ASSET ALLOCATION WITH HEDGE FUNDS ON THE MENU

Phelim Boyle* and Sun Siang Liew†

ABSTRACT

Hedge funds have become an increasingly important asset class in recent years. This paper discusses the asset allocation decision of an investor who is considering investing in hedge funds. We develop a simple procedure that may help in making this decision when the assets consist of a core equity portfolio, the risk-free asset, and a hedge fund. A regime-switching framework is used to model the joint returns of the hedge fund and the equity market. We use monthly intervals so that the regimes can change only at most once a month. Within each regime the returns on the two risky assets are bivariate lognormal with constant parameters. These parameters are estimated from the empirical data. We show how to determine the optimal allocation of an investor, such as a pension plan, to hedge fund assets. Our procedure is based on the maximization of expected utility, and we use different horizons. We restrict the admissible strategies to buy-and-hold strategies, so we do not allow for portfolio balancing. We illustrate the procedure with examples. We find that bias in the hedge fund expected return has an important impact on the results. We note that some hedge fund strategies have substantial left-tail risk to which investors may be very averse. This type of risk aversion is not adequately captured by the standard expected utility model, but it could be added as a constraint.

1. INTRODUCTION

Hedge funds are now an important, if controversial, feature of financial markets. They have experienced phenomenal growth during the last few years. It is estimated that there are now over 8,000 hedge funds with total current assets of U.S. \$1.2 trillion. There is an ongoing active debate on the role of hedge funds in the financial sector. Hedge funds are very lightly regulated investment vehicles because the typical hedge fund investor is a high-net-worth individual or an institution, and these investors are presumed to be financially sophisticated.

Some observers have suggested that there should be increased regulation of hedge funds to protect investors and reduce the probability of systemic risk. Danielsson, Taylor, and Zigrand (2005) have analyzed the costs and benefits of hedge funds. They suggest that the main cost of hedge funds is their potential to increase systemic risk. On the positive side, they note that hedge funds contribute to economic efficiency by enhancing price discovery and providing additional diversification. Danielsson et al. suggest that some form of increased regulation is likely but argue that existing regulatory methods involving disclosure and activity restrictions are unsuitable for hedge funds.

Despite their growth, many institutional investors remain skeptical of the benefits of hedge fund investments in their portfolios. Kirschner, Mayer, and Kellser (2006) suggest four reasons for this situation:

- Hedge funds typically lack transparency. Investors do not have detailed information on the underlying assets in the fund's portfolios.

---

* Phelim Boyle, PhD, FIA, is a professor in the School of Business, Wilfrid Laurier University, Waterloo, Ontario N2L 3C5, Canada, pboyle@wlu.ca.

† Sun Siang Liew, MSc, was a graduate student in the Department of Statistics and Actuarial Science, University of Waterloo, Waterloo, Ontario N2L 3G1, Canada, when this project was started. Currently he is a catastrophe analyst with Validus Research, Inc., Accelerator Center, Suite 13, 295 Hagey Boulevard, Waterloo, Ontario N2L 6R5, Canada, sunsiang.liew@validusre.bm.

- Some hedge fund strategies are complex, and some use derivatives to increase their leverage and their returns. However, leverage increases risk as well.
- Hedge fund fees are generally significant, and there is often a lock-up period of three months or longer.
- There have been a number of dramatic hedge fund failures. These include Long Term Capital Management (1998), Amaranth (2006), and more recently in 2007 two mortgage-based hedge funds operated by the investment bank Bear Sterns.

However, in the last few years investors have experienced low interest rates and often disappointing equity returns. In this environment investments that promise relatively stable absolute returns have a strong appeal. Hedge fund assets under management have grown at about 20% per annum during the last decade (Kirschner, Mayer, and Kellser (2006)). The investor base of hedge funds has shifted during the last decade away from high-net-worth individuals to institutional investors. The relative share of hedge fund assets owned by individuals has declined from 66% in 1996 to 44% in 2004.[1] However, this share has still increased in dollar value because of the dramatic growth in hedge fund assets over this period. University endowments have increased their allocation to hedge funds in recent years and currently hold about 17% of their assets in hedge funds.

Corporate and public pension plans have adopted a more cautious and gradual move into hedge funds. As a percentage of total pension plan assets it is estimated that hedge fund investments represent about 1.6% (Kirschner, Mayer, and Kellser (2006)). However, the asset allocation to hedge funds is skewed because many plans do not hold any hedge fund investments. Pension plans that decide to invest in hedge funds tend to be the larger and presumably the more sophisticated plans.[2] The California Public Employees Retirement System (CalPERS) is the largest public pension plan in the United States with current total assets of some $250 billion. In June 2007 CalPERS approved new allocation ranges to potentially double its investments in hedge funds from its existing level of $5 billion (2% of total assets) to more than $10 billion (4% of total assets) (CalPERS 2007). In 2002 CalPERS invested $50 million in five hedge funds under its Risk Managed Absolute Return Strategies (RMARS) program, which as of June 2007 had almost $5 billion under management. That includes approximately $4 billion in 21 absolute return funds in eight strategies, and $1 billion with seven funds of funds in Asia, Europe, and emerging markets. The RMARS program has earned a 9.5% annualized return on investment compared with 7.4% return for the benchmark over the past five years. This was accomplished with approximately 4% volatility and a very low correlation to both stocks and bonds.

Many investors including pension plans are confronted with the decision as to whether or not to invest in hedge funds. The aim of this paper is to provide a simple framework and a practical procedure to assist in making this decision. Here is a brief summary of our approach. We assume that the investor already has an existing core equity portfolio in place and is considering investing in a hedge fund. We examine how an investment in hedge funds alters the asset allocation decision. For simplicity we assume that there are just three asset classes: the core equity portfolio, the hedge fund, and the riskless asset. We model the investment returns of the equity portfolio and the hedge funds using a regime-switching model. This model has been shown to provide a very good fit to equity returns; see, for example, Hardy (2003) and Guidolin and Timmermann (2005). Furthermore Chan et al. (2007) have shown that a regime-switching model is well suited to capture the dynamics of hedge fund returns. We use a multivariate regime-switching model to examine the dynamic relationships between the returns on the core portfolio and the hedge fund returns.

---

[1] In June 2005 the Financial Services Authority published a discussion paper (DP05/4) on hedge fund systemic risk.

[2] For example, as of December 2006 the Ontario Teachers Pension Plan had invested $14.5 billion in hedge funds. This amount represented some 14.5% of the total assets of the fund.

The procedure is illustrated by fitting the model to one of the datasets used by Chan et al.[3] As noted by Guidolin and Timmermann (2005) the existence of regimes has an important impact on the optimal asset allocation.

The use of the regime-switching approach offers some special advantages in the context of hedge fund returns. Several authors have shown that hedge fund returns are not normally distributed (e.g., Agarwal and Naik (2004); Chan et al. (2007)). In this paper we assume that the hedge fund returns are lognormal in each regime with fixed mean and variance. The regimes are determined by a hidden Markov process. The regime-switching model provides a flexible way to capture the higher moments of the empirical distribution of hedge fund returns. Furthermore it enables us to model the joint dynamics of the equity market and any specific hedge fund strategy. In particular the regime-switching approach captures the time variation in these dynamics. This provides a more realistic framework than the assumption of stationary independent and identically distributed returns.

The traditional approach to portfolio optimization is based on the single-period Markowitz approach. In this case the investor makes asset allocation decisions based on the mean and variance of the return distribution. This is a reasonable approach when asset returns are normal or approximately normal. Since hedge fund returns can have significant skewness and kurtosis the use of the standard mean variance model to make asset allocation decisions about hedge fund investments is questionable. To overcome this problem we assume that the investor has a power utility function. This means that the decision maker will care about the entire distribution of returns and not just the first two moments as in the standard mean variance approach. We use a pragmatic procedure to determine plausible parameters for this utility function. This calibration procedure provides a risk aversion level that is based on the investor's current equity allocations relative to bonds. Many institutional investors such as pension plans have guidelines on the fraction of their portfolios that can be invested in equities. For example, the range might be 40–80%. We show how to derive an implied level of risk aversion that is consistent with these types of limits. Based on this analysis we estimate that risk aversion coefficients ranging from 2 to 5 are plausible in this context.

We assume that the core equity portfolio is a broadly based equity market index such as the S&P 500. For the hedge fund data we used the Credit Suisse/Tremont hedge fund indexes (www.hedgeindex.com) from their inception (January 1994) to June 2006. There are a number of problems with hedge fund databases and we discuss them in the paper. The Credit Suisse/Tremont hedge fund database contains monthly data on 14 different hedge fund indexes. These indices correspond to different styles of hedge fund investing or combinations of these styles. We fitted a bivariate regime-switching model to the S&P 500 and each hedge fund index one at a time. We note here and we discuss later why the use of hedge fund index returns will underestimate the volatility of a typical individual hedge fund.

We note that the hedge fund return data used in this paper are net of fees and expenses. These fees can be quite significant[4] and are generally composed of two parts: a management fee and an incentive fee. The basic management fee is expressed as a percentage of the assets under management and is payable irrespective of the fund's performance. Normally these fees range from 1% to 2%. The incentive fees are payable only if the fund is profitable, and there is usually a high-water-mark restriction. The high-water-mark feature means that the hedge fund manager does not receive incentive fees unless the value of the fund exceeds the highest net asset value it has previously achieved.

We make two simplifying assumptions about the admissible investment strategies. We consider only buy-and-hold strategies. This assumption makes the analysis more tractable. We also assume that short sales are not allowed. The short

---

[3] It is well documented in the literature that there are several biases in most hedge fund datasets. These include the self-selection bias, the survivorship bias, the sample selection bias, and the backfill bias. We discuss these biases in the paper and suggest simple ways to minimize their impact.

[4] Based on a recent sample (July 2007) of 125 event-driven hedge funds the median management fee was 1.5%, the median incentive fee was 20%, and most of the funds had a high water mark.

sale restriction is operative for many institutional investors including pension plans and most mutual funds. It is not clear how an investor could short a hedge fund unless the hedge funds shares were traded.[5] We assume the investor has a universe of three assets: the core equity portfolio, the hedge fund, and the risk-free asset. As an application we use this model to determine the optimal allocation when the core equity portfolio is the S&P 500 Index, and the hedge fund is one of the 14 Credit Suisse/Tremont hedge fund indexes. To deal with various biases in the hedge fund database we use different approximations. We find the optimal allocation in each of the three assets for different values of investor risk aversion over different holding periods.

Our results show how the fraction allocated to the hedge fund depends on the investor's preferences, the time horizon, and hedge fund style. There is another interesting horizon effect in the regime-switching models that one does not see in stationary models. Suppose we have two regimes and just two assets: the equity portfolio and the risk-free security. Regime 1 occurs most frequently (say, 90% of the time), and, in this regime, equities have high expected returns and low volatility. Regime 2 is the less frequent one with negative expected returns and relatively high volatility. At inception, the economy can be in either regime 1 or regime 2. There is a dramatic difference in the optimal allocation for short horizons between the two regimes. Assuming the investor knows that the economy is in regime 1, the optimal allocation will favor equities. However, if the investor knows that the economy is in regime 2, the investor should invest in the risk-free security over a short horizon. In the case of a buy-and-hold strategy, as the horizon becomes longer even if the economy starts in regime 2, it becomes more likely it will spend an increasing fraction of time in regime 1 and so equities will start to become attractive again.

