Bayesian Modeling of Hedge Fund Return Characteristics

Weiren Chang[1], PhD, CFA
JP Morgan Global Access Portfolios, 270 Park Avenue, 5th Floor, New York, NY 10017

Abstract
The objective of this work is to show the rich and intuitive inference capabilities of the Bayesian framework and its applications to the financial services industry. Parametric and nonparametric methods were used to analyze returns characteristics of a representative fund of hedge funds[2] (FoHF). Regression models were used to capture market factor sensitivities of the FoHF. Multiple Bayesian models were proposed; results by Markov Chain Monte Carlo simulations were compared and discussed. Actual market data and hedge fund returns were used in this study.

Key Words: Hedge fund, Fund of hedge funds, Bayesian, Markov Chain Monte Carlo, MCMC

1. Introduction

Hedge funds (HFs) are privately placed investment vehicles that are free to invest in a wide range of securities and derivatives, are typically more actively managed, take both long and short positions, and use leverage to increase returns. Many hedge funds aim to achieve more risk-controlled or asymmetric return profiles compared to long-only mutual funds. Hedge funds are only available to accredited investors and qualified purchasers.

Many investors access hedge funds via a multi-strategy fund of hedge funds (FoHF). A typical FoHF dynamically allocates to underlying hedge funds of various styles and the weights are not available to outside investors. This form of investment represents a large portion of hedge fund investments. Even if an institution invests in single-strategy hedge funds directly, it typically diversifies into various strategies, thus forming a portfolio similar to a multi-strategy FoHF. It is important for hedge fund investors to understand and analyze FoHF return characteristics.

The typical FoHF investor would ask the following questions:
1. Is the return good enough? (or, average monthly return $\geq$ 0.6%?)
2. Is it not more volatile than desired? (or, monthly standard deviation $\leq$ 2%?)
These questions are roughly equivalent to asking whether the fund's annualized return is above 7.2% (close to the typical Libor + 7% target) and its annualized volatility is below 6.9%. In the current zero interest rate environment, this implies a desired Sharpe ratio of about one.

Figure 1 shows the monthly return histogram of a representative FoHF, which has an

1 The views expressed in this article are those of the author and do not necessarily reflect the views of his employer. Email: weiren.chang@jpmorgan.com
2 This FoHF is not affiliated with the author's employer.

empirical probability density with a heavy left tail (due to the 2008 financial crisis). The fund's long-term average monthly return is 0.28% with a monthly standard deviation of 1.53%, skewness of -1.4 and extra kurtosis of 3.6.

Most practitioners in the financial services industry use classical statistical methods to analyze return characteristics of hedge funds and fund of hedge funds. Common techniques include return distribution analysis (e.g. mean, standard deviation, skew, drawdown, Sharpe ratio), single and multiple regressions, rolling-window statistics, clustering analysis, and principal component analysis.

This work analyzes return characteristics of major hedge fund styles and a representative multi-strategy FoHF focused on Bayesian methods. We also try to better understand the FoHF's underlying drivers and market factor sensitivities, as well excess returns over benchmarks. This paper's objective and scope are different from that of [Dewaele 2011], which used the classical bootstrap method to assess the average excess return of a large group of FoHF's. Market data from January 2004 to April 2014, are used in the study. This time period covers an economic expansion period, followed by a financial crisis and deep recession, then an ongoing uneven recovery period.

