Market-neutral hedge funds and asset markets: tail or two-state dependence?*

Julio A. Crego
CEMFI
<jacrego@cemfi.edu.es>

Julio Gálvez
CEMFI
<galvez@cemfi.edu.es>

Abstract

We reconcile opposing evidence found in previous literature about the tail neutrality of market-neutral hedge funds (MNHFs) using US data from 2003 to 2013. In particular, we estimate a regime-switching copula model to show the existence of a common macroeconomic regime that affects the distributions of both MNHF returns and the market index; this, in turn, creates a non-linear dependence, which can be confounded with tail dependence. Moreover, we provide evidence of positive (negative) linear correlation between the market index and MNHFs during bull (bear) periods that coincide with the US business cycle. We show with simulated data from our model that sample tail-based tests do not reject the tail dependence hypothesis, even if the tail dependence parameter is set to zero.

Keywords: Hedge funds, market neutrality, regime-switching models, copula, tail dependence.

JEL: G11, G23.

*We would like to thank Dante Amengual, Ramiro Losada, Javier Mencía, Guillermo Pérez-Tellechea, Rafael Repullo, Enrique Sentana, Javier Suarez, seminar audiences at CEMFI, and participants at the 22nd Finance Forum (Zaragoza) and the 39th Simposio de la Asociación Española Economía (Palma de Mallorca) for helpful comments and suggestions. Financial support from the Santander Research Chair at CEMFI is gratefully acknowledged. All remaining errors and omissions are our own responsibility.
Postal Address: Centro de Estudios Monetarios y Financieros (CEMFI), Casado del Alisal 5, Madrid, 28014, Spain.

1 Introduction

It is widely thought of that some hedge funds offer an advantage over other investment vehicles in that they are immune from market fluctuations, which makes them attractive to individual and institutional investors, especially during uncertain times. Indeed, numerous empirical studies have found low correlations between hedge fund returns and market returns (see for example, Agarwal and Naik (2004) and Fung and Hsieh (1999)). This characteristic has propelled growth in an industry whose size was estimated to have swollen from $100 billion in 1997 to $2.63 trillion in 2013 in the United States alone.[1]

Knowing the dependence that exists between hedge funds and asset markets is particularly important, as crashes in this large but seemingly opaque industry might lead to potentially devastating effects in financial markets, given the leveraged positions hedge fund managers take. Indeed, policymakers and academics have implicated hedge funds as having had a role in the 1992 crisis that led to major exchange realignments in the European Monetary System, the 1994 crisis in bond markets, and the 1997 East Asian financial crisis. The best known event involving this industry was the near-collapse of Long Term Capital Management (LTCM) in 1998, which was precipitated by events following the financial crises in East Asia and Russia.[2]

The role of hedge funds in the Great Recession is still not well understood, however. On the one hand, as evidenced by reports such as the Turner Review in the United Kingdom and the High-Level Group on Financial Supervision in the European Union, policymakers seem to have exonerated hedge funds as having had no influence on the financial instability that occurred during this period. For instance, as was stressed by the Turner Review, hedge fund leverage is typically below that of banks. Moreover, hedge funds did not perform a maturity transformation function equivalent to that of commercial and investment banks in the run-up to the Great Recession. Policymakers acknowledge, however, that hedge fund activity in aggregate can have a procyclical systemic impact, and as such, macroprudential regulation is required.

[1]Source: "ETFs/ETPs Grew At A Faster Rate Than Hedge Funds In 2013", nasdaq.com, February 18, 2014.

[2]LTCM mainly took advantage of fixed income arbitrage deals with US, Japanese and European government bonds. Due to the East Asian and Russian crises, however, the value of these bonds diverged, as investors sold Japanese and European bonds and bought US Treasury bonds. This, together with losses from the divergence between Royal Dutch and Shell share prices, led to a substantial decrease in LTCM's equity from $4.72 billion at the beginning of 1998 to $400 million at the end of September 1998. See Lowenstein (2000) for a more comprehensive account of the fall of LTCM.

The academic literature, on the other hand, has found mixed evidence. Some empirical papers have found that hedge funds took action during the Great Recession to prevent losses from occurring. For example, Ben-David et al. (2012) document that hedge funds in the US got rid of their equity holdings during the 2007-2009 financial crisis. Cao et al. (2013), meanwhile, provide evidence that hedge funds were able to anticipate liquidity in the market and adjust their market exposure to hedge changes in aggregate market liquidity. Other papers, meanwhile, posit that hedge funds, rather than commercial or investment banks, may have been the most important transmitter of shocks in periods of financial crises. Adams et al. (2014) finds that spillovers from hedge funds to other financial markets increased in periods of financial distress. These conflicting findings, hence, underscore the need for a more thorough understanding of the relationship, and in particular, the dependence that exists between hedge funds and asset prices.

Hedge funds are usually classified by their investment styles. One such investment style is called market neutral hedge funds (MNHF), which refer to "funds that actively seek to avoid major risk factors, but take bets on relative price movements utilising strategies such as long-short equity, stock index arbitrage, convertible bond arbitrage, and fixed income arbitrage" (Fung and Hsieh (1999), p. 319). As Fung and Hsieh (2001) and Patton (2009) note, they are not only one of the largest, but also are among the fastest-growing investment styles in the industry. As such, recent empirical literature has investigated the "neutrality" of MNHFs to the market index, of which there are numerous definitions. The most prominent one, which is the focus of this empirical study, is the exposure of these funds to market tail risk. Loosely defined, tail risk is the probability that a hedge fund is exposed to extremely negative events; "tail neutrality", hence, implies that the probability of an extremely low return on the hedge fund is not affected by conditioning on the event that an extremely low return on the market is also observed (Patton (2009), p. 2510).[3]

While numerous studies have found that there is low correlation between MNHFs and the market index, there is no consensus on whether this particular style of hedge funds is exposed to tail risk or not. Brown and Spitzer (2006) propose a tail neutrality measure

[3]Patton (2009) examines four other neutrality concepts: "mean neutrality", "variance neutrality", "Value-at-Risk neutrality" and "complete neutrality", all of which we do not explicitly address in this paper.

which uses a simple binomial test for independence, and find that hedge funds exhibit tail dependence. They also compare their measure with results from logit regressions similar to those used by Boyson et al. (2010), and find that while both techniques are able to capture tail dependence, their tail neutrality measure performs considerably better. Patton (2009), meanwhile, proposes a test statistic using results from extreme value theory and finds that there is no tail dependence between MNHFs and the market index. He acknowledges, however, that the results he obtains might be due to data limitations, and hence, cannot be reconciled with Brown and Spitzer (2006). The analyses performed in the previous papers are developed in an essentially static framework. In contrast, Distaso et al. (2010) use hedge fund index data to model dependence using a time-varying copula, and find that there does not exist tail dependence between hedge fund index returns and market index returns. Finally, Kelly and Jiang (2012) utilise a time-varing tail risk measure and find that the average exposure to tail risk of MNHFs is negative, which they take as evidence of the sensitivity of hedge funds to tail risks.[4]

A common thread of most of these papers, however, is that the results hinge on the assumption that the joint distribution of hedge fund and asset market index returns is fairly static over time. This paper, in turn, departs from the literature by allowing for regime switches, both in the joint distribution of hedge fund and asset market indices, and their corresponding marginal distributions.[5] An advantage that regime-switching models offer over other models is the possibility of capturing business cycle events. While tail dependence measures are frequently used as a metric of financial stability, they have no clear economic interpretation on the importance of tail events. Hence, allowing for regime-switching dependence provides a more intuitive way of linking asset price movements with macroeceonomic events that are linked to the real economy.

