u(r_{it}, r_{mt}) &= \mathcal{U}(w_i r_{it} + w_m r_{mt}), \text{ where } \mathcal{U} \text{ is some utility function} \\ \frac{\partial^2 u(r_{it}, r_{mt})}{\partial r_{it} \partial r_{mt}} &= \mathcal{U}''(w_i r_{it} + w_m r_{mt}) w_i w_m \\ &\leq 0 \text{ if } w_i, w_m \geq 0 \end{aligned}$$

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 $w_i, w_m \geq 0$. The weak inequality above holds strictly if $w_i w_m > 0$; of course if either portfolio weight is zero then the investor is “correlation neutral”. Thus a risk averse investor subject to short selling constraints will be (weakly) “correlation averse”.

We can use the concept of correlation aversion to derive a ranking of dependence functions for risk averse investors. If F_1 and F_2 are two possible joint distribution functions for (r_{it}, r_{mt}) with common marginal distributions, and if

$$E_{F_1}[u(r_{it}, r_{mt})] \leq E_{F_2}[u(r_{it}, r_{mt})]$$

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,

$$F_2 \preceq_u F_1 \Leftrightarrow F_2 \leq F_1 \quad (23)$$

¹²Epstein and Tanny (1980) interpret the condition in equation (21) 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.

In terms of dependence functions, this implies that

$$C_I \preceq_u C^* \Leftrightarrow C_I \leq C^* \quad (24)$$

and so any dependence function that is a CIT away from independence 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,

$$C_I \leq C^* \Rightarrow \rho_S(C_I) \leq \rho_S(C^*) \quad (25)$$

and so we have

$$C_u^{**} \preceq_u C_I \preceq_u C^* \Leftrightarrow C^{**} \leq C_I \leq C^* \Rightarrow \rho_S(C^{**}) \leq \rho_S(C_I) \leq \rho_S(C^*) \quad (26)$$

Thus we may obtain an approximate ordering of the funds for a general risk averse investor by categorising 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 25.7% of funds had significant rank correlation at the 0.05 level, and 24.6% 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 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 not market 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 probably not 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 95th 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. A summary of the results obtained thus far is presented in Table 3.

[INSERT TABLE 3 ABOUT HERE]

At the 0.05 level, we found that 32.7% of funds failed a joint test of market neutrality against general non-neutral alternatives, while 24.4% of funds failed a joint test of market neutrality against alternatives that are disliked by risk averse investors. These figures are roughly consistent with the figures we presented for the individual tests. Overall, we conclude that between one-quarter and one-third of “market neutral” funds exhibit significant deviations from market neutrality. 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 cast some doubt on the label of “market neutral” for these funds, and 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.

Tables 4 to 7 report the results of tests for different versions of neutrality on these four categories of hedge funds. These tables show 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 70% of equity non-hedge funds, for example, exhibit some significant violation of market neutrality at the 0.05 level, and over 87% 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. Hedge funds in the funds of funds category are the most market neutral of these four categories, with about 45% of 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 fund categories.

[INSERT TABLES 4, 5, 6 AND 7 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 analyse whether our results change when we drop the last six 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&P500 index as the market index, and a summary of results for this case are collected in Table 3. Corresponding results when the MSCI World index is instead used are presented in Table 8. Comparing these two tables shows that our results are robust to this choice. We also checked the results (not reported) when the MSCI Europe index was employed and again no substantial differences were found.

[INSERT TABLES 8, 9, AND 10 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 Table 9. There are no major differences in the results.

End-game behaviour. In the months leading up to a fund dropping out of the HFR or TASS databases it is conceivable that the behaviour 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&P500 index and MSCI World index (not reported) as the market index, dropping the last six observations on each fund. The results are presented in Table 10, and are not substantially different from the results in Table 3.

Age of the fund and its market neutrality. In the paper 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&P500 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 fund¹³. 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

¹³The observant 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 coefficient of over 0.9. We re-did the regression without these observations and the relation was still significantly positive.

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 emerges when using rank correlation coefficients. 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.