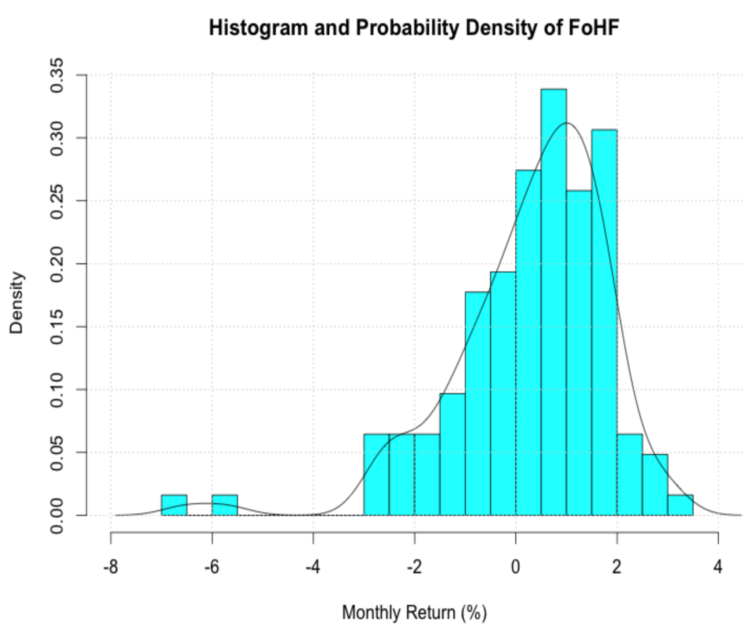

Figure 1: Long-term Return Histogram of FoHF

Bayesian models in this study are coded in the language of BUGS (Bayesian inference Using Gibbs Sampling). Open-source software systems, R and JAGS (Just Another Gibbs Sampler), were used to perform Markov Chain Monte Carlo (MCMC) simulations and statistical analysis.

Market factors used to explain the FoHF returns include US and international equity indices, US Treasury yields, US and emerging market corporate bond indices, commodity prices, and traded weight US dollar index. Market data, including indices and exchangetrade funds (ETF) as proxies, are downloaded (free of charge) from the Federal Reserve Economic Data (FRED) center and yahoo.com.

There are several hedge fund strategy classifications, either asset weighted and equally weighted. The asset-weighted indices are provided by Credit Suisse, formerly known as Dow Jones CS hedge fund indices. The 9 major hedge fund strategies are Long/Short Equity (LSE), Equity Market Neutral (EMN), Event Driven (ED), Fixed Income Arbitrage (FIArb), Convertible Arbitrage (CVArb), Global Macro (GM), Managed Futures (MF), Emerging Markets (EM), and Dedicated Short (DS). Monthly index returns are available for download (free of charge) from www.hedgeindex.com. Another major HF index provider is the Hedge Fund Research (HFR) group, which has different strategy classifications and publishes equally weighted index data. HFR data is available by subscription only.

The paper first highlights the benefits of Bayesian methods in contrast to classical statistics. Various Bayesian models and simulation results are presented next. Results from classical statistical analysis provides additional insights to hedge fund characteristics. The author will also discuss potential applications of Bayesian methods to hedge fund analysis.

2. Brief Introduction to Bayesian Statistics

Probability theory studies random phenomena and processes modeled with supposedly known parameters. A probabilistic model is a data-generating process. Statisticians collect data (observational or experimental), try to discover the process behind the data to gain insight, and make inferences based on the findings. To a statistician, the parameters used to fit a model are always estimated and not actually known. Uncertainties in probability are *aleatory*; while uncertainties faced in statistics are mostly *epistemic* or the combination of both. The nature of this study is statistical and is focused on the Bayesian approach.

Bayesian statistics originated from Bayes' theorem, published posthumously in 1763, and later independently discovered by Laplace in 1812. The theorem links the posterior probability to prior and conditional probabilities of events A and B in a simple form:

$$\Pr(A|B) = \Pr(B|A) \Pr(A) / \Pr(B)$$

Pr(A) and Pr(B) are prior probabilities; Pr(B|A) is the conditional probability of B given A; Pr(A|B) is the posterior (conditional) probability of A given B.

This seemingly simple theorem in probability is deeper than it looks and has wide-ranging applications. *An unknown model parameter is viewed as a random variable by Bayesian statisticians*. In other words, the epistemic uncertainty about a parameter is quantified as probability. By contrast, classical/frequentist statisticians view a model parameter unknown but fixed. Bayesian statistics extend the theorem's probabilistic thinking into a coherent and flexible inferential framework. Box had a nice comment on Bayesian inference: “Bayesian methods allow greater emphasis to be given to scientific interest and less to mathematical convenience” [Box 1973].