To operationalise this, we model the joint distribution of MNHF returns and asset market returns, which we represent by the market index return, through a regime-switching Student-t copula. Meanwhile, we model the marginal distribution as an asymmetric Student-t distribution proposed by Galbraith and Zhu (2010). Previous empirical studies that have employed Markov-switching copula models include Rodriguez (2007),

[4]Patton and Ramadorai (2013) use a dynamic framework to analyse risk exposures of hedge funds to different asset classes. However, they do not explicitly study tail dependence.

[5]Regime-switching models have been extant in the empirical asset pricing literature, prominent examples of which include Perez-Quiros and Timmermann (2001), Ang and Bekaert (2002) and Guidolin and Timmermann (2008).

who analyses contagion between stock markets in Asia during the 1997 financial crisis, and da Silva Filho et al. (2012), who study stock market dependence of the US, United Kingdom and Brazilian stock market indices. We generalise these studies by allowing the parameters of the marginal distributions to depend on the regime.

Our contribution to the literature is threefold. First, we find evidence that the shifts in the marginal distributions of MNHFs and the market are consistent with a model where both returns follow a common regime, creating a non-linear dependence that can neither be captured by models with smooth dynamics in the dependence nor by a model as the one proposed by da Silva Filho et al. (2012) where the copula parameters are regime-dependent, but the parameters of the marginal distributions are not.

Second, and most importantly, we find evidence of small, unconditional correlation between the MNHF and the market indices that is consistent with previous literature; however, we find that there exists an economically significant correlation conditioning on the regime. In particular, we find a negative and significant correlation in the bear regime (coinciding with NBER recession periods), which is consistent with the MNHF cutting positions in response to the market declines as found in Patton and Ramadorai (2013) and Ben-David et al. (2012), and a positive and significant correlation in the bull regime (coinciding with NBER expansion periods).

Third, we reconcile previous evidence on tail dependence by showing that tests that reject the null of no unconditional tail dependence using MNHF data (e.g., Brown and Spitzer (2006)) also reject the null hypothesis with simulated data from our model without tail dependence. Intuitively, as these tests are based on sample tails, if the sample is not extremely large they cannot differentiate between the non-linear dependence due to regime switches and tail dependence. Nevertheless, when the tests are based on extrapolation methods, such as copula methods or extreme value theory, they do not reject the null hypothesis.

The rest of the paper is organised as follows. Section 2 reviews the definition of tail dependence. Section 3 discusses the data used in the estimations. The modelling approach used in this paper, and the estimation results, are discussed in Section 4. Section 5 reconciles our results with Brown and Spitzer (2006)'s findings of tail dependence between MNHFs and asset markets. Section 6 concludes.

2 Tail dependence and copulas

Dependencies between (extreme) financial asset returns have gained increasing attention from academics and market practitioners in recent times, particularly after the global financial crisis of 2007-2009. In particular, a positive correlation between MNHF and market returns does not seem to be desirable, as negative events might weaken the stability of the financial system (given the size of the hedge fund industry). This imposes a cost to the government, as it needs to institute policies not only to restore financial stability, but also to prevent these events from occurring (Acharya et al. (2009)).

The concept of tail dependence has been used to measure extreme financial asset dependence, as it is closely related to contagion effects. More formally, given a bivariate random vector (X, Y) with marginal cumulative distribution functions (c.d.f.) F_X and F_Y, we say that X has lower tail dependence with respect to Y if and only if

$$\lambda_L \equiv \lim_{u \to 0^+} Prob(X < F_X^{-1}(u) | Y < F_Y^{-1}(u)) > 0,$$

where λ_L is defined as the lower tail dependence coefficient.[6] Clearly, tail neutrality (i.e., $\lambda_L = 0$) is heavily preferred by risk-averse investors over one of positive tail dependence ($\lambda_L > 0$), as underscored by Patton (2009); as opposed to the latter case, tail neutrality implies that the probability that both the market and the hedge fund will experience a negative return is zero.

One way of estimating tail dependence is through extrapolation methods, which are fully parametric procedures that require the entire data history to recover the tail dependence measure. A prominent example of these methods are copulas, which have an advantage in that they allow the characterisation of the dependence structure of the joint distribution of financial asset returns independently of their marginal distributions by construction. In particular, given a bivariate random vector (X, Y) with marginal c.d.f.'s (F_X, F_Y) and copula function $C(F_X(x), F_Y(y))$, the lower tail dependence is given by:

$$\lambda_L = \lim_{u \to 0^+} \frac{C(u, u)}{u}. \tag{1}$$

In this vein, we estimate the tail dependence between MNHFs and the asset market through a regime-switching copula, which uses the whole data set in order to identify

[6]The upper tail dependence coefficient is given by: $\lambda_U \equiv \lim_{u \to 1^-} Prob(X > F_X^{-1}(u) | Y > F_Y^{-1}(u))$.

the parameters of the copula function. Through the specific parametric form of the copula, we then take the limit and obtain the tail dependence measure.

We also consider another broad class of estimation procedures used in the literature, which we label as sample tail-based estimation. These techniques drop most of the sample and focus only on the relationship between the $p\%$ lowest values in the sample. An example of this type of technique are the logit regressions considered in Section 5. Specifically, we test how accurate the inference with sample tail-based techniques is when the true data generating process has some non-linear dependence different from tail dependence, and in particular, a Markov-switching common factor.

3 Data and summary statistics

The dataset we use for the study is collected from three sources. We obtain daily hedge fund return data from the Hedge Fund Research (HFR) database from April 1, 2003 to November 8, 2013. We only consider hedge fund indices that are of the "equity market neutral" type, which is indicated in the substrategy type of the database. Although we have data for global indices, we cannot clearly identify from the database the type of the hedge fund; hence, we only work with data from the US. Moreover, the data might be susceptible to selection and "instant history" biases; that is, (i.) the sample of hedge funds used to construct the MNHF index might not correspond to the universe of MNHFs, and that (ii.) the index returns might not accurately reflect the true performance of the hedge funds in the sample as hedge funds might have "filled" the past history with returns that are presumed to be high. However, as noted in the introduction to the Hedge Fund Indices document by HFR, they perform a cluster and representation analysis to mitigate these biases. These biases, however, will only affect our conclusions if they are not independent from the market. The other prominent bias discussed in the literature is the survivorship bias, which Fung and Hsieh (2002) note as a natural bias which generally cannot be rectified. This bias results from the fact that the hedge funds used to calculate the MNHF index are the surviving ones, and does not include those who have either ceased their operations, or those who have stopped reporting their performance. An implication of this is that the results that we present might not be representative of hedge funds that were not included in the computation of the index. While we are unable to test for survivorship bias (as we have index data), Liang (2000) notes that HFR data have fairly small and insignificant survivorship bias.

To represent the simple daily market index return, we obtain NYSE Composite data from Datastream.[7] We finally obtain data on recession dates from the NBER, which corresponds to December 2007 to June 2009 for the data in our sample. We use this data, however, for illustrative purposes.

Table 1 presents summary statistics of the hedge fund index and the NYSE Composite. We observe that for NBER expansion periods, both MNHFs and NYSE indices have positive expected return, have less dispersion, and fatter tails. During NBER recession periods, meanwhile, we find that the mean return is negative, have more dispersion, and thinner tails. We find, however, that the hedge fund index is negatively skewed during both expansion and recession periods; the market index, meanwhile, is positively skewed during recession periods, and negatively skewed during expansion periods. Unconditionally, however, we find that the hedge fund index returns have negative mean return and have distributions that are more dispersed with fatter tails, while the market index has positive mean return, distributions that are less dispersed with less fat tails.