[INSERT FIGURE 1 ABOUT HERE]

When we looked at the test results for mean, variance and VaR neutrality as a function of the number of observations available, an interesting pattern emerged. Consistent with the results for linear and rank correlation, the probability that a fund violates mean neutrality significantly increases with the number of observations available. However, the probability that a fund violates variance neutrality significantly *decreases* with the number of observations available. Thus older funds tend to be less mean neutral but more variance neutral. No significant pattern was found for VaR neutrality or tail neutrality.

5 Conclusions

“Market neutral” hedge fund manage about 20% of the \$650 billion 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 generalising 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 against either a general non-neutral alternative hypothesis, or a non-neutral alternative hypothesis that focusses solely on deviations from neutrality that would be disliked by a risk averse investor. We apply the tests to a combined database of monthly

“market neutral” hedge fund returns from the HFR and TASS hedge fund databases over the period April 1993 to April 2003, using a block bootstrap method to deal with serial correlation, volatility clustering and non-normality of the returns. We use data on 194 live and 23 dead “market neutral” hedge funds to evaluate their neutrality against a market index, the S&P500.

Although the average correlation, across funds, between fund returns and market returns is only 0.016, we found that about one-third of “market neutral” funds are significantly non-neutral in some way, at the 0.05 level, and about one-quarter of these funds are significantly non-neutral in a way that is specifically disliked by risk averse investors. In a series of robustness checks we verified 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 observations on funds that died during the sample period. We compared these results with those obtained by looking at “equity hedge”, “equity non-hedge”, “event driven” and “fund of fund” hedge funds, and found strong evidence that “market neutral” funds are more neutral to market risks than these funds.

Overall, our results suggest that the dependence between hedge fund returns and market returns is often significant and positive, even for “market neutral” funds. The widely-cited diversification benefits from investing in hedge funds thus may not be as great as first thought. Some analysis of a fund’s co-movements with the market is required to determine whether the fund is offering the degree and type of market neutrality desired by the investor.

The work in this paper leaves unanswered many interesting questions. For example, one could examine the return-dependence trade-off, according to the different neutrality measures. Having shown that risk averse investors prefer one particular dependence structure over another, one would expect that a fund that has a less desirable dependence structure with the market would compensate its investors with a higher expected return. Ang, *et al.* (2001) conduct such a study for U.S. equities, and find evidence that stocks with an undesirable dependence structures with the market command a premium of up to 6.5% per year.

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 concern to this sector.

6 Appendix: Details of the bootstrap tests

Concerns about serial correlation, volatility clustering, and non-normality of asset returns prompted us to use a block bootstrap. We used the stationary bootstrap of Politis and Romano (1994), and the algorithm of Politis and White (2003) to determine the 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.

We imposed the condition that the bootstrapped fund returns were independent of the bootstrapped market returns by re-sampling each of these series 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.

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 95th 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 95th percentile would be 1: the *cdf* of a Binomial(5, 0.05) evaluated at 1 is 0.9774.) We conclude that a fund failed the *joint* test of market neutrality if it failed more tests than the 95th 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.

	Median fund	S&P500	MSCI Europe	MSCI World
Mean*	8.0825	10.1602	8.5604	6.6308
Std dev*	7.9488	15.6600	15.9666	14.7250
Skewness	0.0261	-0.5424	-0.3655	-0.5425
Kurtosis	3.6547	3.2561	3.7360	3.2753
Min	-4.4800	-14.4431	-13.1550	-13.3503
Max	5.8800	9.7766	13.4872	9.0228
JB stat	2.5761	5.9932	4.9222	6.0362
p-value	0.2758	0.0500	0.0853	0.0489
Number of obs	42	121	121	121

Notes: This table presents some summary 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 217 funds with more than 6 observations. The mean and standard deviation statistics have been annualised to ease interpretation. ‘JB stat’ refers to the Jarque-Bera (1980) test of normality, the p-value for this test is also reported.

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

Notes: This table presents some summary statistics on the number of observations available on each of the categories of hedge funds used in this paper. The bottom three rows present the number of funds available for analysis in the tests of the various types of market neutrality.