We can extend Bayes' theorem as follows. Let A = Θ (a model parameter or hypothesis), and B = D (data), thus

$$\Pr(\Theta|D) = \Pr(D|\Theta)\ \Pr(\Theta)\ /\ \Pr(D)$$

The Bayesian view considers $\Pr(\Theta)$ as the prior degree of belief about Θ; $\Pr(D|\Theta)$ as the likelihood function of D given Θ; $\Pr(\Theta|D)$ as the posterior degree of belief about Θ given that D has occurred. Probabilities can be further extended to probability density functions if Θ and D have continuous values.

$$f(\Theta|D) = f(D|\Theta)\ f(\Theta)\ /\ f(D)$$

That is, posterior density of Θ is proportional to its likelihood function multiplied by its prior density. It is important to note that $f(\Theta|D)$, i.e. probability density function of a model parameter given observed data, is what statistics are concerned with. However, classical statistics and inferences are based on $f(D|\Theta)$, assuming a known functional form in f and a fixed model parameter Θ, then use data to fit the unknown value of Θ. The mis-interpretations of p-value are discussed in [Berger 1988].

Classical/frequentist statistics dominated most of the second half of the twentieth century, though the Bayesian approach has been advocated as a sound alternative by prominent statisticians [Box 1973] such as Jeffreys, De Finetti, Lindley, Box, etc. Prior to 1990, conjugate probability density functions were commonly used to facilitate analytical solutions of posterior densities in order to perform statistical inference. Simulation-based Bayesian methods started to gain tractions in the mid-1990's due to efficient MCMC algorithms for posterior sampling combined with availability of high-speed computing. The Nobel Prize in Economics in 2011 was awarded to Christopher Sims due to his analysis of monetary policy effects using the Bayesian approach [Sims 2011]. The Bayesian paradigm has been gaining acceptance by theorists and practitioners of various fields in recent years. In recent years, many researchers applied hierarchical Bayesian modeling to complex problems successfully.

One major concern about the Bayesian approach is its subjectivity due to use of prior probability. However, it is advantageous to incorporate expert knowledge to guide parameters' prior probabilities, which can improve the accuracy of posterior probability estimates. The “illusion of objectivity” is also discussed in [Berger 1988]. Lack of long-term return history when analyzing hedge funds is common, the Bayesian method can come to the rescue since it does not always require long-run data to make inference. It is also important to note that the posterior probability is the weighted average of prior and data. When there are enough data, the potential “bias” in prior is generally not an issue. Also, when non-informative priors are used, the data usually dominate the posterior estimates, so the bias in prior is not an issue.

The Bayesian statistical inference framework is based on probabilistic reasoning. It provides a natural way to calculate a parameter's credible interval (or Bayesian inference interval) and its associated probability, hence providing richer inferential properties. This concept of probability associated with a parameter's credible interval does not exist in the classical/frequentist statistical paradigm.

This flexible and coherent framework facilitates intuitive and deeper understanding to statistical inference problems. The approach does not resort to ad hoc tricks often proposed in classical statistics. It also allows iterative and sequential updating of one's knowledge about model parameters, as D.V. Lindley said: “Today's posterior is tomorrow's prior” [Lindley 2000].

The four-step process in Bayesian analysis consists of:

1. Modeling: Model parameters of interest as random variables, specify prior (density) functions and likelihood.
2. Updating: Update posterior (density) functions of parameters using data, priors, and likelihood.
3. Sampling: Generate samples from posterior (density) functions.
4. Inference: Summarize on parameters from posteriors samples using point or interval estimates.

3. Bayesian Modeling and Analysis of FoHF Returns

In the financial service industry, most statistical analyses rely on the ubiquitous Normal distribution model using the classical/frequentist approach. The author believes the Bayesian paradigm offers an attractive alternative.

The following sections discuss the Bayesian models used to model the FoHF returns in this study. Markov Chain Monte Carlo (MCMC) simulations were done using open-source R and JAGS, which is a convenient tool for generating posterior probability samples based on Gibbs Sampling algorithm.

The following are simulation parameters throughout this study.

1. Number of steps in adaptation: 1000
2. Number of steps in burn-in: 2000
3. Number of MCMC chains: 3
4. Number of thinning steps: 5
5. Number of posterior simulation steps per chain: 5000

It is advisable to make sure posterior samples do not show significant autocorrelations.

3.1. Normal Likelihood of Two Parameters, Non-informative Priors

We first use a simple Bayesian model as a baseline to demonstrate the modeling process. Monthly FoHF return data from the boom period of 2004 to 2007 are used. Table 1 shows the BUGS code of the model. The likelihood function is a Normal density function with two parameters: average monthly return (mu) and the monthly standard deviation (sigma). Non-informative (or “diffuse”) Normal and Gamma prior density functions are specified for mu and prec (“precision”) respectively. The inverse of variance is called “precision”. The standard deviation, sigma, is the inverse of square root of precision. Two types of posterior inferences are done: point and interval estimates. The latter, called credible interval, has a (useful) probability associated with it.