Table 2, meanwhile presents dependence statistics between the hedge fund index and the market. We present three dependence statistics: the Pearson correlation coefficient, the Kendall-τ coefficient, and the left tail dependence statistic implied by a Clayton copula.[8] We find that during NBER recession periods, the market index and the hedge fund index exhibit negative correlation. During NBER expansion periods, however, we find that the market index and the hedge fund index exhibit positive correlation. The Pearson correlation coefficient and the Kendall-τ differ, however, when we take into

[7]We also tried representing the market index with the S&P 500 return and the Nasdaq return indices, since, as Fung and Hsieh (2002) note, it might be the case that the correlations between MNHFs and the market are sensitive to the index used to represent the market. The results we obtain, however, are similar.

[8]More formally (as discussed in Okimoto (2008)), the Pearson correlation coefficient can be computed from copulas through Spearman's ρ, considered as the linear correlation between $F_X(X)$ and $F_Y(Y)$:

$$\rho_S = 12 \int_0^1 \int_0^1 C(u,v)dudv - 3.$$

Kendall's-τ, meanwhile, is defined as the difference between the probability of two random corcondant pairs and the probability for two random discordant pairs for two iid vectors (X_1, Y_1) and (X_2, Y_2), which can be calculated from copulas through:

$$\tau_K = 4 \int_0^1 \int_0^1 C(u,v)dC(u,v) - 1.$$

account unconditional dependence. The Pearson correlation coefficient shows that there does not exist any dependence, while the Kendall-τ exhibits positive dependence. These statistics indicate that there does not exist unconditional correlation but some non-linear dependence between the MNHF and the Market that could explain a positive Kendall-τ. This fact motivates the use of a regime-switching model that could capture non-linear dependence. Interestingly, though, we do not observe any dependence from the left tail parameter of the Clayton copula.

4 Tail or two-state dependence?

The main objective of this paper is to analyse the dependence that exists between MNHFs and the market index. This is particularly important because the type of dependence that exists has different implications for risk management. On the one hand, the existence of tail dependence implies that hedge funds are sensitive to extreme left tail events. The existence of two-state dependence, on the other hand, implies that there is a persistent, common latent factor that drives the dependence between hedge funds and the market index; therefore, the occurrence of "extreme left" tail events become more predictable. In this section, we discuss the model specification for the empirical analysis we pursue, and the subsequent results.

4.1 Model specification

To analyse the dependence between MNHFs and the market index, we appeal to the model proposed by Rodriguez (2007) that extends the conditional copula model of Patton (2006) by introducing a hidden Markov chain to capture unobserved regime-switching. More formally, let $\{(x_{Ft}, x_{Mt})\}_{t=1}^{T}$, $t = 1, \ldots, T$ be the MNHF and market returns, respectively. To model the dependence between these two variables, we represent the joint distribution through a regime-switching copula as follows:

$$F(x_{Ft}, x_{Mt}|s_t, \mathcal{I}_{t-1}) = C_{\theta_{ct}}\left(F_F(x_{Ft}|s_t, \mathcal{I}_{t-1}), F_M(x_{Mt}|s_t, \mathcal{I}_{t-1})|s_t, \mathcal{I}_{t-1}\right), \qquad (2)$$

where $C_{\theta_{ct}}$ is the conditional copula with time-varying parameters θ_{c,s_t}, s_t is the state, $\mathcal{I}_{t-1}$ is the set of all possible information, and $F_i(x_{it}|s_t, \mathcal{I}_{t-1})$ $(i = M, F)$ are the marginal c.d.f's of x_{it}.

The marginal distributions are assumed to be:

$$x_{it} = \mu_{s_t} + \rho_{s_t}(x_{it-1} - \mu_{s_t}) + \sigma_{s_t}\varepsilon_{it}, \tag{3}$$

where $\varepsilon_{it} \sim f_{AST}(\varepsilon; \alpha_{s_t}, \nu_{1,s_t}, \nu_{2,s_t}\mu_{s_t}, \sigma_{s_t})$. We assume that ε_{it} follows an asymmetric Student-t distribution proposed by Galbraith and Zhu (2010), an extension of the two-piece method by Hansen (1994), that allows for an additional skewness parameter. This distributional assumption has two main advantages: First, it is extremely flexible and, second, it has a c.d.f. that can be efficiently computed which reduces the estimation time since we need to compute the quantile of each observation and time period. The density has the following form:

$$f_{AST}(x;\theta) = \begin{cases} \frac{\alpha}{\alpha^*}K(\nu_1)\left[1 + \frac{1}{\nu_1}(\frac{x}{2\alpha^*})^2\right]^{-\frac{\nu_1+1}{2}}, & x \leq 0 \\ \frac{1-\alpha}{1-\alpha^*}K(\nu_2)\left[1 + \frac{1}{\nu_2}(\frac{x}{2(1-\alpha^*)})^2\right]^{-\frac{\nu_2+1}{2}}, & x > 0 \end{cases}, \tag{4}$$

where $\theta = (\alpha, \nu_1, \nu_2)^T$, $\alpha \in (0,1)$ is the skewness parameter, $\nu_1 > 0$, $\nu_2 > 0$, are the left and right tail parameters respectively, $K(\nu) \equiv \Gamma((\nu+1)/2)/[\sqrt{\pi\nu}\Gamma(\nu/2)]$, and α^* is defined as:

$$\alpha^* = \alpha K(\nu_1)/[\alpha K(\nu_1) + (1-\alpha)K(\nu_2)]. \tag{5}$$

Denoting by μ the location parameter and σ the scale parameter, the general form of this density that we estimate is: $\frac{1}{\sigma}f_{AST}(\frac{x-\mu}{\sigma};\theta)$.

The copula model chosen is the Student-t copula, which has the following parameterisation:

$$\begin{aligned} C_{\theta_{ct}}(u_1, u_2 | s_t, \mathcal{I}_{t-1}) = &\int_{-\infty}^{\tau^{-1}_{\delta_{s_t}}(u_1)} \int_{-\infty}^{\tau^{-1}_{\delta_{s_t}}(u_2)} \frac{1}{2\pi\sqrt{1-\delta^2_{S_t}}} \\ &\times \left(1 + \frac{r^2 - 2\delta_{s_t}rs + s^2}{\eta^{-1}_{s_t}(1-\delta^2_{s_t})}\right)^{-\frac{\eta^{-1}_{s_t}+1}{2}} drds, \end{aligned} \tag{6}$$

where the copula parameters δ_{s_t} and $\eta^{-1}_{s_t}$ are the correlation (dependence) and degrees of freedom parameters respectively, which are allowed to change with the state. This parameterisation provides the following two advantages. First, it allows the series to be negatively correlated. Second, it allows for different degrees of tail dependence depending on the state. However, it has some limitations since by construction upper and lower tail dependence are equal and non-negative.[9]

[9]As can be observed in Figure 3 where the empirical joint density is plotted, these assumptions do not seem far fetched.

The hidden regime is modeled as a first-order Markov chain with two different states that correspond to a "bull" and a "bear" market:[10]

$$P(s_t = Bull\,|s_{t-1} = Bull, s_{t-2}, ..., s_0) = p_{11} \;\; \forall t$$

We estimate the parameters of the model via maximum likelihood estimation as in Hamilton (1989). Standard errors are obtained by inverting the Hessian matrix.