Our optimization procedure takes all the moments of the joint return distribution into account. If we use the parameters based on the historical distribution, the hedge fund becomes a significant component of the optimal portfolio, especially for higher levels of risk aversion. This result, that investors who are more risk averse should invest more in the hedge fund, is counterintuitive. This phenomenon has been observed before by Cvitanic et al. (2003). We suggest some possible reasons for this counterintuitive result relating to the significance of the left-tail risk and certain aspects of the data and model specification.

The layout of the rest of the paper is as follows. Section 2 describes the main types of hedge funds. We describe hedge fund return statistics and discuss some of the biases that are a feature of many hedge fund databases. We document the return characteristics of the Credit Suisse/Tremont hedge fund indexes. This is one of the datasets used by Chan et al. (2007), and we updated it to June 2006. We confirm that the return data are highly nonnormal.

In Section 3 we model the hedge fund returns using a regime-switching model. Initially we fit local univariate regime-switching models to the individual hedge fund time series of returns. Section 4 uses multivariate regime-switching models to capture the joint dynamics of the core equity portfolio and each hedge fund strategy one at a time. Section 5 explains our optimization procedure and discusses its implementation. We present our numerical results in Section 6 and discuss their implications. To simplify the presentation we discuss only one generic case in some detail. Section 7 concludes the paper, and the Appendix contains some technical equations.

2. HEDGE FUNDS

In this section we describe the common types of hedge fund styles and discuss certain features of hedge fund returns. Because of the nature of most hedge fund databases the returns suffer from certain biases. We analyze characteristics of a sample of hedge funds and confirm that the distributions of these returns exhibit significant skewness and sometimes excessive kurtosis.

Hedge fund managers provide their performance information on a voluntary basis. Some funds may cease reporting their results because they are doing poorly. These funds have returns that are below the average returns of other funds, and their omission creates an upwards bias. At

---

[5] Of course, hedge funds themselves can take short positions, and some are highly levered. In this paper we take the hedge fund return as an asset class.

the other extreme, there are certain funds that have become very successful. They may have grown to a size where they do not wish to attract new investors. These funds may also decide to leave the database for a very different reason. Presumably their performance will be superior to that of the average fund. While it is difficult to get accurate estimates of these two effects, many observers believe that the reported returns are biased upwards.

A second limitation of hedge fund databases is that typically they report data only on funds still in existence or that they are new and rapidly growing. Funds that are no longer active are dropped from the database entirely. This practice imparts an upward (or survivorship) bias to performance statistics, since funds that are closed are likely to have poor performance.

A third type of bias is the backfill bias (or instant history bias). This occurs when a fund first joins the database and is permitted to backfill its historical returns. This bias could be estimated by averaging the returns since inception and comparing them to the average returns since the fund joined the database.

There are different estimates of the magnitude of these biases in the literature. Liang (2000) estimated that the survivorship bias is about 2% per annum based on data for the period 1994–98. Fung and Hsieh (2000), using data for the same period, estimated the survivorship bias to be 3% per annum. They also estimated the backfill bias to be 1.4% per annum. Malkiel and Saha (2005), using data for the period 1996–2003, computed that the bias as measured by the difference between the expected returns of all hedge funds and those that survived is 3.79% per annum. De los Rios, Diez, and Garcia (2006) show that the extent of the bias varies with the strategy and the source of the data.

2.1 Hedge Fund Styles

Hedge funds are lightly regulated investment entities that employ a wide range of investment strategies. They have few investment restrictions, and they can take short positions, employ leverage, and use derivatives. There are many different types of hedge funds, each characterized by its specific style or strategy. Here are brief descriptions of the common ones:[6]

Convertible Arbitrage: Convertible arbitrage funds invest in convertible securities employing both single-security and portfolio-hedging strategies. The managers take a position in the convertible and hedge out the stock price risk.

Dedicated Short Bias: Short-biased managers invest mostly in short equity positions either directly or through derivatives.

Emerging Markets: These funds invest in emerging markets with less developed economies and aim to profit from market growth or economic conditions.

Equity Market Neutral: These funds use quantitative strategies to profit from pricing inefficiencies among securities while hedging away market risk.

Event Driven: Event-driven strategies aim to profit from corporate events related to particular companies. Subcategories of event-driven strategies include merger arbitrage, distressed securities, and corporate actions.

Merger Arbitrage: In this strategy the fund aims to profit from mergers and acquisitions by buying shares of the target and in some cases going short the bidder's shares.[7]

Distressed Securities: These funds exploit mispricing in the securities of corporations that are in financial distress.

Fixed-Income Arbitrage: This strategy exploits price anomalies among related fixed-income securities.

Global Macro: Global Macro is a macroeconomics-based strategy. It aims to profit from shifts in global economic conditions such as inflation, interest rates, currencies, and other macroeconomic factors.

Long/Short Equity: Long/short funds reduce risk by taking long positions in securities the manager thinks are undervalued and short positions in securities the manager believes to be overvalued.

Managed Futures: This strategy invests in financial and commodity futures markets and currency markets around the world.

---

[6] New hedge fund strategies are being invented all the time.

[7] If the merger involves an exchange of stock, then the fund can reduce market risk by shorting the common stock of the bidder.

Multi-Strategy: These funds allocate capital across different strategies depending on their perception of their relative risk reward profile.

Fund of Funds: Fund of funds invest in other hedge funds to obtain diversification benefits.

2.2 Hedge Fund Returns

For our empirical analysis we used the Credit Suisse/Tremont hedge fund indexes, from their inception (January 1994) to June 2006. The weight of each fund in an index is given by the relative size of its assets under management. To qualify for inclusion in the Credit Suisse/Tremont Indexes, a hedge fund must satisfy certain criteria:

1. A minimum of U.S. $50 million assets under management
2. A minimum one-year track record, and
3. Current audited financial statements that meet certain reporting requirements in terms of disclosure and transparency.

Funds are reselected on a quarterly basis as necessary, and the indexes are calculated and rebalanced monthly.

The pre-1999 hedge fund returns in this database suffer from both backfill and survivorship bias. Since then Credit Suisse/Tremont indicates that funds in the process of liquidation are not removed from the index and therefore captures all of the potential negative performance before a fund ceases to operate. In addition, Credit Suisse/Tremont states that fund managers contributing to the index are not allowed to backfill their historical returns. This suggests that the post-1999 data should be free of backfill and survivorship bias.

In Table 1 we show the summary statistics for the 150 monthly log-returns of the Credit Suisse/Tremont indexes from January 1994 to June 2006.[8] These numbers contain no correction for bias,[9] but they are net of fees and expenses. We observe a considerable degree of variation in the risk-and-return characteristics of the different hedge fund strategies. For comparison purposes we also include the corresponding statistics for the S&P 500 index. Here and throughout this paper the S&P 500 index returns have been adjusted to include dividend reinvestment since this makes for fairer comparisons.

In particular, the annualized mean return ranges from −0.26% for Dedicated Short to

Table 1

Summary Statistics for Monthly Log Returns of the Credit Suisse/Tremont Hedge Fund and S&P 500 Indexes for January 1994 to June 2006

Index	Sample Size	Annual Mean	Annual SD	Corr. with S&P 500	Min. Monthly	Med. Monthly	Max. Monthly	Skew	Excess Kurtosis	Max. Drawdown
CSFB Indexes										
Hedge Funds	150	10.58%	7.79%	48.72%	−7.55%	0.83%	8.53%	0.11	2.31	−14.86%
Convertible Arbitrage	150	8.65	4.71	13.97	−4.68	1.03	3.57	−1.32	3.14	−12.82
Dedicated Shortseller	150	−0.26	17.16	−75.67	−8.69	−0.39	22.71	0.82	2.07	−62.68
Emerging Markets	150	9.64	16.27	48.23	−23.03	1.38	16.42	−0.66	4.65	−60.03
Equity Market Neutral	150	9.69	2.93	37.25	−1.15	0.80	3.26	0.31	0.35	−3.60
Event Driven	150	11.18	5.62	56.08	−11.77	1.03	3.68	−3.39	24.54	−17.49
Distressed	150	12.94	6.35	54.86	−12.45	1.21	4.10	−2.89	19.28	−15.46
Event-Driven Multi-Strategy	150	10.27	6.04	48.84	−11.52	0.90	4.66	−2.51	16.75	−20.50
Risk Arbitrage	150	7.53	4.15	45.31	−6.15	0.61	3.81	−1.25	6.71	−7.91
Fixed Income Arbitrage	150	6.37	3.72	3.57	−6.96	0.73	2.05	−3.06	16.62	−13.33
Global Macro	150	13.50	10.95	23.81	−11.55	1.07	10.60	0.02	2.94	−31.18
Long/Short Equity	150	11.76	10.22	59.25	−11.43	0.82	13.01	0.22	3.85	−16.30
Managed Futures	150	6.82	11.98	−14.21	−9.35	0.22	9.95	0.04	0.37	−19.53
Multi-Strategy	147	9.28	4.33	9.86	−4.76	0.84	3.61	−1.19	3.32	−7.37
S&P 500	150	10.94	14.52	100.00	−14.46	1.28	9.78	−0.58	0.70	−59.30

---

[8] There are only 147 months for the Multi-Strategy index.

[9] We will introduce a bias correction when we use these data to implement our asset allocation model.

13.50% for Global Macro, and the annualized standard deviation ranges from 2.93% for Equity Market Neutral to 17.16% for Dedicated Short. However, many strategies have had quite low volatilities over this period. This would make them very attractive from a mean variance perspective. Figure 1 shows the mean-volatility plot for the 14 hedge funds and S&P 500 index. The plot seems to indicate that most hedge fund strategies offer superior risk-return characteristics over the S&P 500. There are at least two reasons why this conclusion is premature. The first is that the hedge fund returns are often biased. The second is that the higher moments beyond the second have not been taken into account.

From Table 1 we see that the skewness can be large and negative for some strategies such as Event Driven, Distressed, Risk Arbitrage, and Fixed Income Arbitrage. As might be expected the kurtosis for these strategies is also quite large. These features indicate a significant left-tail risk for some strategies. It seems clear that these returns are not normal, and we confirm this in the next subsection.

2.3 Normality Test

We use the Jarque-Bera statistic to test the normality of the series of hedge fund returns. The following table shows the results of Jarque-Bera test performed on the Credit Suisse/Tremont indexes. At the 5% significant level, the null hypothesis is rejected if the test statistic is greater than 5.99. The tests indicate that except for Equity Market Neutral and Managed Futures, all other funds' returns do not follow a normal distribution. This provides additional evidence against the log-normal distribution and indicates that this distribution is not adequate to model hedge fund returns.