Table 2 summarizes the posterior inference of the baseline model. It is important to note that both mu and sigma are considered as random variables. Both have posterior probability distributions, so it is possible to make statistical inference (using probabilistic reasoning) on both parameters. Point estimates from Bayesian and classical approaches gave similar results. Interpretations and the inference of interval estimates are meaningfully different.

The Bayesian inference results on the baseline model showed low probabilities of the FoHF having a negative average monthly return and being more volatile than the desired monthly standard deviation of 2%. However, the probability of having good (≥0.6%)

average monthly return is only 67%. The benefit of using the credible interval is that an intuitive probability measure is given. By contrast, classical statistics would state “the null hypothesis can not be rejected”, which does not provide any further insight. The Bayesian posterior probability is also different from p-value. The former is the desired $\Pr(\Theta|D)$; the latter is $\Pr(D|\Theta)$, which assumed Θ is known and fixed but this is simply a mathematical convenience and not our view.

Table 1: BUGS Code: Baseline Bayesian Model (#1). Monthly returns data from 2004 to 2007

```
model
{
  for( i in 1:nData )
  { # likelihood f(D|Θ)
    y[i] ~ dnorm(mu, prec)
  }
  # priors f(Θ)
  mu ~ dnorm(0, 0.01)
  prec ~ dgamma(0.001, 0.001)
  sigma <- 1/sqrt(prec)
}
```

Table 2: Posterior Inference: Baseline Bayesian Model. Monthly returns data from 2004 to 2007

Average Monthly Return (mu)	Monthly Standard Deviation (sigma)
Posterior mean of mu = 0.70% 95% credible interval [0.28%, 1.10%]	Posterior mean of sigma = 1.44% 95% credible interval [1.18%, 1.79%]
Probability of having good avg. return: $\Pr(mu \geq 0.6\% \mid D) \approx 67\%$ Probability of having negative avg. return: $\Pr(mu < 0\% \mid D) < 0.1\%$	Probability of being more volatile than monthly 2%, $\Pr(\sigma \geq 2\% \mid D) \approx 0.2\%$
Sample mean of mu = 0.7%	Classical point estimate of sigma = 1.25%

3.2. Normal Likelihood of Two Parameters, Subjective Priors

The second model has the same structure as the baseline except prior probabilities. What if at the end of 2007, the FoHF manager expected an uncertain year ahead even though the fund had been experiencing positive returns? The Bayesian framework allows one to inject expert opinions into prior probabilities. Many criticize the subjectivity of prior expectations. However, this can be beneficial when used prudently.

Table 3 shows the BUGS code of the second model. It reflects the manager's new belief of having negative average return and higher volatility. Instead of using 5 years of data, only 2 years of monthly data are used in the model—this reduced the influence of past data and gave more weight to the prior probability/belief.

Table 3: BUGS Code: Bayesian Model (#2) with Subjective Priors. Monthly returns data from 2006 to 2007

```
model
{
  for( i in 1:nData )
  { # likelihood f(D|Θ)
    y[i] ~ dnorm(mu, prec)
  }
  # priors f(Θ)
  mu ~ dnorm(-0.025, 10000)
  prec ~ dgamma(15, 0.01)
  sigma <- 1/sqrt(prec)
}
```

Table 4 summarizes the posterior inference of the model with subjective priors. The probability of having good average monthly return dropped to 21% from 67% of the baseline model. The probability of being more volatile than 2% increased to 81% from 0.2% of the baseline model. The author obviously had the *benefit of hindsight* about 2008 in building this model. This is done to show how subjective/informative priors can help Bayesian modeling. The Bayesian posterior inference gives much more insight to the researcher. It also allows direct comparison between models. This Bayesian model also showed how to incorporate forward-looking view as priors. The posterior 95% credible intervals of mu and sigma contain the realized/empirical values of average monthly return and monthly standard deviation from 2007 to 2008. This implies the subjective priors are in line with the data.