4.2 Results

Tables 3 and 4 present the estimation results for the marginal distributions. We focus first on Table 3, which shows results for the marginal distribution of the hedge fund indices. We find that during "bull" periods, the marginal distributions of hedge funds have a positive location and scale shift. The skewness parameter is also positive. During "bear" periods, hedge funds have negative location shift and positive scale shift parameters, and a positive skewness parameter, though it is less than that of the "bull" period. Meanwhile, from Table 4 it can be observed that the market index return exhibits positive location and scale shift, and skewness parameters both during "bull" and "bear" periods.

Table 5 presents the moments implied by the regime-switching model, which were obtained through simulations. We find that during "bull" periods, both MNHFs and market index exhibit a positive mean return, while during "bear" periods, they exhibit negative return, albeit small in magnitude. The marginal distributions are less dispersed in "bull" periods than in "bear" periods. This is consistent with empirical evidence on the countercyclicality of return volatilities found in previous studies (see for example, Brandt and Kang (2004)). More interestingly, the marginal distributions during "bear" periods exhibit fatter tails than those in "bull" periods, implying that tail returns were more likely to have occurred during the crisis. Moreover, marginal distributions of MNHFs exhibit positive skewness during "bull" periods and negative skewness in "bear" periods, implying that abnormal returns tend to be positive in the "bull" periods and negative in "bear" periods. This is not the case for stock returns, where the signs are opposite, which might be due to the non-linear investments that hedge funds usually carry out.

[10]We also considered three states (which we can interpret, as in Billio et al. (2009), as "bull", "bear" and "tranquil"), and we obtain similar results.

Table 6 presents results from the estimated conditional copulas. Focusing on the dependence parameter, we find that hedge funds and asset markets exhibit positive correlation during "bull" periods, and negative correlation during "bear" periods. The tail dependence parameter is not significantly different from zero, both during the "bull" and the "bear" periods.[11] Additionally, the model finds a high positive correlation in the "bull" periods and a low negative correlation in the "bear" periods. The positive linear correlation is an interesting feature of the data that is robust to different specifications, and it might be driven because of demand reasons (during bull times investors prefer higher returns at the cost of bearing a positive correlation with the market) and/or by supply reasons (during "bull" periods, it might be more costly or even impossible to hedge some market risks). The result of negative correlation between MNHFs and the market index during "bear" periods is also consistent with recent empirical results (e.g., Ben-David et al. (2012) and Patton and Ramadorai (2013)) which assert that hedge funds have cut their exposures to the market during unfavorable periods. This might be due to either one of two prominent hypotheses. The former, which is called in the theoretical literature as "limits-to-arbitrage" (see e.g., Brunnermeier and Pedersen (2009)), asserts that arbitrageurs (in our case, hedge funds managers) cannot benefit from mispricing in asset markets and, as a consequence, cannot monetise the illiquidity premium because investors have cut off their access to capital. The latter reason offered in the literature is that, to avoid losses from the market, hedge funds move their capital away from equity markets to alternative investment opportunities, which are possibly less liquid. However, as we do not have data on the portfolio holdings of hedge funds, we are not able to test the motives of MNHFs to reduce their exposures to the market.[12]

We then evaluate the economic significance of confounding tail dependence with two-state dependence from a risk management perspective. To assess this, we compute the Exposure CoVaR measure defined in Amengual et al. (2013), which is the α-VaR of the hedge fund index conditional on the market being at its α-VaR level.[13] Figure 1

[11]Note that the parameter under the null hypothesis of zero tail dependence is on the boundary of the parameter space (see Andrews (1999)). To further analyse statistical significance, we conduct a parametric bootstrap. The results indicate that our estimates correspond to the 0.41 and 0.43 percentile of the bootstrapped sample; therefore, we cannot reject the null hypothesis of no tail dependence. Moreover, the implied tail dependence is economically insignificant.

[12]Ben-David et al. (2012), meanwhile, provide evidence that is consistent with the "limits-to-arbitrage" hypothesis. However, they consider nine classes of hedge fund styles, MNHFs being one of them.

[13]More formally, $CoVaR_{\alpha}^{i|X^j \geq VaR_{\alpha}^j}$ is implicitly defined by the α-quantile of the conditional proba-

presents the results for CoVaR computed from the 90th to the 99th percentiles of the distribution of MNHF returns with simulated data from the regime-switching model. We calculate the CoVaR under each regime, and compare them with the unconditional CoVaR. The figure illustrates that the unconditional CoVaR is always in the middle of the risk measures conditional on each regime. An implication of this result is that if the hedge fund manager is able to identify which regime the economy is in, he will be able to compute more accurately risk measure, and in particular, Exposure CoVaR.

Figure 2 presents the smoothed probabilities of the regimes. We find that the smoothed probabilities closely approximate the actual occurrence of the "bull" and "bear" periods during the periods specified, including the European turmoil. Moreover, as in previous literature, we find that our "bull" and "bear" periods coincide with NBER-dated recession periods. An implication of this is that we can link the behavior of hedge funds as correlated with the recession and expansion periods in the US economy. That is, hedge fund managers reduce their exposures to the market during recession periods. Finally, Figure 3, presents the joint distribution of the data and our fitted model by their deciles. We observe that the model is able to capture most of the dependence structure; moreover, we find that fat tails resemble to tail dependence in the 10th and 90th deciles. Additionally, the negative correlation and the flexibility of our distribution are able to capture some of the negative dependence in the 10th and 90th deciles.[14]

The model in this paper is flexible enough to capture most of the features of returns (e.g., excess kurtosis, skewness, asymmetries); however, it comes at the cost of not having closed-form expressions. While not relevant for the results presented in this section, this hampers testing the null hypothesis of only one state. In particular, we cannot make use of the optimal test presented in Carrasco et al. (2014). Therefore, we test if the parameters are different between both NBER periods using an standard likelihood ratio test and we reject the null hypothesis at any level of significance.. Intuitively, if there exists one unique state there is no reason why the returns should be different in both periods, therefore, testing the null of one state is equivalent to testing that the MNHF returns and the market returns are different in NBER recessions and expansions.[15] We

bility distribution: $\Pr[X^i \geq CoVaR_\alpha^{i|X^j \geq VaR_\alpha^j} | X^j \geq VaR_\alpha^j] = \alpha$, X being the loss.

[14]We use negative dependence to refer to the fact that $Prob(x < q_{x,10\%} | y < q_{y,90\%}) > 0.1$, where $q_{z,p}$ is defined by $Prob(z < q) = p$.

[15]We conduct the test under the assumption that the NBER recession periods conditional on the state are independent of the market neutral hedge funds and the market. This might seem unreasonable at

interpret the result of this test, combined with the difference in magnitude of the coefficients between states and the consistency with previous literature, as evidence of the presence of two states.

In sum, the results suggest that two-state dependence, in contrast with tail dependence, is extremely important from a risk management perspective. As previous empirical papers show, hedge fund managers seem to consider macroeconomic regimes when they evaluate the risks faced by the fund.

5 Reconciliation with previous literature

The insight gained from the previous section is that hedge funds and asset returns do not exhibit significant tail dependence, which is in line with most of the previous literature. Instead, we find that hedge funds and asset markets have a non-linear dependence that is dependent on a common regime. In this section, meanwhile, we test if the methodology used in Brown and Spitzer (2006) to reject the hypothesis of no tail dependence is accurate when the true DGP corresponds to the regime-switching model we have earlier outlined. In their paper, they use monthly individual hedge fund data and illustrate the dependence with the market using rank-rank plots; similarly, we simulate a similar sample of 70 hedge funds with 37 months of data from our model, and construct the rank-rank plot. Figure 4 compares the rank-rank plot obtained from the model, with that obtained by Brown and Spitzer (2006). In these plots, each bin has a color that ranges from black, which implies that the bin is (nearly) empty, to white, which implies that this bin has most of the observations. As expected, the rank-rank plots look similar, with the model being able to replicate the observation that the lower tails contain most of the observations.