Figure 1

Mean-Standard Deviation Plot of the Hedge Fund and S&P 500 Indexes

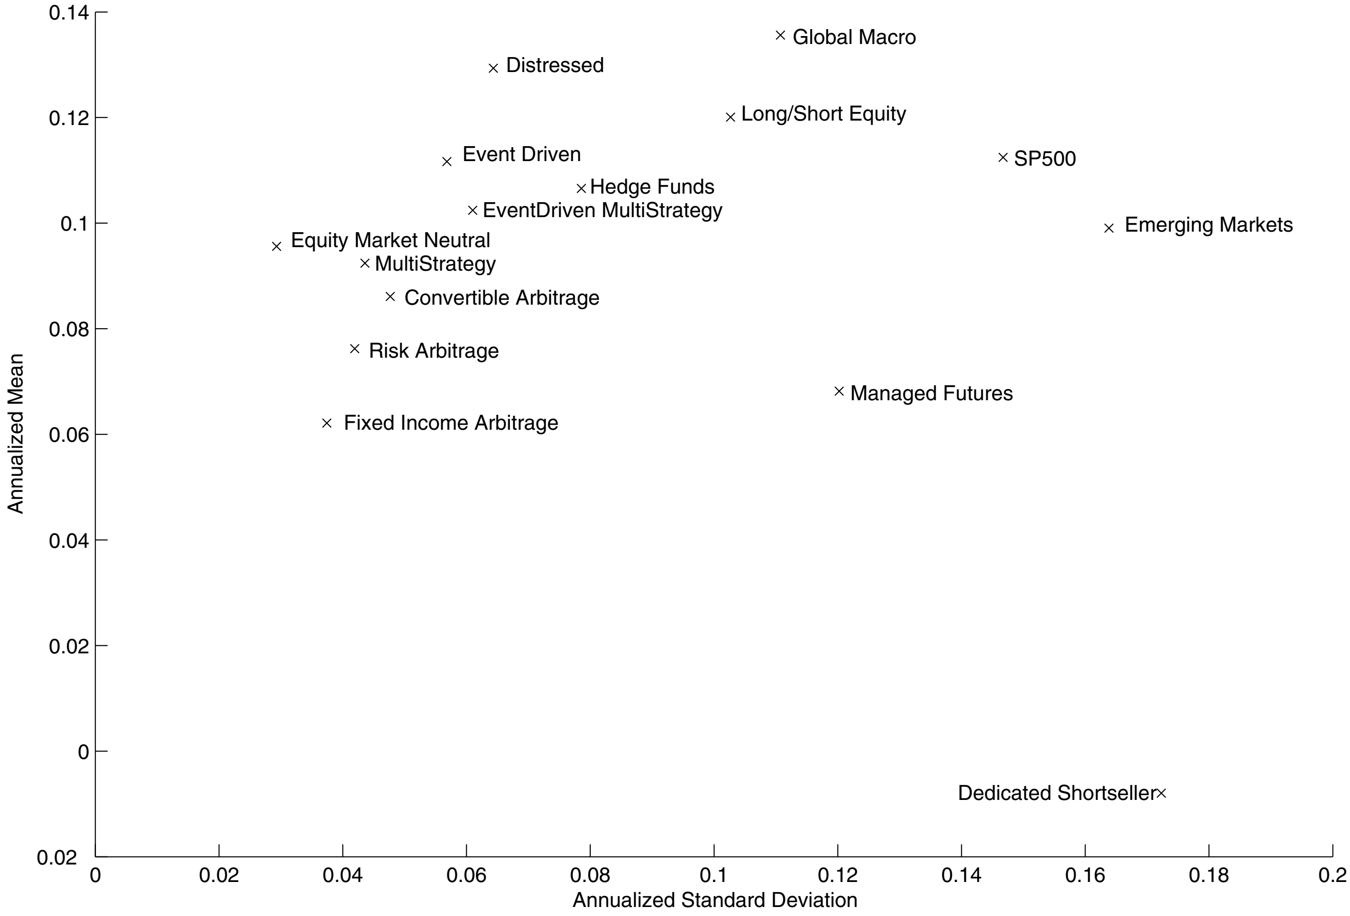

Table 2

Results of Jarque-Bera Test for Normality on CSFB Hedge Fund Indexes and S&P 500 Based on Monthly Data

Index	Sample Size	Skewness	Excess Kurtosis	Jarque-Bera	Normality Assumption
CSFB Indexes					
Hedge Funds	150	0.11	2.31	28.63	Rejected
Convertible Arbitrage	150	−1.32	3.14	95.50	Rejected
Dedicated Shortseller	150	0.82	2.07	38.82	Rejected
Emerging Markets	150	−0.66	4.65	129.38	Rejected
Equity Market Neutral	150	0.31	0.35	2.78	Not Rejected
Event Driven	150	−3.39	24.54	3,685.20	Rejected
Distressed	150	−2.89	19.28	2,301.37	Rejected
Event-Driven Multi-Strategy	150	−2.51	16.75	1,735.68	Rejected
Risk Arbitrage	150	−1.25	6.71	288.16	Rejected
Fixed Income Arbitrage	150	−3.06	16.62	1,785.02	Rejected
Global Macro	150	0.02	2.94	46.38	Rejected
Long/Short Equity	150	0.22	3.85	81.76	Rejected
Managed Futures	150	0.04	0.37	0.51	Not Rejected
Multi-Strategy	147	−1.19	3.32	92.27	Rejected
S&P 500	150	−0.58	0.70	10.42	Rejected

3. MODELS OF HEDGE FUND RETURNS

There is strong evidence that regime-switching models are able to capture the main features of the empirical distribution of equity returns. Relevant papers include Ang and Bakaert (2002a,b), David and Veronesi (2001), Guidolin and Timmermann (2005), Hardy (2003), Turner, Startz, and Nelson (1989), and Whitelaw (2000). For monthly stock market indexes, the regime-switching log-normal (RSLN) model produces good results. Recently Chan et al. (2007) and Billio, Getmansky, and Pelizzon (2006) have used regime-switching models for hedge fund returns. In this section we fit a parsimonious set of regime-switching models based on the joint distribution of the S&P 500 and the different hedge fund strategies.

We now describe a basic regime-switching model following the notation of Hardy. Assume that the asset price at time t is S_t and assume there are two regimes. If the economy is in regime s_t in the interval $[t, t + 1)$, the returns over this period have the following distribution:

$$\ln \frac{S_{t+1}}{S_t} |s_t \sim N(\mu_{s_t}, \sigma^2_{s_t}), \qquad s_t = 1, 2.$$

Note that within each regime the mean and variance are assumed to be constant. In this case we have six parameters:

$$\boldsymbol{\theta} = \{\mu_1, \mu_2, \sigma_1, \sigma_2, p_{1,2}, p_{2,1}\},$$

where $p_{i,j} = \text{Prob}[s_{t+1} = j|s_t = i]$ $(i, j = 1, 2)$, is the transition probability from regime i to regime j. These parameters can be estimated from the historical returns using maximum likelihood. For example, we fitted this model to the S&P 500 monthly returns for the period February 1956 to June 2006. The results are given in Table 3.

Note that under our regime-switching model, we assume that the return is lognormally distributed in each month. However, the parameters of the lognormal distribution are conditional on the current regime. As time passes the system moves between the regimes, and so the total return after a period of several months ends up being a mixture of lognormals. This tends to produce a distribution that does a better job of fitting the tails than the standard lognormal distribution.

For some applications it would be useful if we could tell (or assign a probability to) which regime we are in at a particular time. The estimation of the regime can be carried out using the MLE procedure. In calculating the likelihood at time t, we have four possible cases with different combinations of the regime at time $(t - 1)$ and the regime at time t:

Table 3

Maximum Likelihood Estimates of the Univariate RSLN Model for S&P 500, with the Corresponding Standard Errors in Parentheses

Index	Monthly μ_1	Monthly μ_2	Monthly σ_1	Monthly σ_2	Annual μ_1	Annual μ_2	Annual σ_1	Annual σ_2
S&P 500 (standard error)	0.0124 (0.0018)	−0.0049 (0.0160)	0.0328 (0.0021)	0.0608 (0.0068)	0.1592 (0.0243)	−0.0567 (0.1209)	0.1135 (0.0072)	0.2106 (0.0236)
Index	p_{11}	p_{12}	p_{21}	p_{22}	**Log-L**			
S&P 500 (standard error)	0.9568 (0.0339)	0.0432 (0.0339)	0.1411 (0.1152)	0.8589 (0.1152)	1088.95 (29.16)			

Figure 2

Monthly Returns of S&P 500 Together with Estimation of Each Regime

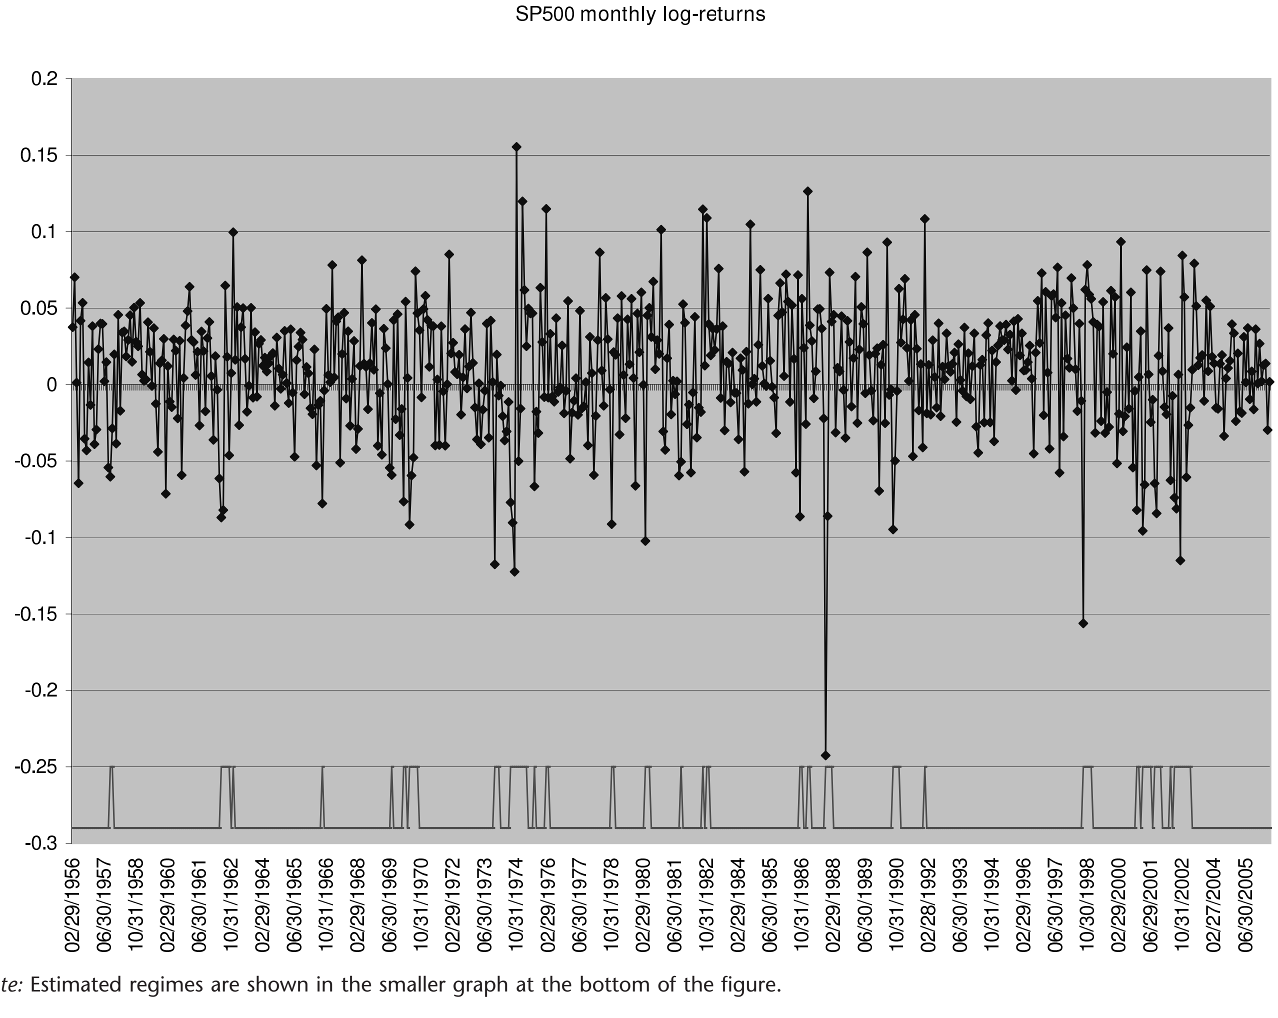

Note: Estimated regimes are shown in the smaller graph at the bottom of the figure.

