Journal of Banking & Finance 37 (2013) 139–149

Contents lists available at SciVerse ScienceDirect

Journal of Banking & Finance

journal homepage: www.elsevier.com/locate/jbf

Dynamic hedge fund portfolio construction: A semi-parametric approach

Richard D.F. Harris *, Murat Mazibas

Xfi Centre for Finance and Investment, University of Exeter Business School, Streatham Court, University of Exeter, Exeter EX4 4ST, UK

ARTICLE INFO

Article history:
Received 23 October 2011
Accepted 19 August 2012
Available online 30 August 2012

JEL classification:
G11
G12
C14
C15
C61

Keywords:
Funds of hedge funds
Portfolio optimization
Copula
Extreme value theory
Monte Carlo simulation

ABSTRACT

In this article, we evaluate alternative optimization frameworks for constructing portfolios of hedge funds. We compare the standard mean–variance optimization model with models based on CVaR, CDaR and Omega, for both conservative and aggressive hedge fund investment strategies. In order to implement the CVaR, CDaR and Omega optimization models, we propose a semi-parametric methodology, which is based on extreme value theory, copula and Monte Carlo simulation. We compare the semi-parametric approach with the standard, non-parametric approach, used to compute CVaR, CDaR and Omega, and the benchmark parametric approach, based on both static and dynamic mean–variance optimization. We report two main findings. The first is that the CVaR, CDaR and Omega models offer a significant improvement in terms of risk-adjusted portfolio performance over the parametric mean–variance model. The second is that semi-parametric estimation of the CVaR, CDaR and Omega models offers a very substantial improvement over non-parametric estimation. Our results are robust to the choice of target return, risk limit and estimation sample size.

1. Introduction

Hedge funds have attracted much interest not only for their ability to generate relatively high average returns, but also for the large losses that they can incur, a risk that is exemplified by the rise and fall of Long Term Capital Management in the late 1990s. In spite of such risk, the hedge fund industry witnessed rapid growth in the 2000s, with assets under management reaching $1.93 trillion by early 2008. During the recent credit crisis, there was a significant reduction both in the number of hedge funds and in assets under management, which resulted from a combination of trading losses and asset withdrawals by investors. However, it is estimated that by April 2011 hedge fund assets had recovered to their pre-crisis level (see Strasberg and Eder, 2011). An important contributing factor to this recent growth has been the availability of funds of hedge funds, which enable investors to access hedge fund alpha with lower risk, albeit at the expense of an additional layer of fees. Another contributing factor to the growth in the hedge fund industry was the launch in the early 2000s of investable hedge fund indices. These have generated further interest from small-and medium-sized investors, who would otherwise be precluded from investing in the hedge fund market. Central to both of these developments is the role of portfolio optimization in order to construct portfolios of individual hedge funds or investible hedge fund indices.

A number of studies have examined portfolio optimization in a hedge fund context. However, the optimal portfolio allocation across individual hedge funds is complicated by the fact that owing to the strategies that hedge fund managers typically adopt, hedge fund returns are far from normally distributed, often exhibiting very significant negative skewness and excess kurtosis (see, for example, Amin and Kat, 2001; Lo, 2001; Brooks and Kat, 2002; Fung and Hsieh, 1997a, 2001; Agarwal and Naik, 2004; Hudson et al., 2006; Wegener et al., 2010). Portfolio optimization in the presence of such non-normality generally leads to very different portfolio allocations than those implied by mean–variance (MV) analysis (see, for example, McFall Lamm, 2003; Fung and Hsieh, 1997b; Cvitanic et al., 2003; Terhaar et al., 2003; Popova et al., 2003; Glaffig, 2006; Wong et al., 2008), which is a parametric approach. Motivated by the well established volatility clustering in hedge fund returns, Giamouridis and Vrontos (2007) show that the use of multivariate conditional volatility models improves portfolio performance and provides a more accurate tool for tail-risk measurement. Harris and Mazibas (2010) provide further evidence on the use of multivariate conditional volatility models in the context of dynamic hedge fund risk measurement and portfolio allocation, and show that simple volatility models, such as the RiskMetrics EWMA model of Morgan (1996), provide the biggest

* Corresponding author. Tel.: +44 1392 723215; fax: +44 1392 262475.
E-mail addresses: R.D.F.Harris@exeter.ac.uk (R.D.F. Harris), M.Mazibas@exeter.ac.uk (M. Mazibas).

0378-4266/$ - see front matter © 2012 Elsevier B.V. All rights reserved.
http://dx.doi.org/10.1016/j.jbankfin.2012.08.017

improvements in performance. The non-normality in hedge fund returns has prompted the use of alternative measures of risk in the optimization framework. Agarwal and Naik (2004) and Giamouridis and Vrontos (2007) compare mean–variance and mean-CVaR portfolios constructed using HFR hedge fund strategy indices. Krokhmal et al. (2003) compare the CVaR and CDaR approaches for minimum risk portfolios of individual hedge funds, while Hentati et al. (2010) compare the CVaR and Omega approaches. These alternative approaches rely on non-parametric estimation, in which the moments and quantiles of the density function of portfolio returns are estimated empirically, and these are used to compute the various risk measures used in the optimization process. The non-parametric approach, while straightforward to implement, relies on a large data sample to generate sufficiently accurate estimates of the various measures. Moreover, it does not readily lend itself to incorporating the well established dynamic characteristics of hedge fund returns, such as autocorrelation and volatility clustering.