Brown and Spitzer (2006) also propose two tail neutrality tests based on the binomial test and a logit regression. We discuss each test in turn. The first test divides the distribution into 4 quadrants, which we illustrate in Figure 5:

1. LL for observations where the market and the hedge fund are both below the median;

first glance; however, we can observe that there are several market downturns which are not classified as NBER crisis. The results suggest that NBER periods are classified according to the "real" economy. Of course, the fact that MNHFs and the market might share a state that is related with the NBER provides the power of the test.

2. LW for observations where the fund is below the median, but the market is not;
3. WL for observations where the market is below the median, but the hedge fund is not; and
4. WW for observations where the neither the market nor the hedge fund are below the median.

Brown and Spitzer (2006) calculate the standard odds ratio, which is provided by the following formula:

$$Odds = \frac{n(LL) \times n(WW)}{n(LW) \times n(WL)}, \tag{7}$$

where $n(\cdot)$ signifies the number of hedge fund and asset market pairs that are in each quadrant. As implied by the formula, there exists tail dependence if the ratio is greater than 1, which implies that the LL and WW quadrants have more observations than the LW and WL quadrants. This test, however, assumes a centred and stable distribution over the whole sample which it is not consistent with the empirical results presented in the previous sections. To test if the non-linear dependence implied by the regime-switching model is accepted by the binomial test as evidence of tail dependence, we simulate 10,000 samples of a bivariate time series with a length of 2,600 days (similar to our sample and Brown and Spitzer (2006) sample) from our estimated model without tail dependence, where we change the tail dependence parameter to zero. For each time series, we calculate the odds ratio, and compare the resulting statistic with the p-value implied by independence, and a bivariate Normal distribution. We also calculate results from a chi-squared test of independence, which Brown and Spitzer (2006) note as a stronger test of tail dependence. We finally calculate the proportion of simulated data that amounts to a rejection of the null hypothesis of tail dependence, and find that we are able to reject the null hypothesis of no tail dependence for all of the simulations performed. These results are robust to the sample size and the number of simulations. Note, however, that an implicit assumption of the tests proposed by Brown and Spitzer (2006) is that the joint distribution of hedge fund and market returns are "steady", while the returns we have simulated come from a regime-switching model.[16]

[16]In principle, if the model is stationary, then the tests conducted by Brown and Spitzer (2006) should work on the unconditional distribution; however, they have insufficient data.

The second test, which is based from Boyson et al. (2010), is a logit regression of a dummy that takes value 1 if the hedge fund return is below its p percentile, on the market return and the same dummy for the market index, that is,

$$\Pr(x_F < q_{F,p}) = \Lambda(\mu + \rho_1 x_M + \rho_0 \mathbb{1}\{x_M < q_{M,p}\}), \tag{8}$$

where x_F and x_M are the hedge fund and market returns respectively, $q_{Z,p}$ is defined by: $\Pr(x_Z < q_{Z,p}) = p$, $Z = \{M, F\}$; $\Lambda(\cdot)$ is the logit c.d.f and $\mathbb{1}\{\cdot\}$ is the indicator function. If there exists tail dependence the coefficient of the market dummy should be positive. The main issue in this test is the selection of the percentile p. In particular, Brown and Spitzer (2006) use 15%. We do several regressions using different values of p from 20% to 0.1%. We perform these regressions for each of the simulated time series, and calculate, for each tail cut-off, the proportion of rejection of the null hypothesis of no tail dependence. Similarly, Boyson et al. (2010) do not assume different regimes for the economy.

Figure 6 presents the results of the logit regressions. We find that when the tail cut-off is large, almost all of the simulations have significant coefficients, which Brown and Spitzer (2006) take as evidence of tail dependence. However, as the tail cut-off becomes smaller, we find that the rate by which we can reject the null hypothesis of no tail dependence becomes smaller. In fact, looking at the extreme tails, we find that the rejection rate whittles to 5% of the simulations.[17] This result suggests that indeed, there is no tail dependence between the market index and hedge fund returns. Note, however, that the results we obtain may be due to the model's parameters, which we calibrate to the data.

We consider this empirical evidence as a reconciliation between two opposing results in the literature. On one hand, studies using interpolation techniques like copula methods to compute tail dependence measures find there was no evidence of such; on the other hand, studies using the sample tails to make inference consistently reject the hypothesis of no tail dependence. With simulated data from our model, we show that indeed, if the hedge fund and the market follow a two-state Markov-switching model with no tail dependence, we are able to conclude that there exists tail dependence between the hedge fund and the market.

[17] Repeating the same exercise with a logit regression where we correct for linear dependence yields similar results.

6 Conclusions

A distinctive feature of hedge funds is its seemingly low correlation with the market. Previous empirical literature has investigated the "neutrality" of market-neutral hedge funds by looking at the dependence between their return with those of the market index; in particular, the literature has addressed the question of whether there exists tail dependence between hedge funds and financial market returns. Though the empirical literature has agreed on the low correlation between hedge funds and asset markets, there has been no consensus on the question of tail dependence. This paper revisits this question by employing a regime-switching copula, which recognises the well-known fact that the risks faced by hedge funds are non-linear, and hence, more complex than those faced by traditional asset classes (see for example, Agarwal and Naik (2004) and Chan et al. (2005)). We depart from the previous empirical literature that has used regime-switching copulas in other contexts, however, by not only allowing the copula, but also the marginal distributions of hedge fund and market indices, respectively, to depend on the regime.

Results from our estimation indicate that conditional on the state of the economy, MNHFs and market returns exhibit either positive or negative dependence. Moreover, the results suggest that there exists a (latent) regime-dependent factor that drives hedge fund returns. This common latent factor can be interpreted as a non-linear, highly persistent risk factor that is non-hedgeable, for example, liquidity risk. Moments and risk measures computed from simulations implied by the model also suggest that the presence of macroeconomic regimes yields a better assessment of the risks faced by hedge funds. Moreover, these results are consistent with the empirical evidence on the role of hedge funds during the recent financial crisis.

We then reconcile the results we obtain with previous literature by conducting the same tests developed to detect tail dependence with simulated data from the regime-switching model we have earlier proposed and find that, indeed, the model is able to generate the tail dependence observed in previous studies. This result suggests that by not taking into account regime-switching dependence, one might inaccurately conclude that hedge funds and asset markets indeed exhibit positive left tail dependence.

Knowledge about the type of dependence that pervades between MNHFs and market returns has important implications for risk management. For instance, the existence of two-state dependence implies that extremely negative events become more predictable. One can also interpret left tail evidence as suggestive evidence of moral hazard in the hedge fund industry; for example, hedge fund managers might invest in very risky portfolios, knowing that they are protected by limited liability.[18] In contrast, the existence of two-state dependence might be driven by incomplete markets; that is, there exists a non-hedgeable risk (e.g. liquidity risk as pointed out by Ben-David et al. (2012)) that impedes MNHF managers from creating a fully neutral portfolio. While the first economic reason will imply the convenience of introducing tighter regulation of the industry, the second might justify public liquidity support in periods where liquidity is scarce.[19]

An interesting avenue for future research is to investigate which systematic risk factors are important during "bull" and "bear" periods. While previous literature has found that hedge funds are indeed exposed to systematic risk factors (see Agarwal and Naik (2004) and Patton (2009) for examples), taking into account which factors are important in different regimes allows for a more precise estimation of risk modelling. Through this way, we might be able to understand more fully the mechanisms underlying the cyclicality of hedge fund returns.