	Proportion rejected: Neutral	Proportion rejected: Neutral on downside
Linear correlation	0.2924	0.2807
Mean neutrality	0.2200	0.2533
Variance neutrality	0.0933	0.0933
VaR neutrality	0.0339	0.0339
Tail neutrality	-	0.0392
Rank correlation	0.2573	0.2456
Joint test	0.3272	0.2442

Notes: This table presents a summary of the results of tests of various types of market neutrality. The column “Proportion rejected: Neutral” reports the proportion of funds for which a null hypothesis of market neutrality (using the various versions of this concept) could be rejected at the 0.05 level in favour of a general alternative. The column “Proportion rejected: Neutral on the downside” reports the proportion of funds for which a null hypothesis of market neutrality (using the various versions of this concept) could be rejected at the 0.05 level in favour of a dependence structure of a type that is disliked by risk averse investors.

	Proportion rejected: Neutral	Proportion rejected: Neutral on downside
Linear correlation	0.5579	0.5687
Mean neutrality	0.5190	0.5687
Variance neutrality	0.0664	0.1066
VaR neutrality	0.0165	0.0165
Tail neutrality	-	0.0053
Rank correlation	0.5879	0.5751
Joint test	0.5201	0.05131

Notes: This table presents a summary of the results of tests of various types of market neutrality, applied to “equity hedge” funds. Returns are in US dollars and the market index used is the S&P500. The column “Proportion rejected: Neutral” reports the proportion of funds for which a null hypothesis of market neutrality (using the various versions of this concept) could be rejected at the 0.05 level in favour of a general alternative. The column “Proportion rejected: Neutral on the downside” reports the proportion of funds for which a null hypothesis of market neutrality (using the various versions of this concept) could be rejected at the 0.05 level in favour of a dependence structure of a type that is disliked by risk averse investors.

Table 5: Results for “equity non-hedge” funds		
	Proportion rejected: Neutral	Proportion rejected: Neutral on downside
Linear correlation	0.8571	0.8701
Mean neutrality	0.7922	0.8052
Variance neutrality	0.1039	0.1688
VaR neutrality	0.0517	0.0517
Tail neutrality	-	0.4167
Rank correlation	0.8831	0.8701
Joint test	0.7158	0.7158

Notes: This table presents a summary of the results of tests of various types of market neutrality, applied to “equity non-hedge” funds. Returns are in US dollars and the market index used is the S&P500. The column “Proportion rejected: Neutral” reports the proportion of funds for which a null hypothesis of market neutrality (using the various versions of this concept) could be rejected at the 0.05 level in favour of a general alternative. The column “Proportion rejected: Neutral on the downside” reports the proportion of funds for which a null hypothesis of market neutrality (using the various versions of this concept) could be rejected at the 0.05 level in favour of a dependence structure of a type that is disliked by risk averse investors.

Table 6: Results for “event driven” hedge funds		
	Proportion rejected: Neutral	Proportion rejected: Neutral on downside
Linear correlation	0.6444	0.7000
Mean neutrality	0.7093	0.7442
Variance neutrality	0.1279	0.2442
VaR neutrality	0.0638	0.0638
Tail neutrality	-	0.0769
Rank correlation	0.5667	0.6667
Joint test	0.6273	0.6546

Notes: This table presents a summary of the results of tests of various types of market neutrality, applied to “event driven” hedge funds. Returns are in US dollars and the market index used is the S&P500. The column “Proportion rejected: Neutral” reports the proportion of funds for which a null hypothesis of market neutrality (using the various versions of this concept) could be rejected at the 0.05 level in favour of a general alternative. The column “Proportion rejected: Neutral on the downside” reports the proportion of funds for which a null hypothesis of market neutrality (using the various versions of this concept) could be rejected at the 0.05 level in favour of a dependence structure of a type that is disliked by risk averse investors.