In this paper, we propose a semi-parametric approach to hedge fund portfolio optimization that addresses the shortcomings of the parametric and non-parametric approaches. In the semi-parametric approach, we first standardize the returns of each portfolio constituent in order to filter out the predictable dynamics related to autocorrelation and volatility clustering. We model the marginal density of each standardized return series using a combination of extreme value theory (for the tails of the density) and a piecewise polynomial (for the center of the density), and construct the joint density of hedge fund index returns using a copula-based approach. We then simulate hedge fund returns from this joint density in order to compute the relevant moments and quantiles required for portfolio optimization. Using monthly index return data from the HFR database for the period 1990–2011, we compare our proposed semi-parametrically estimated models with their non-parametrically estimated counterparts. We examine the performance of the different estimation approaches for both conservative (i.e. minimum risk) and aggressive (i.e. maximum return) investors. For the aggressive investment strategy, we consider three different formulations of the optimization problem, each representing a different portfolio on the efficient frontier: minimization of risk subject to a target return, maximization of return subject to a risk constraint, and maximization of return per unit of risk. We also compare the semi-parametrically and non-parametrically estimated CVaR, CDaR and Omega models with a number of commonly used benchmarks, including a naïve (i.e. equally weighted) portfolio, a benchmark fund of hedge funds index and three parametric models (i.e. a constant volatility MV model and two time-varying volatility MV models). We report two main findings. The first is that the CVaR, CDaR and Omega optimization models offer a significant improvement in terms of risk-adjusted portfolio performance over the mean–variance optimization models and benchmark portfolios. The second is that semi-parametric estimation of the CVaR, CDaR and Omega models offers a very substantial improvement over non-parametric estimation. Our results are robust to the choice of target return, risk limit and estimation sample size.

The outline of the paper is as follows. Section 2 describes the optimization framework, the estimation methods and the evaluation criteria. Section 3 describes the data used in the empirical analysis. Section 4 reports the empirical results. Section 5 provides a summary and some concluding remarks.

2. Methodology

In this section, we first set out the generic optimization problem for two types of investor: conservative and aggressive, and define the different measures of risk that we use to construct the objective function in the optimization problem. We then describe the semi-parametric and non-parametric approaches to estimating the optimal portfolio in each case. Finally we define the evaluation criteria that are used to compare the different models.

2.1. Optimization framework

Consider an investor who allocates their wealth among m individual hedge funds or investible hedge fund indices, with portfolio weight vector $x = [x_1, \ldots, x_m]'$. For a conservative investor, the portfolio optimization problem is given by

$$\begin{aligned} &\min_{\mathbf{x}} \Phi_p(\mathbf{x}) \\ &\text{subject to } \mathbf{x} \geqslant 0, \mathbf{x}'\mathbf{1}_m = 1 \end{aligned} \tag{1}$$

where $\Phi_p(\mathbf{x})$ is the portfolio risk function to be minimized (i.e. standard deviation, CVaR, CDaR or lower partial moment) and $\mathbf{1}_m$ is an m-vector of ones. The budget constraint and non-negativity constraints yield an unleveraged long-only portfolio. For an aggressive investor, we consider three separate formulations of the portfolio optimization problem. The first minimizes portfolio risk subject to a target expected portfolio return, r_0:

$$\begin{aligned} &\min_{\mathbf{x}} \Phi_p(\mathbf{x}) \\ &\text{subject to } \mathbf{x} \geqslant 0, \mathbf{x}'\mathbf{1} = 1, E(r_{p,t}) \geqslant r_0 \end{aligned} \tag{2}$$

where $\mathbf{r}_t = [r_{1,t}, \ldots, r_{m,t}]'$ is the m-vector of hedge fund returns at time t and $r_{p,t} = \mathbf{x}'\mathbf{r}_t$ is the portfolio return at time t. The second formulation of the portfolio optimization problem for an aggressive investor maximizes portfolio return subject to a portfolio risk constraint:

$$\begin{aligned} &\max_{\mathbf{x}} E(r_{p,t}) \\ &\text{subject to } \Phi_p(\mathbf{x}) \leqslant \omega C, \quad \mathbf{x} \geqslant 0, \quad \mathbf{x}'\mathbf{1} = 1 \end{aligned} \tag{3}$$

where ω is the risk limit and C is the invested capital, which is set arbitrarily to 1. The third formulation of the portfolio optimization problem for an aggressive investor maximizes portfolio expected excess return per unit of portfolio risk:

$$\begin{aligned} &\max_{\mathbf{x}} \frac{E(r_{p,t}) - r_f}{\Phi_p(\mathbf{x})} \\ &\text{subject to } \mathbf{x} \geqslant 0, \quad \mathbf{x}'\mathbf{1} = 1 \end{aligned} \tag{4}$$

where r_f is risk-free rate of return. This formulation is a generalization of the tangency portfolio in expected return-risk space for the different risk measures.

2.2. Optimization models

We now define the different risk measures, Φ_p, that are used in the optimization problems described above.

2.2.1. Mean–variance optimization model

As a benchmark, we use the standard mean–variance model of Markovitz (1952), in which the risk measure is portfolio standard deviation, given by

$$\sigma(\mathbf{x}) = [\mathbf{x}'\mathbf{H}\mathbf{x}]^{1/2} \tag{5}$$

where $\mathbf{H}$ is the $m \times m$ covariance matrix of hedge fund index returns. We consider two versions of the mean–variance model. In the static mean–variance model, $\mathbf{H}$ is estimated using the sample covariance matrix. In the dynamic mean–variance model, $\mathbf{H}$ is estimated using the multivariate RiskMetrics EWMA model of Morgan (1996) and the multivariate DCC-GARCH model of Engle and