[18] Gorton and Ordoñez (2012) state that hedge fund and other well-informed traders may have played a more aggressive role in triggering the Great Recession in that they might have taken advantage of their private information and engaged in "predatory trading", that is, they hoarded high-quality collateral and traded low-quality collateral.

[19] For example, the New York Federal Reserve staged a bailout of LTCM by striking out a restructuring deal with the creditors of the company, leading to a recapitalisation of the hedge fund (Haubrich (2007)).

References

Acharya, V., M. Richardson, et al. (2009). *Restoring financial stability: how to repair a failed system*, Volume 542. John Wiley & Sons.

Adams, Z., R. Füss, and R. Gropp (2014, 5). Spillover effects among financial institutions: A state-dependent sensitivity value-at-risk approach. *Journal of Financial and Quantitative Analysis FirstView*, 1–48.

Agarwal, V. and N. Y. Naik (2004). Risks and portfolio decisions involving hedge funds. *Review of Financial Studies 17*(1), 63–98.

Amengual, D., G. Fiorentini, and E. Sentana (2013). Sequential estimation of shape parameters in multivariate dynamic models. *Journal of Econometrics 177*(2), 233–249.

Andrews, D. W. K. (1999). Estimation when a parameter is on a boundary. *Econometrica 67*(6), pp. 1341–1383.

Ang, A. and G. Bekaert (2002). International asset allocation with regime shifts. *Review of Financial Studies 15*(4), 1137–1187.

Ben-David, I., F. Franzoni, and R. Moussawi (2012). Hedge fund stock trading in the financial crisis of 2007 - 2009. *Review of Financial Studies 25*(1), 1–54.

Billio, M., M. Getmansky, and L. Pelizzon (2009). Crises and hedge fund risk. *University CaFoscari of Venice, Department of Economics Research Paper Series* (10/08).

Boyson, N. M., C. W. Stahel, and R. M. Stulz (2010). Hedge fund contagion and liquidity shocks. *The Journal of Finance 65*(5), 1789–1816.

Brandt, M. W. and Q. Kang (2004). On the relationship between the conditional mean and volatility of stock returns: A latent var approach. *Journal of Financial Economics 72*(2), 217–257.

Brown, S. J. and J. F. Spitzer (2006). Caught by the tail: Tail risk neutrality and hedge fund returns. *WP*.

Brunnermeier, M. K. and L. H. Pedersen (2009). Market liquidity and funding liquidity. *Review of Financial Studies 22*(6), 2201–2238.

Cao, C., Y. Chen, B. Liang, and A. W. Lo (2013). Can hedge funds time market liquidity? *Journal of Financial Economics 109*(2), 493 – 516.

Carrasco, M., L. Hu, and W. Ploberger (2014). Optimal test for markov switching parameters. *Econometrica 82*(2), 765–784.

Chan, N., M. Getmansky, S. M. Haas, and A. W. Lo (2005, March). Systemic risk and hedge funds. Working Paper 11200, National Bureau of Economic Research.

da Silva Filho, O. C., F. A. Ziegelmann, and M. J. Dueker (2012). Modeling dependence dynamics through copulas with regime switching. *Insurance: Mathematics and Economics 50*(3), 346 – 356.

Distaso, W., M. Fernandes, and F. Zikes (2010). Tailing tail risk in the hedge fund industry. *WP*.

Fung, W. and D. A. Hsieh (1999). A primer on hedge funds. *Journal of Empirical Finance 6*(3), 309 – 331.

Fung, W. and D. A. Hsieh (2001). The risk in hedge fund strategies: theory and evidence from trend followers. *Review of Financial Studies 14*(2), 313–341.

Fung, W. and D. A. Hsieh (2002). Hedge-fund benchmarks: Information content and biases. *Financial Analysts Journal*, 22–34.

Galbraith, J. W. and D. Zhu (2010, August). A generalized asymmetric Student-t distribution with application to financial econometrics. *Journal of Econometrics 157*(2), 297–305.

Gorton, G. B. and G. Ordoñez (2012). Collateral crises. Technical report, National Bureau of Economic Research.

Guidolin, M. and A. Timmermann (2008). International asset allocation under regime switching, skew, and kurtosis preferences. *Review of Financial Studies 21*(2), 889–935.

Hamilton, J. D. (1989). A new approach to the economic analysis of nonstationary time series and the business cycle. *Econometrica*, 357–384.

Hansen, B. E. (1994, August). Autoregressive Conditional Density Estimation. *International Economic Review 35*(3), 705–30.

Haubrich, J. G. (2007). Some lessons on the rescue of Long-Term Capital Management. *Policy Discussion Papers* (Apr).

Kelly, B. T. and H. Jiang (2012, November). Tail risk and hedge fund returns. *Chicago Booth Research Paper No. 12-44; Fama-Miller Working Paper.*.

Liang, B. (2000). Hedge funds: The living and the dead. *Journal of Financial and Quantitative Analysis 35*(03), 309–326.

Lowenstein, R. (2000). *When genius failed: The rise and fall of Long-Term Capital Management*. Random House LLC.

Okimoto, T. (2008, 9). New evidence of asymmetric dependence structures in international equity markets. *Journal of Financial and Quantitative Analysis 43*, 787–815.

Patton, A. J. (2006). Modelling asymmetric exchange rate dependence. *International Economic Review 47*(2), 527–556.

Patton, A. J. (2009). Are market neutral hedge funds really market neutral? *Review of Financial Studies 22*(7), 2495–2530.

Patton, A. J. and T. Ramadorai (2013). On the high-frequency dynamics of hedge fund risk exposures. *The Journal of Finance 68*(2), 597–635.

Perez-Quiros, G. and A. Timmermann (2001). Business cycle asymmetries in stock returns: Evidence from higher order moments and conditional densities. *Journal of Econometrics 103*(1), 259–306.

Rodriguez, J. C. (2007). Measuring financial contagion: A copula approach. *Journal of Empirical Finance 14*(3), 401 – 423.

7 Tables and Figures

Table 1: Summary statistics of hedge funds and the NYSE

	Unconditional	NBER Recessions	NBER Expansions
		Market Neutral Fund	
Mean	-0.002	-0.012	0.000
Std	0.261	0.372	0.238
Skewness	-0.107	-0.198	-0.011
Kurtosis	20.135	4.326	29.644
		NYSE	
Mean	0.000	-0.001	0.001
Std	0.013	0.026	0.010
Skewness	-0.177	0.082	-0.351
Kurtosis	13.868	6.128	6.812

Note: The table describes summary statistics of market neutral hedge fund index returns and the NYSE Composite index return. These are observed on a daily frequency from April 2003 to November 2013.

Table 2: Dependence statistics between the market and the hedge fund index

	Unconditional	NBER Recessions	NBER Expansions
Pearson	0.0147	-0.2344	0.1811
	(0.4481)	(0.0000)	(0.0000)
Kendall	0.0501	-0.1558	0.1015
	(0.0001)	(0.0000)	(0.0000)
Left Tail Dep.	0.0000	0.0079	0.0000
(Clayton Copula)	(1.0000)	(0.9845)	(1.0000)

Note: The table describes dependence statistics between market neutral hedge fund (MNHF) index returns and the NYSE Composite index return. We compute three dependence statistics: the Pearson correlation coefficient, the Kendall-τ dependence statistic, and the left tail dependence parameter computed from a symmetrised Clayton-Joe copula. p-values are in parentheses.