$$A = f(s_t = 1, s_{t-1} = 1, y_t | y_{t-1}, \ldots, y_1, \boldsymbol{\theta}),$$

$$B = f(s_t = 1, s_{t-1} = 2, y_t | y_{t-1}, \ldots, y_1, \boldsymbol{\theta}),$$

$$C = f(s_t = 2, s_{t-1} = 1, y_t | y_{t-1}, \ldots, y_1, \boldsymbol{\theta}),$$

$$D = f(s_t = 2, s_{t-1} = 2, y_t | y_{t-1}, \ldots, y_1, \boldsymbol{\theta}).$$

We can then normalize to get the estimated/filtered probabilities as follows:

$$\text{Prob}(\text{Observation } y_t \text{ is in Regime } 1 | \boldsymbol{\theta}) = \frac{A + B}{A + B + C + D}, \quad (3.1)$$

$$\text{Prob}(\text{Observation } y_t \text{ is in Regime } 2 | \boldsymbol{\theta}) = \frac{C + D}{A + B + C + D}. \quad (3.2)$$

The larger of these probabilities can be used to indicate which state is more likely.

These probabilities (3.1) and (3.2) in turn imply the predicted probabilities for the next observation y_{t+1} with

$$\text{Prob}(y_{t+1} \text{ is in Regime } j | y_t; \boldsymbol{\theta}) = \sum_{i=1}^{2} p_{ij} \, \text{Prob}(y_{t+1} \text{ is in Regime } i | y_t; \boldsymbol{\theta}). \quad (3.3)$$

Figure 2 shows the monthly returns of the S&P 500 over the sample period together with the estimation of each regime.

This enables us to group the data according to their likely regimes. Figure 3 plots the distribution of returns within each regime. We see that these distributions (within each regime) appear to be approximately normal.

Figure 3

Density of S&P 500 Observations in the Two Different Regimes

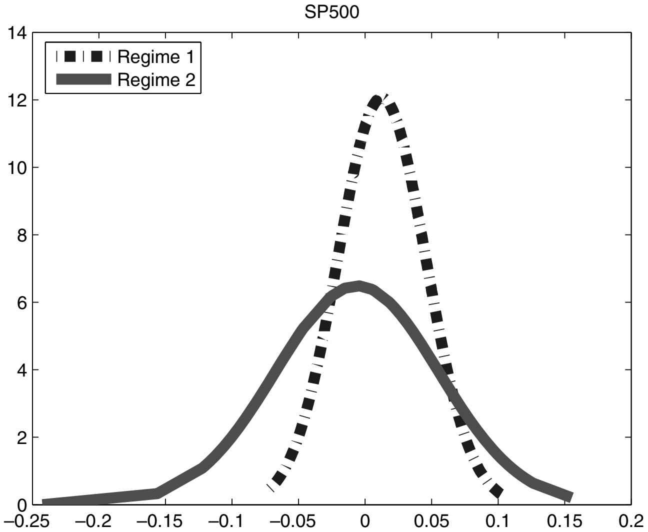

We fitted univariate regime-switching models to each of the hedge fund strategies discussed in the previous section. Chan et al. (2007) also fitted univariate regime-switching models to hedge fund returns for a slightly shorter dataset than ours,[10] and our results are broadly similar to theirs. We give our parameter estimates in Table 4. In the case of Dedicated Short Selling[11] and Managed Futures the maximum likelihood algorithm had trouble converging, indicating that these returns are not well fitted by a regime-switching model. For eight of the remaining strategies we see that the low-volatility state corresponds to the higher expected return. This is what we tend to find for long-only equity portfolios and corresponds to the situation in the case of the S&P 500. However, in four of the strategies corresponding to Hedge Funds, Equity Market Neutral, Global Macro, and Long Short the regime with the highest volatility also has the highest expected return.

In Figure 4 we show the distributions of each hedge fund index in each regime. The horizontal axis denotes the monthly return. We see that these distributions appear to be approximately normal. This is confirmed if we apply the Jarque-Bera test within each regime.

---

[10] Chan et al. (2007), Table 27. Their data period covered January 1994 until August 2004. Ours runs from January 1994 until June 2006.

[11] The pattern of returns for the Dedicated Short Selling Strategy has some unusual features. The average monthly return over the 150 months is almost zero (0.02%) with a standard deviation of 4.95% per month. However, there is a significant number of returns with large negative values of roughly the same magnitude. We can illustrate this feature as follows. The number of months with returns less than −5% is 29. The mean of these 29 returns is negative 6.56%, and their standard deviation of these 29 values is only 0.61%. This represents an unusual pattern of returns.

Table 4

Maximum Likelihood Estimates of the Univariate Regime-Switching Lognormal Model for the CSFB Hedge Fund Indexes

Index	p_{12}	p_{21}	Annual μ_1	Annual μ_2	Annual σ_1	Annual σ_2	Log-L
CSFB Indexes							
Hedge Funds	0.84%	0.72%	8.37%	12.64%	3.48%	9.84%	386.06
Convertible Arbitrage	11.09	17.03	16.76	−2.52	1.92	4.19	469.37
Dedicated Shortseller	76.25	11.26	−55.98	10.26	2.76	15.58	248.79
Emerging Markets	1.16	0.93	14.60	4.04	8.04	20.76	267.15
Equity Market Neutral	3.28	3.02	5.00	14.45	2.04	3.00	516.55
Event Driven	1.67	46.73	13.97	−39.66	3.84	15.00	445.74
Distressed	1.76	58.17	15.95	−46.10	4.56	16.68	421.24
Event-Driven Multi-Strategy	1.18	45.24	12.59	−46.44	4.56	15.96	426.21
Risk Arbitrage	7.61	27.58	8.99	3.10	2.64	7.20	468.83
Fixed Income Arbitrage	6.70	39.59	9.99	−12.23	1.92	6.12	513.62
Global Macro	0.78	0.68	13.60	13.65	3.36	14.04	354.66
Long/Short Equity	0.95	2.98	7.86	21.15	6.36	15.24	340.48
Managed Futures	67.49	17.07	−6.74	9.80	4.32	12.96	295.49
Multi-Strategy	2.56	24.40	11.33	−8.11	3.24	8.64	454.46

Note: The estimates in this table are computed using monthly returns from January 1994 to June 2006.

Figure 4

Densities of Hedge Fund Indexes Grouped according to the Different Regimes

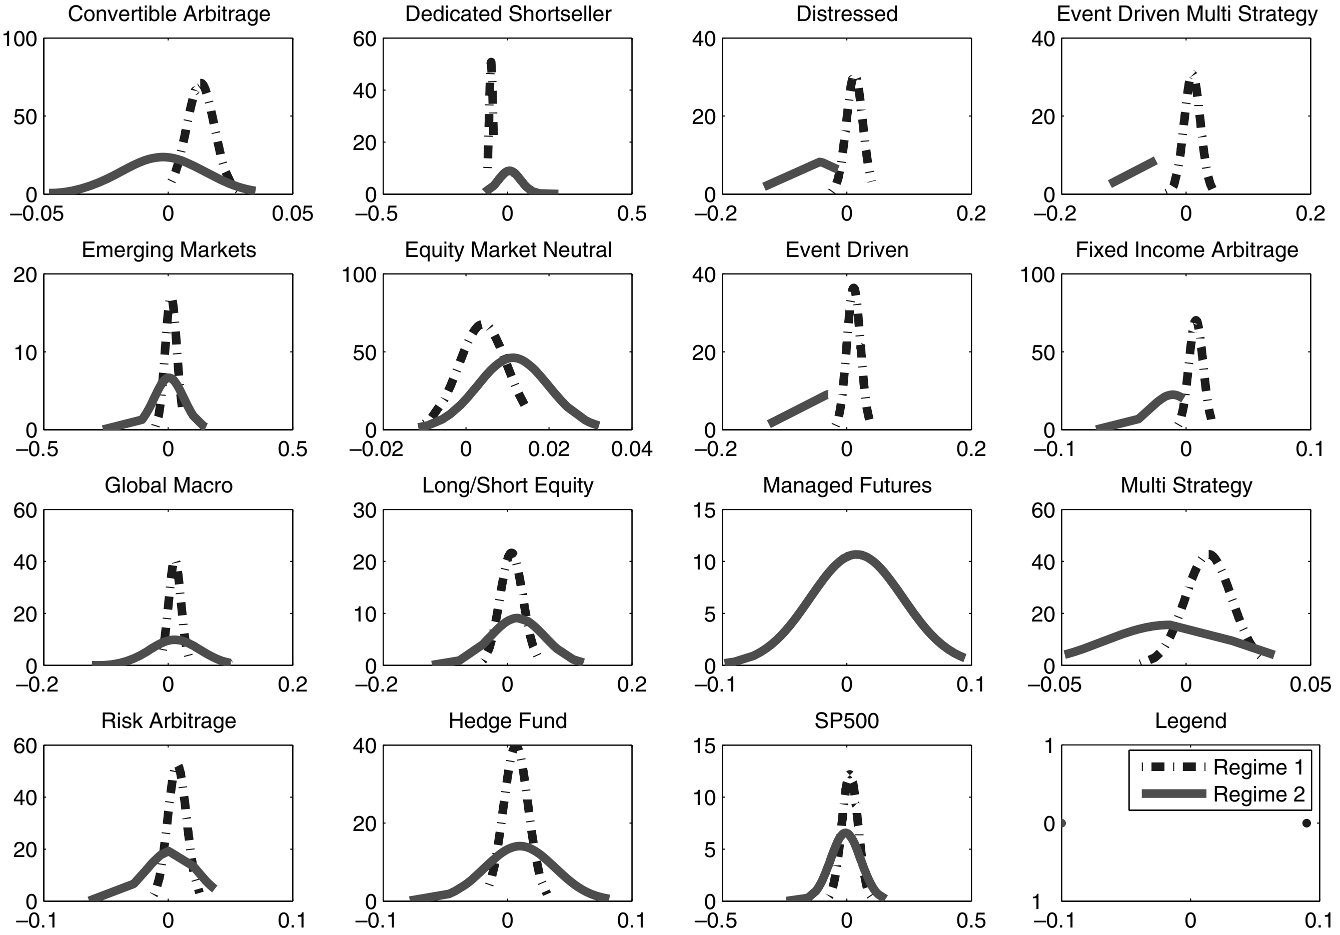

Note: The horizontal axis corresponds to the monthly returns.