Table 4: Posterior Inference: Bayesian Model #2 with Subjective Priors. Monthly returns data from 2006 to 2007

Average Monthly Return (mu)	**Monthly Standard Deviation (sigma)**
Posterior mean of mu = 0.25% 95% credible interval [-0.58%, 1.04%]	Posterior mean of sigma = 2.20% 95% credible interval [1.82%, 2.70%]
Probability of good avg. return ≥ 0.6%: Pr(mu ≥ 0.6% | D) ≈ 21% (cf. 67%) Probability of having negative avg. return: Pr(mu < 0 | D) ≈ 26% (cf. <0.1%)	Probability of being more volatile than 2%: Pr(σ ≥ 2% | D) ≈ 81% (cf. 0.2%)
Realized average monthly return (2006 to 2008) ≈ -0.098%	Realized monthly standard deviation (2006 to 2008) ≈ 2.23%

There are two immediate possible follow-up models:

1. Inject new expectation at the end of 2008 (or 2009) and specify subjective priors of positive average return and lower volatility. The use of subjective priors must be justified based on careful research or expert knowledge.
2. Introduce auto-correlations in monthly returns. For example, model mu as a random walk with noise. This type of model is useful during periods of extreme market sentiments (2008 to 2009).

3.3. Bayesian Single-Factor Regression Model, t-distribution Likelihood, Non-informative Priors

It's important to understand which market factors drive hedge fund and FoHF returns. Common factors include equity indices, benchmark interest rates, investment grade credit and high yield indices, commodities, and currencies. US large-cap equity index S&P 500 is the most followed in the world. Many asset classes have high correlations with it. Our third Bayesian model uses the exchange traded fund (ETF) of the index as the covariate. The objective is to find the sensitivity and excess return of the FoHF vs. the benchmark equity market.

Table 5 shows the BUGS code of the single-factor regression model. To allow fatter tails of monthly returns, Student's t-distribution likelihood is used. Non-informative priors are used for model parameters (mu, sigma, df, alpha, beta). The parameter df is the degree of freedom in the t-distribution density function.

Table 5: BUGS Code: Bayesian Single Regression Model. Monthly returns data from 2004 to 2014

```
model {
  for( i in 1:nData ) {
    y[i] ~ dt(mu[i], prec, df)
    mu[i] <- alpha + beta*x[i]
  }
  df ~ dunif(1,5)
...
sigma <- 1/sqrt(prec)
}
```

Table 6 shows inference results of the Bayesian regression model. Assuming a monthly fee of 0.1%, the after-fee alpha of the FoHF has only 63% probability of being positive, which does not give a high-degree of comfort to the investor. By contrast, classical inference cannot reject the null hypothesis—this is not intuitive and provides limited further insight. It's notable that the probability of having positive beta is close to 100%. The probability of having beta > 0.2 is 86%. This means the investor of the FoHF should expect meaningful but moderate sensitivity to equities. The posterior mean of the degree of freedom is 4.55, which indicates fatter tails than Normal distribution.

3.4. Considerations in Multi-Factor Analysis

Though we advocate Bayesian methods, many classical statistical techniques are useful for exploratory analysis. Prior to using 9 market factors as regression covariates, principal component analysis (PCA) and hierarchical clustering analysis (HCA) were done on the 9 HF strategy indices and 15 directional market factors.

Some of the hedge fund styles are highly correlated, for example, LSE and ED indices have correlation of 0.9. These indices are not suitable as explanatory variables for FoHF or HF returns. Market factors are more suitable, especially as explanatory variables for hedge fund returns. The author started from a set of 15 market factors and gradually reduced to 9 adjusted factors, using PCA, HCA, and matrix condition numbers for guidance.

PCA (on correlations) showed that first 4 principal components (PCs) explained over

90% of return variations in the group of 15 factors plus FoHF; 4 PCs explained over 90% of variations in the 9 HF strategy group. HCA (on correlations) showed roughly 5 distinct styles: 4 less correlated strategies and a group of 5 directional strategies. This type of nonparametric analysis is applicable for analyzing a group of hedge funds being considered for investment.