Table 3: Estimates of the marginal distribution of the hedge fund index

	Bear		Bull	
α (skewness)	0.452	(0.037)	0.570	(0.026)
ν_1 (left tail)	3.598	(0.835)	24.654	(21.167)
ν_2 (right tail)	5.205	(1.357)	4.121	(0.790)
μ (location)	-0.067	(0.035)	0.041	(0.013)
ρ (autocorrelation)	0.063	(0.033)	0.004	(0.024)
σ (scale)	0.268	(0.012)	0.164	(0.005)

Note: The table presents the parameters of the marginal distribution of the hedge fund index returns, which is assumed to be an asymmetric student-t proposed by Galbraith and Zhu (2010). Standard errors are in parentheses, and were computed using the numerical Hessian. Estimates are from NYSE Composite market index data from April 2003 to November 2013.

Table 4: Estimates of the marginal distribution of the market index

	Bear		Bull	
α (skewness)	0.572	(0.035)	0.552	(0.026)
ν_1 (left tail)	3.964	(0.933)	6.921	(2.062)
ν_2 (right tail)	3.096	(0.689)	5.704	(1.459)
μ (location)	0.003	(0.002)	0.002	(0.001)
ρ (autocorrelation)	-0.092	(0.033)	-0.061	(0.022)
σ (scale)	0.015	(0.001)	0.006	(0.000)

Note: The table presents the parameters of the marginal distribution of market index returns, which is assumed to be an asymmetric student-t proposed by Galbraith and Zhu (2010). Standard errors are in parentheses, and were computed using the numerical Hessian. Estimates are from NYSE Composite market index data from April 2003 to November 2013.

Table 5: Moments implied by the model

	Unconditional	Bear	Bull
	(a.) Market Neutral Fund		
Mean	-0.001	-0.023	0.009
Std	0.261	0.368	0.194
Skewness	-0.282	-0.371	0.5444
Kurtosis	22.098	16.296	9.118
	(b.) NYSE Composite		
Mean	0.000	-0.001	0.001
Std	0.014	0.023	0.008
Skewness	0.3454	0.4128	-0.195
Kurtosis	97.409	47.414	5.544

Note: The table presents the moments implied by the estimated model. Moments are obtained through simulation.

Table 6: Estimates of the conditional copula parameters

	Bear		Bull	
p_{jj}	0.990	(0.004)	0.996	(0.002)
δ	-0.048	(0.042)	0.218	(0.026)
η	0.133	(0.046)	0.067	(0.030)
Tail Dep.	0.014	(0.000)	0.006	(0.000)
Kendall τ	-0.031	(0.001)	0.140	(0.000)

Note: The table presents the parameters of the conditional copula of hedge fund and market index returns, which is assumed to be an student-t copula proposed by da Silva Filho et al. (2012). Standard errors are in parentheses, and were computed using the numerical Hessian. Estimates are from HFR and NYSE Composite market index data from April 2003 to November 2013.

Figure 1: Exposure CoVaR

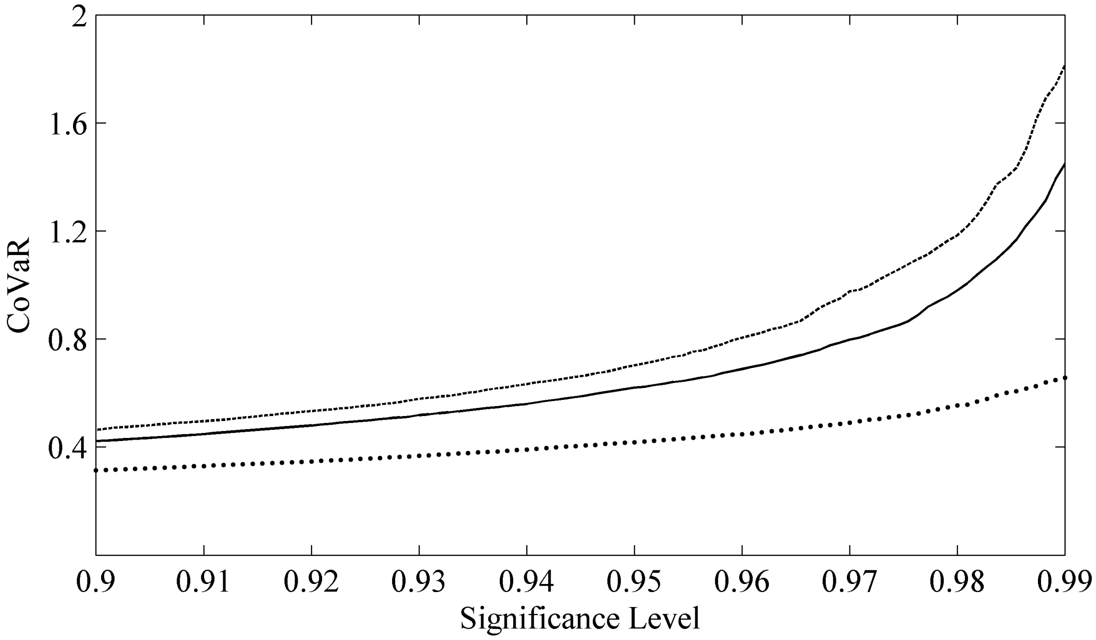

Note: This figure compares the Exposure CoVaR of the market neutral hedge fund index return (MNHF), where the event considered is the market being at its VaR level. The dark line corresponds to the unconditional CoVaR, while the dashed line and the dotted line correspond to the CoVaR conditional on the "bear" and "bull" regimes, respectively.

Figure 2: Smoothed Probabilities

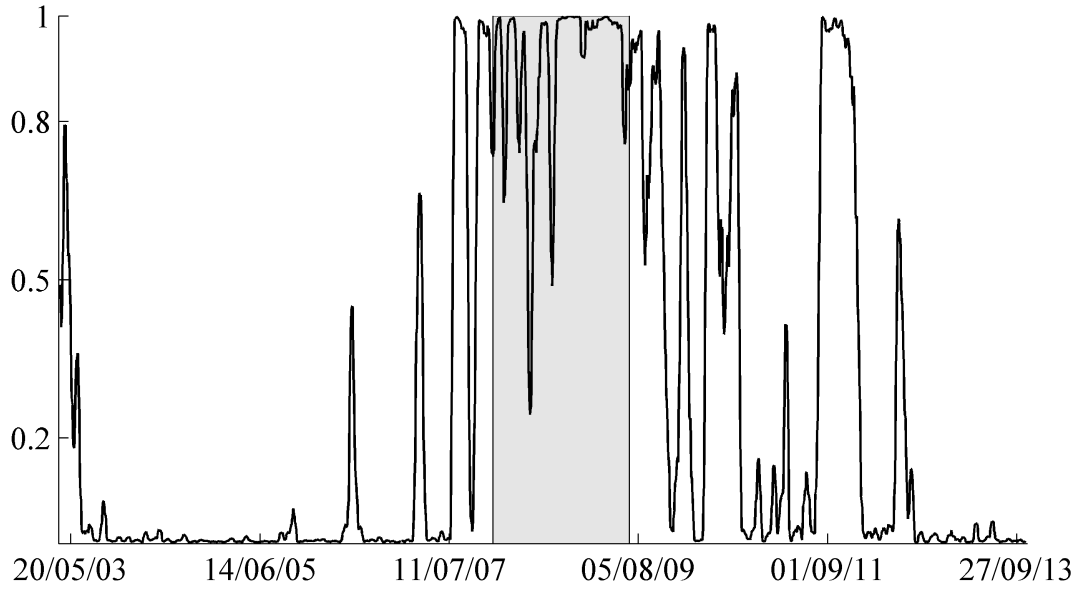

Note: The solid line presents the 10-days moving average smoothed probability of being in the state labeled as "bear". The region inside the vertical lines correspond to the NBER recession period.