Table 5

Application of Jarque-Bera Test within Regimes for Each Hedge Fund Index

Index	No. of Observations	Regime 1 Normality Assumption	No. of Observations	Regime 2 Normality Assumption
CSFB Indexes				
Hedge Funds	63	Not Rejected	87	Not Rejected
Convertible Arbitrage	94	Not Rejected	56	Not Rejected
Dedicated Shortseller	26	Not Rejected	124	Rejected
Emerging Markets	61	Rejected	89	Rejected
Equity Market Neutral	71	Not Rejected	79	Not Rejected
Event Driven	146	Not Rejected	4	Not Rejected
Distressed	146	Not Rejected	4	Not Rejected
Event-Driven Multi-Strategy	148	Not Rejected	2	Not Rejected
Risk Arbitrage	128	Not Rejected	22	Not Rejected
Fixed Income Arbitrage	134	Not Rejected	16	Rejected
Global Macro	62	Not Rejected	88	Not Rejected
Long/Short Equity	106	Not Rejected	44	Not Rejected
Managed Futures	0	N.A.	150	Not Rejected
Multi-Strategy	137	Not Rejected	10	Not Rejected

4. MULTIVARIATE REGIMES

In principle, the portfolio allocation decision will involve the joint returns of the S&P 500 and the different hedge fund strategies. Ideally we would like to model all the strategies jointly together with the core equity portfolio. This would require a very large number of parameters, and we do not have enough empirical data to estimate such a model. Hence we adopt a simpler pragmatic approach. We model each hedge fund strategy and the S&P 500 on a pairwise basis using a bivariate regime-switching model. We assume there is a global regime that applies to both the S&P 500 and the hedge fund in question.[12] This means that both the S&P 500 and the hedge fund are in the same regime at the same time. In this case we have 12 parameters to estimate:

$$\boldsymbol{\theta} = \{\mu_1^e,\ \mu_1^h,\ \mu_2^e,\ \mu_2^h,\ \sigma_1^e,\ \sigma_1^h,\ \sigma_2^e,\ \sigma_2^h,\ \rho_1,\ \rho_2,\ p_{1,2}\ p_{2,1}\},$$

where

Table 6

Maximum Likelihood Estimates of the Bivariate RSLN Model for the CSFB Hedge Fund and S&P 500 Indexes

Index	ρ_1	ρ_2	p_{12}	p_{21}	Annual μ_1	Annual μ_2	Annual σ_1	Annual σ_2
CSFB Indexes								
Hedge Funds	66.20%	42.82%	16.30%	38.86%	12.76%	5.14%	4.72%	12.17%
Convertible Arbitrage	15.65	−0.28	12.91	18.52	15.71	−0.49	1.91	6.08
Dedicated Shortseller	−74.69	−84.26	1.16	2.82	0.64	1.68	15.15	21.71
Emerging Markets	61.88	50.57	10.41	14.74	24.54	−15.41	8.93	23.04
Equity Market Neutral	45.40	50.98	0.96	2.48	8.71	11.43	3.19	2.45
Event Driven	56.22	43.59	3.88	29.16	14.97	−6.83	3.64	10.93
Distressed	52.10	46.70	4.24	28.49	16.86	−2.71	4.34	11.93
Event-Driven Multi-Strategy	51.22	38.74	6.81	34.82	14.12	−4.27	3.90	10.86
Risk Arbitrage	60.17	49.68	1.89	2.55	7.96	5.94	2.80	5.45
Fixed Income Arbitrage	14.80	−38.99	6.11	45.70	9.51	−14.53	2.12	6.56
Global Macro	0.13	40.71	0.93	0.80	13.27	7.55	3.37	14.52
Long/Short Equity	80.66	52.35	4.01	9.60	12.15	7.91	7.96	14.31
Managed Futures	29.49	−0.66	2.68	9.62	1.55	12.54	10.67	15.04
Multi-Strategy	24.28	−47.16	11.32	57.03	12.68	−8.29	3.00	6.68
S&P 500					15.92	−5.67	11.35	21.06

---

[12] Boudreault and Panneton (2006) also recommend the use of a global regime in their analysis of international equity returns.

Figure 5

Monthly Returns for the S&P 500 and the Event-Driven Strategy Shown Together with the Estimation of the Different Regimes

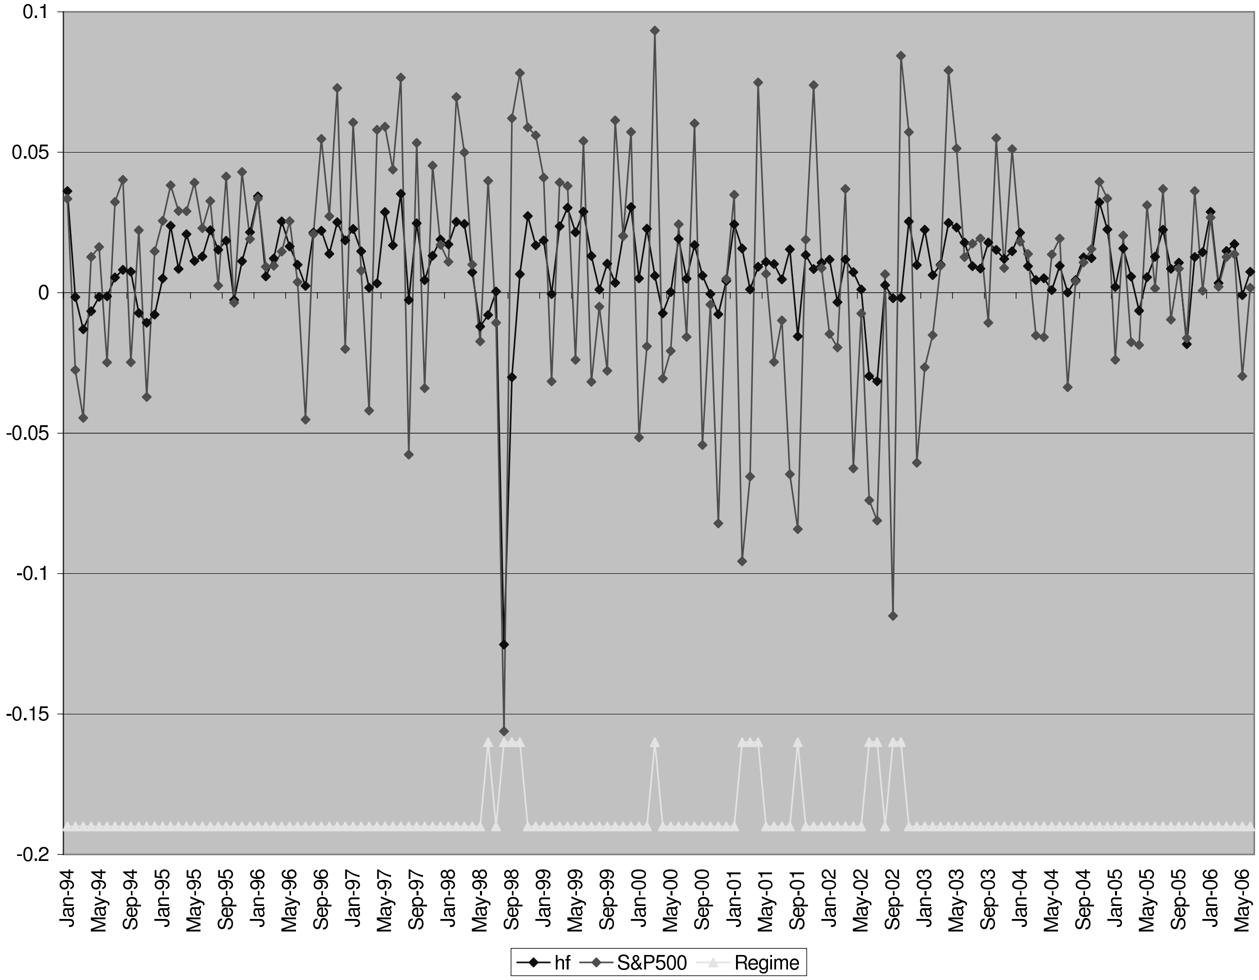

Note: The regimes are shown in the bottom part of the figure just above the horizontal axis.

$\mu_{s_t}^{(\cdot)}$ = mean of asset (·) in regime s_t,
$\sigma_{s_t}^{(\cdot)}$ = standard deviation of asset (·) in regime s_t,
ρ_{s_t} = correlation between the assets in regime s_t,
$p_{i,j}$ = transition probability.

Rather than allowing the assets to be in different regimes at the same time, the global regime gives us more degrees of freedom to estimate the parameters. We do not have enough data points to estimate more elaborate regime-switching structures. Hence we form a global regime between the S&P 500 and each hedge fund considering the funds one at a time. We assume that the regimes are common to the S&P 500 and the hedge fund, and we estimate the model from the historical data. To reduce the numbers of degrees of freedom we estimate the S&P 500 returns and variances in each regime from the univariate model and use these four estimates as known parameters in the estimation of the bivariate model for the hedge fund and the S&P 500. This means that we estimate the remaining eight parameters of the bivariate distribution from the two data series over the period January 1994 until June 2006.

Table 6 gives the estimates of the parameters of the bivariate regime-switching model when we consider each hedge fund and the S&P 500 on a pairwise basis. As in the univariate case we had problems with the convergence of the estimates in the case of Dedicated Short Selling strategy and the Managed Futures strategy. We note that the bivariate estimates in Table 6 differ from the univariate estimates in Table 4 as we would expect. Because of the constraints we imposed by fixing the S&P 500 returns there is an in-built consistency with respect to the S&P 500 returns across the different pairings. However, we allowed the data to select the regimes in each pairwise comparison. That is why the estimates for the regime transition probabilities vary across each pair.

We illustrate the joint returns of the S&P 500 and the Event-Driven Index in Figure 5. We also provide our estimation of the periods in each regime.

5. OPTIMIZATION

In this section we explain the optimization procedure. We use a very simple approach to model the asset allocation decision of a risk-averse investor. We first discuss the problem for general distributions and utility functions and later introduce our specific assumptions. We assume there are three available asset classes:

1. The core equity portfolio
2. The hedge fund
3. The risk-free asset.

We assume that the rates of return on these assets at the end of a fixed time horizon are r^e, r^h, and R, where the first two are random. Assume an investor has initial wealth W_0. The investor's end of period wealth will be

$$\begin{aligned} W &= W_0[x_1(1 + r^e) + x_2(1 + r^h) + x_3(1 + R)] \\ &= W_0[1 + R + x_1(r^e - R) + x_2(r^h - R)], \end{aligned} \tag{5.1}$$

since $x_1 + x_2 + x_3 = 1$. We assume there is no short selling so that

$$0 \leq x_1 \leq 1,\ 0 \leq x_2 \leq 1,\ 0 \leq x_3 \leq 1. \tag{5.2}$$

Assume the investor's preferences are represented by a strictly increasing, strictly concave function u. The investor's decision problem is to maximize

$$\mathrm{E}[u(W)] \tag{5.3}$$

subject to the constraints given by (5.2), where W is given by equation (5.1). The investor selects the investment proportions x_1, x_2, x_3 to maximize expected utility of end-of-period wealth. In view of the constraints on the $x's$ there are only two independent choice variables in this optimization. We will use x_1 and x_2 as the investor's decision variables.