Table 6: Posterior Inference: Bayesian Single Regression Model. Monthly returns data from 2004 to 2014

Average Monthly Excess Return (alpha)	**Equity Market Sensitivity (beta)**
Posterior mean of alpha = 0.15% 95% credible interval [-0.16%, 0.44%]	Posterior mean of beta = 0.24 95% credible interval [0.16,0.32]
Probability of positive alpha: Pr(a > 0 | D) ≈ 84% Probability of positive after-fee alpha: Pr(a > 0.1% | D) ≈ 63%	Probability of positive beta: Pr(b > 0 | D) ≈ 100% Probability of beta higher than 0.2: Pr(b > 0.2 | D) ≈ 86%
Classical regression: Cannot reject null hypothesis that alpha is 0 (p-value > 0.2)	Classical regression: beta = 0.248 with p-value << 0.05 ⇒ Reject null hypothesis
Other Inferences	
Posterior mean of sigma = 1.68% ⇒ Monthly standard deviation is higher than sample estimate of 1.53%	Mean of df = 4.55 ⇒ Slightly fatter tails than Normal distribution

3.5. Bayesian Multi-Factor Regression Model for FoHF, t-distribution Likelihood, Non-informative Priors

We used 9 adjusted market factors (monthly data from 2004 to 2014) as covariates in our multi-factor regression study.

- Equities: X1=US Large Cap, X2=(Small Cap – Large Cap), X3=(Growth – Value), X4=(Emerging Market – Developed Market)
- Rates: X5=US 10-year Treasury Yield
- Credit: X6=(High Yield – Investment Grade)
- Commodities: X7=Oil, X8=Gold
- Currency: X9=US Dollar Index

After factor reductions, there is still slight collinearity between X1 (Large Cap) and X6 (HY—IG) with correlation of 0.69, though much more benign than the group of 15 directional factors. A classical step-wise regression was first run using p-value of 0.05 as the threshold. The intercept (alpha) is not "significant". Significant betas are Large Cap (b1=0.17), Growth–Value (b3=0.12), EM–DM (b4=0.1), US 10-year Yield (b5=0.018), US Dollar Index (b9=-0.28).

Table 8 shows part of the BUGS code of the Bayesian multi-regression model. Table 9 shows the summary of the model's inference. Note that the Bayesian model's betas of (HY–IG) and (USD) have wide credible intervals that contain zero, hence unreliable. The former is partly due to collinearity with Large Cap equity. It's another example of the Bayesian method providing more insight in the inference process.

Table 7: BUGS Code: Bayesian Multi-Factor Regression Model. Monthly returns data from 2004 to 2014

```
model
{
for( i in 1:nData )
{
  y[i] ~ dt(mu[i], prec, df)
  mu[i] <- a + inprod(b[], x[i,])
}
prec ~ dgamma(0.01, 0.01)
sigma <- 1/sqrt(prec)
df ~ dunif(1, 5)
a ~ dunif(-0.02, 0.02)
for(j in 1:nPredictor)
{
  b[j] ~ dt(muB, precB, dfB)
}
muB ~ dnorm(0, 0.1)
precB ~ dgamma(0.01, 0.01)
...
}
```

Table 8: Posterior Inference: Bayesian Multi-Factor Regression Model. Monthly returns data from 2004 to 2014

Average Monthly Excess Return (alpha)	**Equity Market Sensitivity (beta)**
Posterior mean of alpha = 0.13% 95% credible interval [-0.17%, 0.43%]	Posterior means of betas: b1=0.148, b2=0.026, b3=0.071, b4=0.075, b5=0.014, b6=0.064, b7=0.018, b8=0.043, b9=-0.075
Posterior probability $\Pr(a > 0 \mid D) \approx 80\%$ vs. classical: Cannot reject H0 Probability of positive after-fee alpha: $\Pr(a > 0.1\% \mid D) \approx 63\%$	Which betas have 90% credible intervals that do not contain 0? Large cap (b1=0.148) EM–DM (b4=0.075) US 10-Yr Rate (b5=0.014)
Classical regression: Cannot reject null hypothesis that alpha is 0 (p-value > 0.2)	Using p-value of 0.05 as cut-off: Large cap (b1=0.17) Growth–Value (b3=0.12) EM–DM (b4=0.1) US Rate (b5=0.018) US Dollar Index (b9 = -0.28).
Other Inferences	
Posterior mean of sigma = 1.60% ⇒ Monthly standard deviation higher than sample estimate of 1.53%	Posterior mean of df = 4.54 ⇒ Slightly fatter tails than Normal distribution

3.6. Bayesian Multi-Factor Regression Model and Model Comparison

The method used in Section 3.4 can also be applied to modeling hedge fund returns. Let's use Credit Suisse Fixed Income Arbitrage (FIArb) strategy index as the dependent variable. It's useful to use a (classical) hierarchical clustering dendrogram, such as Figure 2, to visualize the correlations (or covariances) between the dependent variable and covariates. In this case, the dependent variable FIArb is closest to LC and (HY—IG), next to (SC—LC) and (EM—DM).