	Proportion rejected: Neutral	Proportion rejected: Neutral on downside
Linear correlation	0.4730	0.5011
Mean neutrality	0.5745	0.6256
Variance neutrality	0.1353	0.1836
VaR neutrality	0.0237	0.0237
Tail neutrality	-	0.0294
Rank correlation	0.4773	0.4989
Joint test	0.4535	0.4503

Notes: This table presents a summary of the results of tests of various types of market neutrality, applied to funds of hedge funds. Returns are in US dollars and the market index used is the S&P500. The column “Proportion rejected: Neutral” reports the proportion of funds for which a null hypothesis of market neutrality (using the various versions of this concept) could be rejected at the 0.05 level in favour of a general alternative. The column “Proportion rejected: Neutral on the downside” reports the proportion of funds for which a null hypothesis of market neutrality (using the various versions of this concept) could be rejected at the 0.05 level in favour of a dependence structure of a type that is disliked by risk averse investors.

	Proportion rejected: Neutral	Proportion rejected: Neutral on downside
Linear correlation	0.2807	0.2749
Mean neutrality	0.2200	0.2600
Variance neutrality	0.0800	0.0867
VaR neutrality	0.0339	0.0339
Tail neutrality	-	0.0392
Rank correlation	0.2807	0.2632
Joint test	0.2133	0.2304

Notes: This table presents a summary of the results of tests of various types of market neutrality. The column “Proportion rejected: Neutral” reports the proportion of “market neutral” funds for which a null hypothesis of market neutrality (using the various versions of this concept) could be rejected at the 0.05 level in favour of a general alternative. The column “Proportion rejected: Neutral on the downside” reports the proportion of funds for which a null hypothesis of market neutrality (using the various versions of this concept) could be rejected at the 0.05 level in favour of a dependence structure of a type that is disliked by risk averse investors.

	Proportion rejected: Neutral	Proportion rejected: Neutral on downside
Linear correlation	0.2164	0.2749
Mean neutrality	0.3467	0.4533
Variance neutrality	0.0200	0.0533
VaR neutrality	0.0000	0.0000
Tail neutrality	-	0.0000
Rank correlation	0.2164	0.2515
Joint test	0.2581	0.2396

Notes: This table presents a summary of the results of tests of various types of market neutrality. The column “Proportion rejected: Neutral” reports the proportion of “market neutral” funds for which a null hypothesis of market neutrality (using the various versions of this concept) could be rejected at the 0.05 level in favour of a general alternative. The column “Proportion rejected: Neutral on the downside” reports the proportion of funds for which a null hypothesis of market neutrality (using the various versions of this concept) could be rejected at the 0.05 level in favour of a dependence structure of a type that is disliked by risk averse investors.

	Proportion rejected: Neutral	Proportion rejected: Neutral on downside
Linear correlation	0.2807	0.2690
Mean neutrality	0.2000	0.2333
Variance neutrality	0.0933	0.1000
VaR neutrality	0.0678	0.0678
Tail neutrality	-	0.0392
Rank correlation	0.2632	0.2573
Joint test	0.3134	0.2489

Notes: This table presents a summary of the results of tests of various types of market neutrality. The column “Proportion rejected: Neutral” reports the proportion of “market neutral” funds for which a null hypothesis of market neutrality (using the various versions of this concept) could be rejected at the 0.05 level in favour of a general alternative. The column “Proportion rejected: Neutral on the downside” reports the proportion of funds for which a null hypothesis of market neutrality (using the various versions of this concept) could be rejected at the 0.05 level in favour of a dependence structure of a type that is disliked by risk averse investors.