In our case the equity portfolio and hedge fund are assumed to follow a bivariate regime-switching lognormal process. This process is characterized by two regimes. Recall that we fitted a global regime to these two asset classes in Section 4. At inception we can assume the process is in regime 1 or regime 2 or that there is a specified probability that it is in either regime. For purposes of exposition, we assume the return-generating process starts in regime 1.

We now discuss some useful connections between the joint distribution of the equity portfolio and the hedge fund at the end of T months and the sojourn times in each regime. If we start the process in regime 1, then there will be a distribution of possible regime visits at time T. During the entire period of T months the number of visits to regime 1 can range from 1 to T. Consider the set of paths that result in a total of exactly t months in regime 1 and hence $(T - t)$ months in regime 2 where $1 \leq t \leq T$. Conditional on these paths the returns have a bivariate lognormal distribution at the end of the horizon. Furthermore we can find the explicit expressions for the parameters of this distribution. It is convenient to work with the corresponding bivariate normal distribution (X, Y), where X corresponds to the core equity return and Y corresponds to the hedge fund return. The parameters of the bivariate normal distribution, conditional on the process spending a total of exactly t months in regime 1 and $(T - t)$ months in regime 2, are as follows:

$$\begin{aligned} \mu_X &= t\mu_1^e + (T - t)\mu_2^e, \\ \mu_Y &= t\mu_1^h + (T - t)\mu_2^h, \\ \sigma_X^2 &= t(\sigma_1^e)^2 + (T - t)(\sigma_2^e)^2, \\ \sigma_Y^2 &= t(\sigma_1^h)^2 + (T - t)(\sigma_2^h)^2, \\ \mathrm{Cov}(X, Y) &= \sigma_X\sigma_Y\rho_{XY} \\ &= t(\sigma_1^e\sigma_1^h\rho_1) + (T - t)(\sigma_2^e\sigma_2^h\rho_2). \end{aligned}$$

These five quantities μ_X, μ_Y, σ_X, σ_Y, ρ_{XY} completely specify the bivariate normal distribution and hence (by taking exponents) the bivariate lognormal distribution.[13] Note that this lognormal distribution is conditional on the the process spending t months in regime 1. Since we are focusing only on the terminal distribution at time T the sequencing of the months spent in regime 1 does not matter.[14] It will help to ease the notation in some later expressions if we let Γ_t denote the set of parameters:

$$\mu_X(t),\ \mu_Y(t),\ \sigma_X(t),\ \sigma_Y(t),\ \rho_{XY}(t),$$

where we now explicitly recognize the dependence of these five parameters on the sojourn time t.

Let p_i be the probability that the process spends a total of exactly i months in regime 1 during the total period (of T months) where $1 \leq i \leq T$. These probabilities can be obtained recursively (see Hardy 2003). Note that

$$\sum_{i=1}^{T} p_i = 1.$$

The investor's expected utility at time T is given by

$$\mathrm{E}[u(W)] = \sum_{i=1}^{T} p_i \mathrm{E}[u(W)|i], \tag{5.4}$$

where we use the notation $\mathrm{E}[u(W)|i]$ to denote the expected utility of terminal wealth given that the process spends exactly i months in regime 1 and $(T - i)$ months in regime 2. We have already mentioned that the distribution of the two risky assets follows a bivariate lognormal distribution in this case, and we also know the parameters of this distribution. Hence we can compute the expectation in equation (5.4) by evaluating the T separate terms where the two risky assets in each term have a bivariate lognormal distribution with known parameters given by Γ_i.

The procedure for finding the optimal investment triplet for a given horizon T is as follows. Assume that the process spends exactly i months in regime 1. Let the random returns on the two risky assets be (r_i^e, r_i^h) in this case. These variables have a bivariate lognormal distribution that is a function of the parameter set Γ_i. Hence using equation (5.1) the expectation in (5.4) becomes

$$\begin{aligned}\mathrm{E}[u(W)] = \sum_{i=1}^{T} p_i \mathrm{E}[u(W_0[1 + R + x_1(r_i^e - R) \\ + x_2(r_i^h - R))] \,|i], \end{aligned} \tag{5.5}$$

where

$$0 \leq x_1, \leq 1,\ 0 \leq x_2 \leq 1,\ 0 \leq (x_1 + x_2) \leq 1.$$

The right-hand side of this last equation consists of a linear sum of T expectations over different bivariate lognormal distributions. When these expectations are evaluated we are left with a function of the two choice variables x_1 and x_2. There is a unique solution because the utility function u is concave. It is straightforward to find the optimal values of x_1 and x_2. Since x_1 and x_2 lie in one half of the unit square, because of the no-short-selling constraint, we used a simple grid search to locate the optimal values.

We can use any suitable numerical method here to compute the various bivariate distribution.[15] For our purposes it is convenient to use quasi random numbers to do the simulation because we can obtain any desired level of accuracy with this method. In two dimensions the best quasi random numbers are constructed from the Fibonacci sequence. For $k \geq 3$, they are defined by the recurrence relation

$$N_k = N_{k-1} + N_{k-2},$$

where

$$N_1 = 1, \qquad N_2 = 1.$$

For each integer m, denote N_m by N_m. For $1 \leq i \leq N_m$,

$$\begin{aligned} u_i &= \frac{i}{N_m} - \frac{1}{2N_m}, \\ v_i &= \frac{\mathrm{mod}(iN_{m-1}, N_m)}{N_m} + \frac{1}{2N_m}. \end{aligned}$$

Recall that the operation $\mathrm{mod}(x, y)$ gives the remainder when x is divided by y. For example, $\mathrm{mod}(25,\ 7) = 4$. The u_i constitute a set of N_m

---

[13] We provide the formula for the bivariate lognormal distribution in the Appendix.

[14] Note we are considering only a buy-and-hold strategy in this paper. If we wanted to analyze portfolio revisions the regime sequencing would matter.

[15] We also could have used numerical integration.

well-dispersed points in the unit square. We use these points to obtain the bivariate normal distribution and hence the corresponding bivariate lognormal distribution.

We assume a power utility function so that

$$u(x) = \frac{x^\alpha}{\alpha}, \qquad \alpha < 1, \alpha \neq 0. \qquad (5.6)$$

It is conventional to include α in the denominator, but we would get the same results if the α term in the denominator was replaced by the sign of α. This utility assumption is often made in the finance literature to model investor preferences; it is further discussed in Panjer (1998). For this utility function the relative risk aversion parameter is constant and equal to $(1 - \alpha)$ (see Panjer 1998, ch. 4). Power utility functions have the property that the *optimal* investment proportions do not depend on the size of initial wealth W_0.

To get some plausible parameter estimates of α, we make some simplifying assumptions. We consider an investor that in the absence of the hedge fund would optimally invest a certain fraction in the core equity portfolio. We use the Merton ratio to relate the fraction invested in equities to the risk aversion parameter.[16] The fraction invested in equities in this case is given by

$$\frac{\bar{\mu} - r}{\sigma^2(1 - \alpha)}.$$

Suppose we assume that $\bar{\mu} = 0.10$, $\sigma = 0.15$, $r = 0.05$. These are close to the historical estimates.[17] Table 7 gives the relation between different values of α and the corresponding values of the Merton ratio.

Table 7

Merton Ratio for Different Levels of Risk Aversion

α	Relative Risk Aversion	Fraction in Equities
−1	2	1.11
−2	3	0.74
−3	4	0.56
−4	5	0.44

Hence, in our numerical work we assume relative risk aversion levels ranging from 2 to 5. We compute the optimal investment allocation when the available assets are the core equity portfolio, the hedge fund and the risk-free rate. We restrict the investment strategies to buy and hold strategies over different time horizons T.

6. RESULTS

In this section we give numerical results for one particular fund.[18] We use the event-driven fund as a representative hedge fund style. We carry out the calculations under three separate assumptions and summarize the results in Table 8. First we base the optimization on the parameters estimated in Section 4. This assumes there is no bias in the hedge fund returns.[19] Then we redo the calculations assuming the hedge fund returns are reduced by 3% per annum in each state to account for the bias in the hedge fund returns. Our third set of calculations assumes the hedge fund returns are reduced by 4.5% per annum in each state.[20]

From the first panel in Table 8, we note that there is a sizeable fraction invested in the hedge fund when there is no deduction for bias. For low levels of risk aversion the S&P 500 is more attractive, while for higher levels of risk aversion the hedge fund becomes more desirable. For example, when the risk aversion equals 4 and the horizon is 24 months, the investor puts 74% in the hedge fund and the balance of 26% in the

---

[16] Merton's result assumes the investor has a power utility function with parameter α and rebalances her portfolio continuously between a risky asset and the risk-free asset. It also assumes that the risky asset's return follows geometric Brownian motion in continuous time, which corresponds to the lognormal distribution in discrete time. Note that Merton's result is based on a continuous-time model with fixed mean and variance and that it is not based on a regime-switching model.

[17] The historical estimates from the S&P 500 are $\mu = 0.1094$, $\sigma = 0.1452$ for the period January 1994 to June 2006 (see Table 1). The returns in Table 1 are based on the log of the total monthly returns (including dividends). The relation with Merton's expected return $\hat{\mu}$ is $\mu = \bar{\mu} - \sigma^2/2$. Hence for this dataset $\bar{\mu} = 0.1199$. However, many observers expect a lower risk premium going forward, so we use $\bar{\mu} = 0.10$.

[18] We obtained results for each hedge fund style but have chosen to report only one to save space.

[19] De los Rios and Garcia (2006) show that the extent of the bias varies by hedge fund style and by the database used.

[20] Cvitanic et al. (2003) also used a reduction of 4.5% to allow for bias in the hedge fund expected returns.

Table 8

Optimal Allocation across Core Equity Portfolio, Hedge Fund, and Risk-Free Asset for Different Holding Periods Assuming We Start in Regime 1

T	1 Month			12 Months			24 Months			36 Months		
RRA	S&P	hdg	rf	S&P	hdg	rf	S&P	hdg	rf	S&P	hdg	rf
1. No Allowance for Bias in the Hedge Fund Returns												
2	0.64	0.36	0.00	0.58	0.42	0.00	0.57	0.43	0.00	0.58	0.42	0.00
3	0.39	0.61	0.00	0.37	0.63	0.00	0.37	0.63	0.00	0.36	0.64	0.00
4	0.27	0.73	0.00	0.26	0.74	0.00	0.26	0.74	0.00	0.26	0.74	0.00
5	0.19	0.81	0.00	0.19	0.81	0.00	0.20	0.80	0.00	0.19	0.81	0.00
2. Assuming Reduction in Hedge Fund Expected Returns of 3% pa												
2	1.00	0.00	0.00	1.00	0.00	0.00	1.00	0.00	0.00	1.00	0.00	0.00
3	1.00	0.00	0.00	1.00	0.00	0.00	1.00	0.00	0.00	1.00	0.00	0.00
4	1.00	0.00	0.00	0.85	0.15	0.00	0.84	0.16	0.00	0.82	0.18	0.00
5	0.82	0.18	0.00	0.67	0.33	0.00	0.65	0.35	0.00	0.64	0.36	0.00
3. Assuming Reduction in Hedge Fund Expected Returns of 4.5% pa												
2	1.00	0.00	0.00	1.00	0.00	0.00	1.00	0.00	0.00	1.00	0.00	0.00
3	1.00	0.00	0.00	1.00	0.00	0.00	1.00	0.00	0.00	1.00	0.00	0.00
4	1.00	0.00	0.00	1.00	0.00	0.00	1.00	0.00	0.00	1.00	0.00	0.00
5	1.00	0.00	0.00	0.90	0.10	0.00	0.88	0.02	0.10	0.85	0.00	0.15

core equity portfolio. There is little variation by horizon.