We expect FIArb to have higher credit market exposure so subjective priors were used to reflect this knowledge. The BUGS code is similar to that in Table 8 except the priors of the beta vs. (HY—IG), b6. Other betas have non-informative priors. Two models used two slightly different priors for b6. The first is t-distribution centered at 1 with precision of 1000; the second is centered at 0.5 with precision of 10. The posterior means of b6 are 0.33 and 0.31 respectively. The 90% credible intervals of b6 are [0.147, 0.518] and [0.142, 0.482]. All others betas' 90% credible intervals contain zero in both models, so not reliable. This implies that the model is not very sensitive to the priors.

The probabilities of FIArb having positive alpha are over 95% in both models; two posterior means of monthly alpha are 0.373% and 0.367% respectively. When comparing multiple Bayesian models, the model with the lowest deviance information criterion (DIC) is preferred. The DIC values of the two models are -706 and -707. The two models are very close by DIC, point estimates, and credible intervals. Classical statistics do not have such a coherent set of methods to compare models.

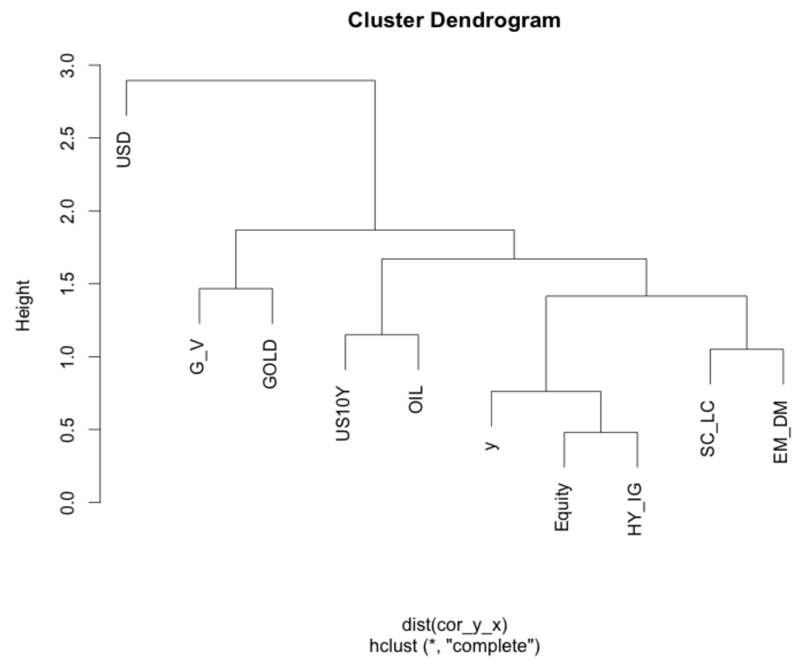

Figure 2: Clustering Dendrogram of CS-FIArb and Market Factors

4. Concluding Remarks

We showed that the Bayesian approach is powerful in modeling hedge fund and FoHF return characteristics. The Bayesian paradigm treats model parameters as random variables, demands more thinking in the modeling process including setting up likelihood and priors, requires more computations, and has intuitive and more meaningful inference potentials. The author expects to continue research on using Bayesian methods to dynamic modeling of hedge funds and market factors.

References

Berger, J.O., and Berry, D.A. (1988), "Statistical Analysis and the Illusion of Objectivity", American Scientist, Vol. 76

Box, G.E.P., and Tiao, G. (1973), Bayesian Inference in Statistical Analysis, Wiley-

Interscience Publication

Dewaele, et al (2011), “Assessing the Performance of Funds of Hedge Funds”, CEB Working Paper No. 1/041, Sep. 2011

Lindley, D.V. (2000), “The Philosophy of Statistics”, The Statistician, Pt. 3, pp. 293-337

Sims, C. A. (2011), Statistical Modeling of Monetary Policy and its Effects, Nobel Prize Lecture