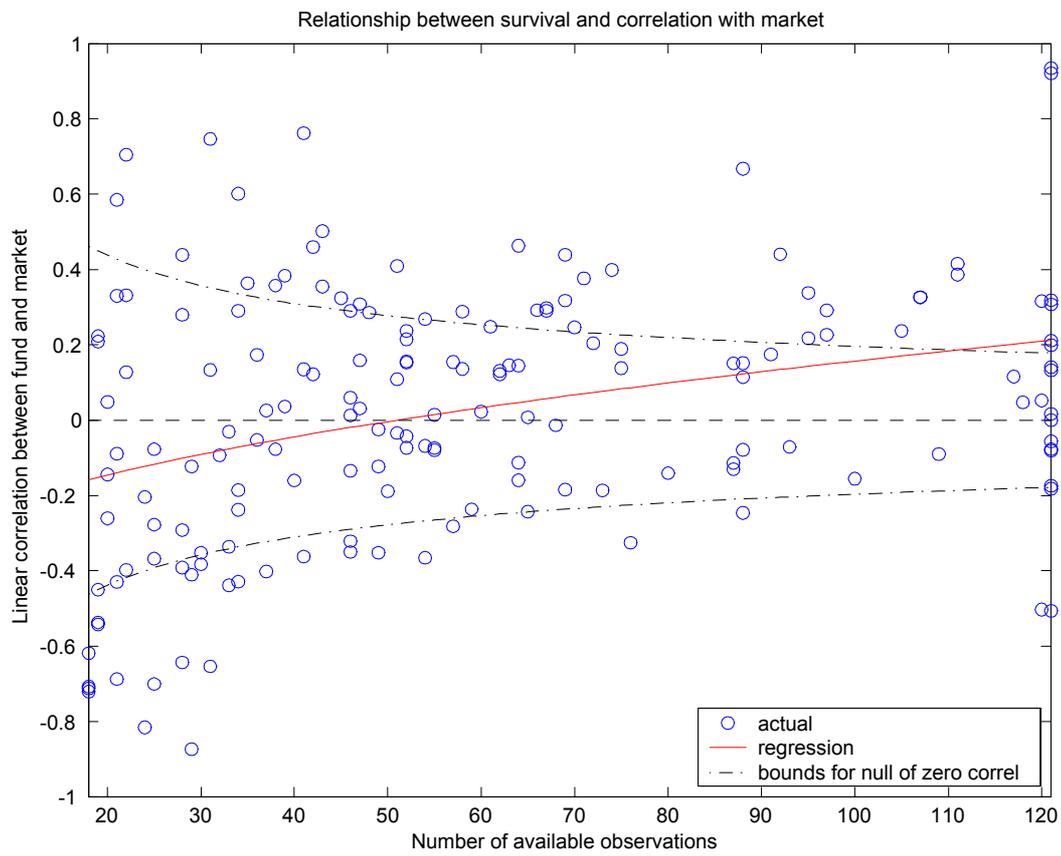


Figure 1: *The relation between linear correlation and the number of available observations on a hedge fund.*

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, Chen, Joseph, and Xing, Yuhang, 2001, Downside Risk and Expected Returns, working paper, Columbia Business School.
- [4] Argarwal, Vikas, and Naik, Narayan Y., 2002, Risks and Portfolio Decisions Involving Hedge Funds, forthcoming in the *Review of Financial Studies*.
- [5] 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.
- [6] Arrow, Kenneth J., 1971, *Essays in the Theory of Risk Bearing*, Markham Publishing Co., Chicago.
- [7] Artzner, P., Delbaen, F., Eber, J.-M. and Heath, D., 1999, Coherent Measures of Risk, *Mathematical Finance*, 9(3), 203-228.
- [8] 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.
- [9] Bansal, Ravi, Hsieh, David A., and Viswanathan, S., 1993, A New Approach to International Arbitrage Pricing, *Journal of Finance*, 48(5), 1719-1747.
- [10] Brown, Stephen J., and Goetzmann, William N., 1997, Mutual Fund Styles, *Journal of Financial Economics*, 43(3), 373-399.
- [11] 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.
- [12] Brown, Stephen J., Goetzmann, William N., and Park, James, 2001, Careers and Survival: Competition and Risk in the Hedge Fund and CTA Industry, *Journal of Finance*, 56(5), 1869-1886.
- [13] Caldwell, Ted, 1995, Introduction: The Model for Superior Performance, in J. Lederman and R. A. Klein eds., *Hedge Funds*, Irwin, New York.
- [14] Chapman, David, 1997, Approximating the Asset Pricing Kernel, *Journal of Finance*, 52(4), 1383-1410.
- [15] Chen, Xiaohong, and Shen Xiaotong, 1998, Sieve Extremum Estimates for Weakly Dependent Data, *Econometrica*, 66(2), 289-314.
- [16] Christoffersen, Peter F., 1998, Evaluating Interval Forecasts, *International Economic Review*, 39(4), 841-862.