In the second panel we have allowed for bias in the hedge fund return by a deduction of 3%. In this case an investor with low levels of risk aversion will prefer the S&P 500 over the hedge fund. We need higher levels of risk aversion to see some investment in the hedge fund. For example, when the risk aversion equals 4 and the horizon is 24 months, the investor puts 84% in the core equity portfolio and the balance of 16% in the hedge fund.

In the third panel we increase the deduction for the hedge fund expected return to 4.5%. In this case there is no investment in the hedge fund for almost all levels of risk aversion and maturity horizons. It is only for high levels of risk aversion and longer times to maturity that the investor begins to put money in the risk-free asset. When the risk aversion is 5 and the horizon is 24 months, the investor puts 88% in the core equity portfolio, 2% in the hedge fund, and 10% in the risk-free asset. To some extent the hedge fund and the risk-free assets are substitutes. A similar point has been made by Cvitanic et al. (2003).

In Table 9 we examine the optimal strategy assuming we start in regime 2. In the top panel we do not adjust the hedge fund returns. For the one-month horizon the risk-free asset dominates for all levels of risk aversion. At the one-year horizon the hedge fund and the core equity portfolio are both attractive for low levels of risk aversion. For the 24- and 36-months horizon the optimal portfolio consists only of equity and the hedge fund. At low levels of risk aversion the equity portfolio dominates, but for higher levels the hedge fund becomes more attractive. Note that for longer horizons the allocations in the top panel of Table 9 are very similar to those in the top panel of Table 8. This implies that the effect of the initial regime wears off after one year or so.

In the next two panels, where we reduce assumed expected returns on the hedge fund, the hedge fund disappears from the optimal portfolio for all horizons and all levels of risk aversion.

Table 10 shows the corresponding allocation when there is a 50% probability of starting in either regime. We first discuss the top panel. As the maturity horizon lengthens, there is an increasing probability that the process will spend time in regime 1 and the optimal portfolio will consist of the S&P 500 and the hedge fund. As risk aversion increases there is an increasing allocation to the hedge fund.

In the second panel when the hedge fund expected returns are reduced by 3% the optimal portfolio includes the hedge fund only at the longer durations and the higher levels of risk aversion. For the two-year horizon and risk aversion of four the optimal allocations are 76% in

Table 9

Optimal Allocation across Core Equity Portfolio, Hedge Fund, and Risk-Free Asset for Different Holding Periods Assuming We Start in Regime 2

T	1 Month			12 Months			24 Months			36 Months		
RRA	S&P	hdg	rf	S&P	hdg	rf	S&P	hdg	rf	S&P	hdg	rf
1. No Allowance for Bias in the Hedge Fund Returns												
2	0.00	0.00	1.00	0.51	0.49	0.00	0.53	0.47	0.00	0.54	0.46	0.00
3	0.00	0.00	1.00	0.33	0.66	0.01	0.35	0.65	0.00	0.35	0.65	0.00
4	0.00	0.00	1.00	0.25	0.49	0.26	0.26	0.74	0.00	0.26	0.74	0.00
5	0.00	0.00	1.00	0.21	0.38	0.41	0.19	0.81	0.00	0.20	0.80	0.00
2. Assuming Reduction in Hedge Fund Expected Returns of 3% pa												
2	0.00	0.00	1.00	0.79	0.00	0.21	1.00	0.00	0.00	1.00	0.00	0.00
3	0.00	0.00	1.00	0.53	0.00	0.47	0.90	0.00	0.10	1.00	0.00	0.00
4	0.00	0.00	1.00	0.40	0.00	0.60	0.68	0.00	0.32	0.75	0.09	0.16
5	0.00	0.00	1.00	0.32	0.00	0.68	0.54	0.00	0.46	0.60	0.08	0.32
3. Assuming Reduction in Hedge Fund Expected Returns of 4.5% pa												
2	0.00	0.00	1.00	0.79	0.00	0.21	1.00	0.00	0.00	1.00	0.00	0.00
3	0.00	0.00	1.00	0.53	0.00	0.47	0.90	0.00	0.10	1.00	0.00	0.00
4	0.00	0.00	1.00	0.40	0.00	0.60	0.68	0.00	0.32	0.78	0.00	0.22
5	0.00	0.00	1.00	0.32	0.00	0.68	0.54	0.00	0.46	0.63	0.00	0.37

the S&P 500 and 24% in the hedge fund. In the third panel where we penalize the hedge fund expected returns by 4.5%, the hedge fund is never included in the optimal portfolio.

Overall we see that the investment in the hedge fund depends heavily on the assumed expected return. If there is no allowance for bias, there is a significant allocation to the hedge fund. In this case the hedge fund weight increases with time to maturity and level of risk aversion. If we reduce the assumed expected return on the hedge fund, the optimal allocation quickly declines.

It is clear that the optimal portfolio allocations are extremely sensitive to our assumptions about the expected return vector. This phenomenon is also a feature of the standard mean variance case where several authors (see Jobson and Korkie (1980), (1981); Heath Windcliff and Boyle

Table 10

Optimal Allocation across Core Equity Portfolio, Hedge Fund, and Risk-Free Asset for Different Holding Periods Assuming There Is an Equal Chance of Starting in Either Regime

T	1 Month			12 Months			24 Months			36 Months		
RRA	S&P	hdg	rf	S&P	hdg	rf	S&P	hdg	rf	S&P	hdg	rf
1. No Allowance for Bias in the Hedge Fund Returns												
2	0.17	0.00	0.83	0.53	0.47	0.00	0.55	0.45	0.00	0.56	0.44	0.00
3	0.11	0.00	0.89	0.34	0.66	0.00	0.35	0.65	0.00	0.35	0.65	0.00
4	0.08	0.00	0.92	0.25	0.75	0.00	0.26	0.74	0.00	0.26	0.74	0.00
5	0.07	0.00	0.93	0.20	0.80	0.00	0.20	0.80	0.00	0.20	0.80	0.00
2. Assuming Reduction in Hedge Fund Expected Returns of 3% pa												
2	0.17	0.00	0.83	1.00	0.00	0.00	1.00	0.00	0.00	1.00	0.00	0.00
3	0.11	0.00	0.89	0.93	0.00	0.07	1.00	0.00	0.00	1.00	0.00	0.00
4	0.08	0.00	0.92	0.71	0.00	0.29	0.76	0.24	0.00	0.77	0.23	0.00
5	0.07	0.00	0.93	0.57	0.00	0.43	0.60	0.24	0.16	0.61	0.33	0.06
3. Assuming Reduction in Hedge Fund Expected Returns of 4.5% pa												
2	0.17	0.00	0.83	1.00	0.00	0.00	1.00	0.00	0.00	1.00	0.00	0.00
3	0.11	0.00	0.89	0.93	0.00	0.07	1.00	0.00	0.00	1.00	0.00	0.00
4	0.08	0.00	0.92	0.71	0.00	0.29	0.86	0.00	0.14	0.90	0.00	0.10
5	0.07	0.00	0.93	0.57	0.00	0.43	0.69	0.00	0.31	0.73	0.00	0.27

(2004)) have shown that the portfolio allocations are very sensitive to the assumed expected return vector. Garlappi, Uppal, and Wang (2007) have proposed a method to deal with estimation risk in the standard mean variance model. It is much more challenging to allow for estimation risk in our more complicated model. Ideally we should consider the entire distribution, which involves all the parameters. In our case the asset with the highest volatility is the S&P 500, but all the adjustments were made to the expected return on the hedge fund, and we ignored estimation risk in the S&P 500 returns. It seems intuitive that when we introduce parameter uncertainty arising from estimation risk this should have a more significant effect on the S&P 500 return than the hedge fund. An investor concerned about parameter uncertainty will therefore tend to reduce the allocation to the S&P 500 relative to the other two asset classes. We plan to explore this issue in subsequent research.

6.1 Why Do More Risk-Averse Investors Put More into the Hedge Funds

It is a somewhat puzzling feature of our results that more risk-averse investors tend to put more in the hedge fund than less risk-averse investors.[21] This is certainly true in the case of Table 8, where the system is assumed to start in regime 1. This somewhat counterintuitive result has been noted before by Cvitanic et al. (2003). There are two factors that may help explain this result.

First, we have used an index of hedge fund returns and not the returns on the individual hedge funds. When we form an index of securities, we reduce the variance even if the returns on the individual securities are positively correlated.[22] This means that the volatility we have used to represent the hedge fund, which is based on the index, is lower than the volatility of a typical fund. The reason why the proportion allocated to hedge funds increases as the investor becomes more risk averse for the event-driven strategy is because this strategy has a low volatility of 5.6%. Just as in the Markowitz model, assets with low volatilities are desirable.

Second, we have seen in Table 2 that hedge fund returns are nonnormal and that in some cases they have considerable negative skewness. The risks associated with these types of return distribution are not well captured by volatility and may be better captured by other types of risk measure such as Value at Risk (VaR) or Conditional Tail Expectation (CTE). One way of highlighting the importance of the left-tail risk is to compare the estimates of VaR computed from the historical experience with the prospective estimates based on the regime-switching models estimated from the very same set of historical data. These comparisons are given in Table 11. We have estimated the one-month 95% VaR based on the historical data and on the regime-switching model assuming we are in regime 2.

We see that when the system is in regime 2 the VaR can be as much as nine times the VaR based on the historical data. An investor with concerns about this left-tail risk would not find this distribution of returns attractive. Our methodology, which is based on the maximization of expected utility, does not focus on this left-tail risk, although the regime-switching model does a better job of incorporating poor returns than the standard lognormal assumption. This left-tail risk may be of great concern to some investors, for example, a pension fund. If we were to rank investors' preferences solely in terms of their aversion to

Table 11

Comparison of Historical VaR and VaR Based on the Regime-Switching Model for Six Hedge Fund Strategies

Fund	Historical VaR	Regime-Switching VaR	Multiple of Historical
Convertible Arbitrage	1.79	2.98	1.66
Event Driven	1.25	11.16	8.93
Risk Arbitrage	1.30	3.16	2.43
Distressed	1.37	12.88	9.40
Fixed Income Arbitrage	1.36	4.01	2.95
Multi-Strategy	1.55	4.91	3.17

---

[21] We thank the referee for highlighting this issue.

[22] Suppose we have N stocks, each with variances σ^2, and all the pairwise correlations are equal to ρ. In this case the variance of the equally weighted portfolio of N stocks is

$$\frac{\sigma^2}{N} + \left(1 - \frac{1}{N}\right) \rho\sigma^2.$$

For example, if $\sigma = 0.10$, $\rho = 0.60$, $N = 100$, then the standard deviation of the equally weighted portfolio is 7.77%, whereas each stock has a standard deviation of 10%.

this type of risk, then our previous conclusions would no longer hold. However, in practice investors also care about positive returns, so there should be tradeoff. It is possible to extend the optimization approach to incorporate a VaR-type constraint (see, e.g., example Basak and Shaprio (2001) and Boyle and Tian (2007)), but we do not pursue this issue here.