Figure 3: Joint distribution

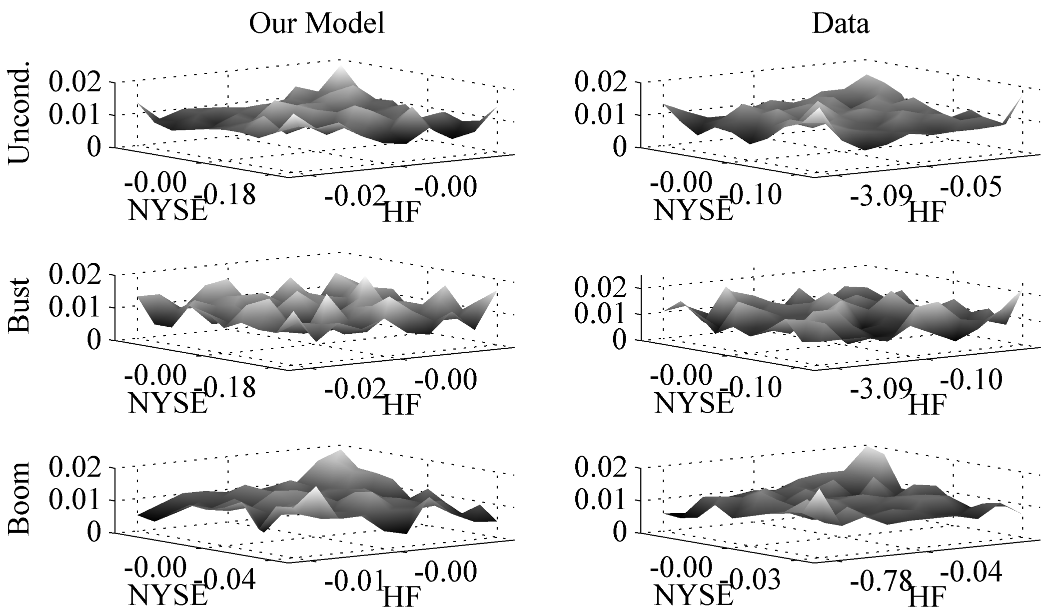

Note: This figure compares one simulation from our model with the empirical joint p.d.f from the data. We classify as "bull" ("bear") periods those days whose probability of being in the "bull" state is higher (lower) than 0.5.

Figure 4: Rank-Rank plot

Note: This figure compares the rank-rank plot for the hedge fund and market joint distribution computed with data from our simulated model with the rank-rank plot in Brown and Spitzer (2006). The rank-rank plot is constructed by dividing each of the marginals in deciles and cross tabulate the two distribution deciles, thus a white square in (0,0)-(0.1,0.1) means that hedge fund returns below the 10 % decile tend to coincide with market returns below the 10% decile.

Figure 5: Binomial test Illustration

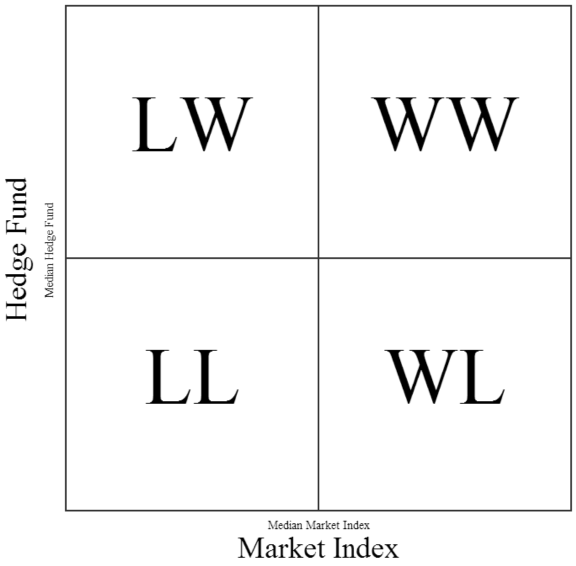

Note: This diagram presents the division of the joint distribution in order to compute the binomial test.

Figure 6: Logit test results

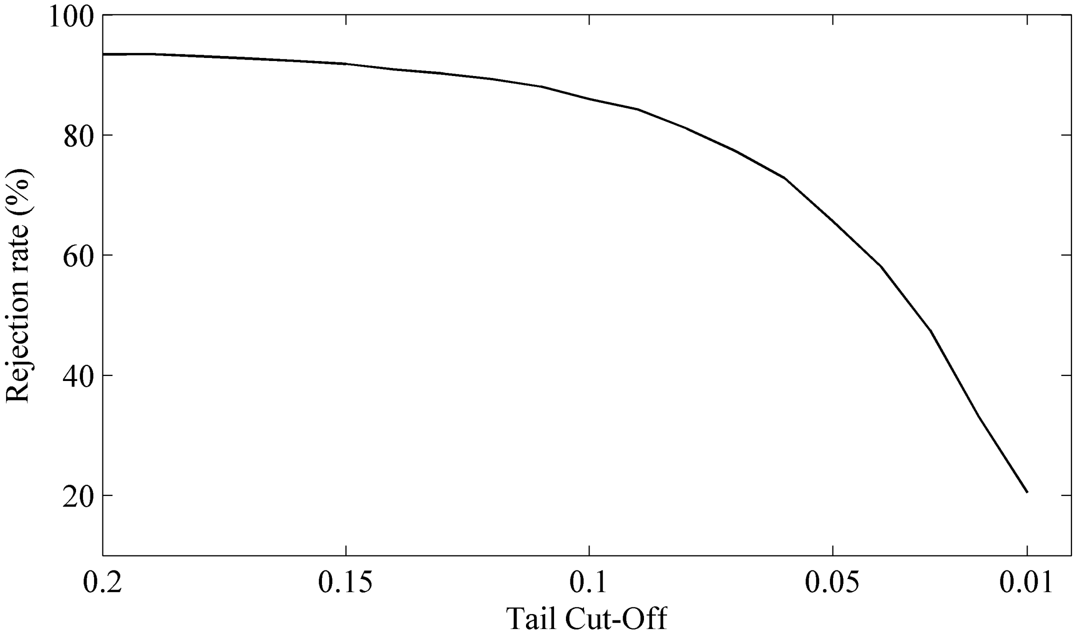

Note: This figure presents the actual size of the test with the null hypothesis $H_0: \ \rho_0 = 0$ in the following model: $Prob(x_F < q_{F,p}) = \Lambda(\mu + \rho_1 x_M + \rho_0 \mathbb{1}\{x_M < q_{M,p}\})$ where x_F and x_M are the hedge fund and market returns respectively, $q_{Z,p}$ is defined by: $Prob(x_Z < q_{Z,p}) = p \ Z = \{M, F\}$, $\Lambda(\cdot)$ is the logit c.d.f and $\mathbb{1}\{\cdot\}$ is the indicator function. In the vertical axis is presented the rejection rate and in the horizontal axis are presented the different values of p.