- [17] 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.
- [18] 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.
- [19] Epstein, Larry G., and Tanny, Stephen M., 1980, Increasing Generalized Correlation: A Definition and some Economic Consequences, *Canadian Journal of Economics*, 13, 16-34.
- [20] 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.
- [21] Fung, William, and Hsieh, David A., 1999, A Primer on Hedge Funds, *Journal of Empirical Finance*, 6, 309-331.
- [22] Fung, William, and Hsieh, David A., 2000, Performance Characteristics of Hedge Funds and Commodity Funds: Natural vs. Spurious Biases, *Journal of Financial and Quantitative Analysis*, 35(3), 291-307.
- [23] 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.
- [24] Fung, William, and Hsieh, David A., 2003, Asset-Based Style Factors for Hedge Funds, forthcoming in the *Financial Analysts Journal*.
- [25] Granger, Clive W. J., 1969, Investigating Causal Relations by Econometric Models and Cross-Spectral Methods, *Econometrica*, 37, 424-438.
- [26] Gupta, Anurag and Liang, Bing, 2003, Do Hedge Funds Have Enough Capital? A Value-at-Risk Approach, working paper, Weatherhead School of Management, Case Western Reserve University.
- [27] Harvey, Campbell R., and Siddique, Akhtar, 2000, Conditional Skewness in Asset Pricing Tests, *Journal of Finance*, 55(3), 1263-1295.
- [28] Hull, John C., 2003, *Options, Futures, and Other Derivatives*, Prentice Hall, New Jersey.
- [29] Jarque, C. M., and Bera, Anil K., 1980, Efficient Tests for Normality, Heteroskedasticity, and Serial Independence of Regression Residuals, *Economics Letters*, 6, 255-259.
- [30] Joe, Harry, 1997, *Multivariate Models and Dependence Concepts*, Monographs on Statistics and Applied Probability 73, Chapman and Hall, London, U.K.
- [31] Kimball, Miles, 1993, Standard Risk Aversion, *Econometrica*, 61, 589-611.
- [32] Koenker, R., and Bassett, G., 1978, Regression Quantiles, *Econometrica*, 46, 33-50.
- [33] Kraus, Alan, and Litzenberger, Robert H., 1976, Skewness Preference and the Valuation of Risk Assets, *Journal of Finance*, 31(4) 1085-1100.

- [34] Longin, François, and Solnik, Bruno, 2001, Extreme Correlation of International Equity Markets, *Journal of Finance*, 56(2), 649-676.
- [35] Nelsen, Roger B., 1999, *An Introduction to Copulas*, Springer-Verlag, New York.
- [36] Nicholas, Joseph G., 2000, *Market-Neutral Investing*, Bloomberg Press, New Jersey.
- [37] Mitchell, Mark and Pulvino, Todd, 2001, Characteristics of Risk and Return in Risk Arbitrage, *Journal of Finance*, 56(6), 2135-2175.
- [38] Politis, Dimitris N., and Romano, Joseph P., 1994, The Stationary Bootstrap, *Journal of the American Statistical Association*, 89(428), 1303-1313.
- [39] Politis, Dimitris N., and White, Halbert, 2004, Automatic block-length selection for the dependent bootstrap, *Econometric Reviews*, 23(1), 53-70.
- [40] Quintos, Carmela, 2003, Extremal Correlation for GARCH Data, working paper, Simon Graduate School of Business, University of Rochester.
- [41] Rothschild, Michael, and Stiglitz, Joseph E., 1970, Increasing Risk: I. A Definition, *Journal of Economic Theory*, 2, 225-243.
- [42] Sharpe, William F., 1992, Asset Allocation: Management Style and Performance Measurement, *Journal of Portfolio Management*, 18, 7-19.
- [43] Sklar, A., 1959, Fonctions de répartition à n dimensions et leurs marges, *Publ. Inst. Statis. Univ. Paris*, 8, 229-231.
- [44] Westfall, Peter H., and Young, S. Stanley, 1993, *Resampling-Based Multiple Testing: Examples and Methods for p-Value Adjustment*, John Wiley & Sons, USA.