As a final comment on this issue, we remark that there is evidence that hedge fund returns in general are becoming lower and more correlated with the market over time (Chan et al. (2007)). We also know from standard mean variance theory that both of these effects, lower expected return and higher correlation with the market, will tend to diminish the attractiveness of hedge fund investments as a substitute for the risk-free asset. These secular changes are consistent with the fact that our estimates for the optimal allocation to hedge funds are generally lower than those obtained by Cvitanic et al. (2003), who used data for the period 1996–2000.

7. CONCLUSION

This paper developed a simple method to estimate the optimal allocation to a hedge fund. We used a bivariate regime-switching model to capture the joint dynamics of the hedge fund return and the return on the market portfolio. We calibrated our model to the empirical data and stressed some limitations of the hedge fund return dataset. We derived a simple optimization procedure based on the expected utility paradigm and developed several sets of numerical results based on different time horizons, different risk aversion assumptions, and varying assumptions about the hedge fund expected returns.

We demonstrated how the optimal allocation depended on the investors's risk aversion and the horizon as well as the starting regime. We also found that the allowance made for bias in the hedge fund expected returns had a huge impact on the portfolio allocations. The optimal weights are extremely sensitive to the assumed hedge fund expected return in the optimization procedure. If we allow for the bias in hedge fund expected returns, then hedge funds tend to disappear from the optimal portfolio. Our optimization approach does not place emphasis on the left-tail risk: a risk that may be of great concern to certain investors. In addition, our model is based on the assumption that the future will be similar to the past. Despite its limitations we hope that the approach presented here may be of practical interest.

8. ACKNOWLEDGMENTS

This paper has benefited from the comments of Anca Drexler, Bin Wei, Bassam Aoun, Feidhlim Boyle, and Xiaofei Zhao. Phelim Boyle acknowledges research support from the Social Sciences and Humanities Research Council of Canada. The authors are grateful to an anonymous referee for constructive comments and the coeditor for many careful and helpful suggestions.

REFERENCES

ACKERMANN, CARL, RICHARD MCENALLY, AND DAVID RAVENSCRAFT. 1999. The Performance of Hedge Funds: Risk, Return and Incentives. *Journal of Finance* 54: 833–74.

ANG, ANDREW, AND GEERT BEKAERT. 2002a. Regime Switches in Interest Rates. *Journal of Business and Economic Statistics* 20(2): 163–82.

———. 2002b. Short Rate Nonlinearities and Regime Switches. *Journal of Economic Dynamics and Control* 26(7–8): 1243–74.

BASAK, SULEYMAN, AND ALEXANDER SHAPIRO. 2001. Value-at-Risk Based Risk Management: Optimal Polices and Asset Prices. *Review of Financial Studies* 14: 371–405.

BILLIO, MONICA, MILA GETMANSKY, AND LORIANA PELIZZON. 2006. Hedge Fund Tail Event Behavior: A Regime-Switching Approach. Working paper, Ca'Foscari, University of Venice.

BOUDREAULT, MATHIEU, AND CHRISTIAN-MARC PANNETON. 2006. Practical Considerations in Multivariate Modelling of Asset Returns for Actuarial Valuation of Investment Guarantees. Paper presented at Canadian Institute of Actuaries Stochastic Modelling Symposium, Toronto, April 3–4.

BOYLE, PHELIM, AND WEIDONG TIAN. 2007. Portfolio Management with Constraints. *Mathematical Finance* 17(3): 319–43.

BOYLE, PHELIM, AND HEATH WINDCLIFF. 2004. The 1/N Pension Investment Puzzle. *North American Actuarial Journal* 8(1): 32–45.

CALPERS. 2007. CalPERS Hikes Potential Allocations to Corporate Governance and Hedge Funds. Press Release, June 18. Sacramento, CA: CalPERS Headquarters.

CHAN, NICHOLAS, MILA GETMANSKY, SHANE M. HAAS, AND ANDREW W. LO. 2007. Systemic Risk and Hedge Funds. In *The Risks of Financial Institutions and the Financial Sector,* ed. M. Carey and R. Stulz, pp. 235–330. Chicago: University of Chicago Press.

CVITANIC, JAKSA, AI LAZRAK, LIONEL MARTELLINI, AND FERNANDO ZAPATERO. 2003. Optimal Allocation to Hedge Funds: An Empirical Analysis. *Quantitative Finance* 3: 29–39.

Danielsson, Jon, Ashley Taylor, and Jean-Pierre Zigrand. 2005. Highwaymen or Heros, Should Hedge Funds be Regulated? *Journal of Financial Stability* 1(4): 522–43.

David, Alexander, and Pietro Veronesi. 2001. Inflation and Earnings Uncertainty and the Volatility of Asset Prices: An Empirical Investigation. Working paper, University of Chicago, Graduate School of Business.

De los Rios, Antonio Diez, and René Garcia. 2006. Assessing and Valuing the Non Linear Structure of Hedge Fund Returns. Working paper 06-31, Bank of Canada, Ottawa, Canada.

Financial Services Authority. 2005. DP05/4: Hedge Funds: A Discussion of Risk and Regulatory Engagement. London: Financial Services Authority.

———. 2006. DP06/2: Hedge Funds: A Discussion of Risk and Regulatory Engagement—Feedback on DP05/4. London: Financial Services Authority.

Fung, William, and David Hsieh. 2000. Performance Characteristics of Hedge Funds and Commodity Funds: Natural versus Spurious Biases. *Journal of Financial and Quantitative Analysis* 35: 291–307.

Garlappi, Lorenzo, Raman Uppal, and Tan Wang. 2007. Portfolio Selection with Parameter and Model Uncertainty: A Multi Prior Approach. *Review of Financial Studies* 20(1): 41–81.

Gilks, W. R., S. Richardson, and D. J. Spiegelhalter. 1995. *Markov Chain Monte Carlo in Practice.* New York: Chapman & Hall.

Global Derivatives. 2003. Overview of Hedge Fund Strategies. http://www.globalderivatives.com/docs/HedgeFundStrategies(Nov03).pdf.

Guidolin, Massimo, and Allen Timmerman. 2005. Size and Value Anomalies under Regime Shifts. St. Louis: Federal Reserve Bank.

Hamilton, James, and Marcelle Chauvet. 2005. Dating Business Cycle Turning Points. NBER working paper no. 11422.

Hardy, Mary. 2001. A Regime-Switching Model of Long-Term Stock Returns. *North American Actuarial Journal* 5(2): 41–53.

———. 2003. *Investment Guarantees: The New Science of Modeling and Risk Management for Equity-Linked Life Insurance.* City, NJ: Wiley.

Jobson, J. Dave, and Bob Korke. 1980. Estimation for Markowitz Efficient Portfolios. *Journal of the American Statistical Association* 75(371): 544–54.

Kirschner, Sam., Eldon Mayer, and Lee Kellser. 2006. *The Investor's Guide to Hedge Funds.* Hoboken, NJ: Wiley.

L'Habitant, François-Serge. 2002. *Hedge Funds: Myths and Limits.* New York: Wiley.

———. 2004. *Hedge Funds: Quantitative Insights.* New York: Wiley.

Liang, Bing. 2000. Hedge Funds: The Living and The Dead. *Journal of Financial and Quantitative Analysis* 35: 309–26.

Malkiel, Burton G., and Atanu Saha. 2005. Hedge Funds: Risk and Return. *Financial Analysts Journal* 61(6): 80–88.

Panjer, Harry H., ed. 1998. *Financial Economics: With Applications to Investments, Insurance and Pensions.* Schaumburg, IL: Actuarial Foundation.

SEC. 2003. Implications of The Growth of Hedge Funds. Washington, DC: Securities and Exchange Commission.

Turner, Christopher M., Richard Startz, and Charles R. Nelson. 1989. A Markov Model of Heteroskedasticity, Risk, and Learning in the Stock Market. *Journal of Financial Economics* 25: 3–22.

Vikas, Agarwal, and Narayan Y. Naik. 2004. Risks and Portfolio Decisions Involving Hedge Funds. *Review of Financial Studies* 17(1): 63–98.

Whitelaw, Robert. 2000. Stock Market Risk and Returns: An Equilibrium Approach. *Review of Financial Studies* 13: 521–47.

Yago, Glennm, Lalita Ramesh, and Noah E. Hochman. 1998. Hedge Funds and Systemic Risk Demystified. Santa Monica, CA: Milken Institute.

APPENDIX

FORMULA FOR LOGNORMAL DISTRIBUTION

This Appendix gives the expression for the the bivariate lognormal distribution used in the paper.

The market index (S) and the hedge fund (H) are assumed to follow two correlated geometric Brownian motions under the $\mathbb{P}$ measure. The stochastic differential equations are

$$dS_t = S_t(\bar{\mu}_S dt + \sigma_S dW_1(t)),$$

$$dH_t = H_t(\bar{\mu}_H dt + \sigma_H dW_2(t)),$$

where the correlation between the Weiner processes is ρ.

The market index S_T and market value of the hedge fund H_T at some future time T follow a bivariate lognormal distribution given by

$$f(S_T, H_T|S_0, H_0) = \frac{1}{2\pi S_T H_T \sigma_S \sigma_H T\sqrt{1-\rho^2}} \times \exp\left\{\frac{-1}{2(1-\rho^2)}\left[\frac{\left(\ln\left(\frac{S_T}{S_0}\right) - \left(\bar{\mu}_S - \frac{\sigma_S^2}{2}\right)T\right)^2}{\sigma_S^2 T}\right.\right.$$

$$\left.\left. - \frac{2\rho\left(\ln\left(\frac{S_T}{S_0}\right) - \left(\bar{\mu}_S - \frac{\sigma_S^2}{2}\right)T\right)\left(\ln\left(\frac{H_T}{H_0}\right) - \left(\bar{\mu}_H - \frac{\sigma_H^2}{2}\right)T\right)}{\sigma_S\sigma_H T} + \frac{\left(\ln\left(\frac{H_T}{H_0}\right) - \left(\bar{\mu}_H - \frac{\sigma_H^2}{2}\right)T\right)^2}{\sigma_H^2 T}\right]\right\}.$$

We have

$$\mathrm{E}\left[\ln\left(\frac{S_T}{S_0}\right)\right] = \left(\bar{\mu}_S - \frac{\sigma_S^2}{2}\right)T,$$

$$\mathrm{E}\left[\ln\left(\frac{H_T}{H_0}\right)\right] = \left(\bar{\mu}_H - \frac{\sigma_H^2}{2}\right)T.$$

In the paper we estimated the mean from the log of one plus the returns. (See title of Table 1.) So the μ's used in the paper are related to the $\bar{\mu}$'s as follows:

$$\mu_S = \left(\bar{\mu}_S - \frac{\sigma_S^2}{2}\right),$$

$$\mu_H = \left(\bar{\mu}_H - \frac{\sigma_H^2}{2}\right).$$

Discussions on this paper can be submitted until April 1, 2008. The authors reserve the right to reply to any discussion. Please see the Submission Guidelines for Authors on the inside back cover for instructions on the submission of discussions.