THE UNIVERSITY OF CHICAGO

THE DIRICHLET PORTFOLIO MODEL: UNCOVERING THE HIDDEN COMPOSITION OF HEDGE FUND PORTFOLIOS

A DISSERTATION SUBMITTED TO
THE FACULTY OF THE UNIVERSITY OF CHICAGO
BOOTH SCHOOL OF BUSINESS
IN CANDIDACY FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

BY
LASZLO FRANK KORSOS

CHICAGO, ILLINOIS
AUGUST 2013

To the Korsos Family: Les, Janet, Alex, and Stefan

Contents

List of Figures

List of Tables

Acknowledgements

First of all, I would like to thank my committee chair, Jeffrey Russell. Without him, none of this would have been possible. His extensive support and advice over the last few years has helped me grow tremendously as a researcher. As well, his keen ability to bridge the statistics and finance worlds has allowed me to carve a niche where my interests fit perfectly. He made every step of this adventure fun, while always encouraging me to constantly develop my skills and interests.

Just as importantly, I want to thank the rest of my committee, Nicholas Polson, Matt Taddy, Ralph Koijen, and Bryan Kelly for all their guidance at every point over the past few years. I attribute much of my love for the Bayesian world of statistics to Nick's influences. He was a major part of my decision to pursue a PhD at Chicago Booth, and has always been a valuable mentor through all of our coffee shop meetings. From the moment I arrived, Matt welcomed me into the group and gave me an open forum to talk about anything and ask any question, even potentially stupid ones! This unquestionably contributed largely to my confidence and ability to succeed. Ralph is to thank for his insight into the finance world and giving direction on how to produce well developed finance research. Bryan's advice and perspective in finance has also been invaluable, especially in developing my ideas into more cogent and complete propositions. I consider all of you close friends, and cannot thank you all enough for everything.

Without a doubt, I give considerable thanks to former dean, Edward Snyder. His excitement and love of the University of Chicago was one of the main reasons why I decided to initially attend the school's MBA program, and eventually pursue a PhD here.

I am also grateful to Barclays Capital Indices for providing access to their fixed income index database. Their generosity has made much of this work possible.

My family has played an invaluable role through every part of my years of education and research. I thank my parents, Les and Janet who allowed me to pursue the educational goals I so arduously laid out for myself. As well, I am indebted to my wonderful brothers, Alex and Stefan who on innumerable occasions were able to help me see the big picture when I, so frequently, was concentrated on such a focused perspective.

Finally, I want to acknowledge my grandfather, Carmine Iannacone Sr., who, although he passed on before my years of college education, was one of the most momentous influences on me. A retired New York City police officer and teacher, he always encouraged education and made learning an exciting endeavor. His inspiring support and encouragement is missed, but will always be remembered.

Abstract

This dissertation is a compilation of three papers developing a class of compositional state space models for modeling a latent set of time-varying portfolio compositions on the simplex as well as net leverage values given a time series of portfolio return observations. Estimation techniques incorporating particle filtering and particle learning concepts are exhibited. These estimation techniques are motivated by the estimation of asset class weights and net leverage values on the aggregate hedge fund industry from 1995 to 2012.

In the first paper, I present a compositional state space model for estimation of an investment portfolio's unobserved asset allocation weightings on a set of candidate assets when the only observed information is the time series of portfolio returns and the candidate asset returns. I exhibit both sequential Monte Carlo numerical and conditionally Normal analytical approaches to solve for estimates of the unobserved asset weight time series. Furthermore, I show how to implement the results as predictive investment weightings in order to construct hedge fund replicating portfolios.

In the second paper, I demonstrate how to implement joint estimation of net portfolio leverage dynamics into the previous paper's setup. By incorporating recent work in parameter learning in state space models, I also show how to not only sequentially estimate the time-varying latent values, but also the parametrization of their generative distributions. Using this technique, I estimate net portfolio leverage on a set of hedge fund indices representing the returns on different broad classifications

of funds. Finally, I exhibit the accuracy of these techniques by estimating asset class level regressions on the asset class return values against the same-period changes in the portfolio investment weights in order to demonstrate investment effects on same-period prices.

In the third paper, I identify that since the complete picture of hedge fund holdings is not observable, this presents a significant analytical hurdle for more detailed analysis because while the average hedge fund does not outperform benchmarks on the whole, it is possible that they exhibit skill in certain asset classes. Using the previously developed decomposition techniques, I discover that net leverage levels in the hedge fund industry are smaller than popular belief due to netting both internally and across different funds. As well, using these estimates, we confirm previous findings that hedge funds do not contribute to herding behavior in most asset classes, and in fact exhibit negative-feedback trading behavior in oil and municipal bonds. As well, the accuracy of these techniques is demonstrated on a set of actively managed diversified equity mutual funds where true industry allocation compositions are readily observable.

1

The Dirichlet Portfolio Model: Uncovering the Hidden Composition of Hedge Fund Investments

1.1 Abstract

Hedge funds have long been viewed as a veritable "black box" of investing since outsiders may never view the exact composition of portfolio holdings. Therefore, the ability to estimate an informative set of asset weights is highly desirable for analysis. We present a compositional state space model for estimation of an investment portfolio's unobserved asset allocation weightings on a set of candidate assets when the only observed information is the time series of portfolio returns and the candidate asset returns. In this paper, we exhibit both sequential Monte Carlo numerical and conditionally Normal analytical approaches to solve for estimates of the unobserved asset weight time series. This methodology is motivated by the estimation of monthly asset class weights on the aggregate hedge fund industry from 1996 to 2012. Furthermore, we show how to implement the results as predictive investment

weightings in order to construct hedge fund replicating portfolios.

1.2 Introduction

Over the past 20 years, the financial world has seen an enormous increase in demand for hedge fund products, thereby contributing to the estimated size of this global industry at approximately $2.13 trillion as of April 1, 2012, according to Hedge Fund Research (HFRI). These products intend to not only maximize returns on the assets under management during times of market boom, but also protect against losses during economic downturns.

High demand for access to these hedge fund products is manifested in the high value of fees charged to investors. On average, these fees come in the form of a 1-2% management fee assessed on the total assets under management, in addition to a 15-25% incentive fee on all capital gains. Since it can be difficult for investors to assess the spectrum of individual hedge fund managers' skill, most investors tend to make their investments though a vehicle called a "fund of funds". The purpose of these intermediaries is to evaluate individual hedge fund managers and then allocate an investors' assets across a broad spectrum of managers. This division of investments is intended to diversify away risk associated with individual managers, and instead provide exposure to the returns of the hedge fund industry as a whole. For these fund of funds services, a median 1.5% management fee plus a 10% capital gains fee is charged on top of the existing individual managers' fees. These fees quickly add up and can easily eat away at any real profits that arise from capital appreciation on the invested funds. Due to the combination of these high fee schedule layers, it is desirable to decompose and analyze the investment portfolios of these funds to determine if they are truly adding value for investors, or if similar strategies can be constructed with a much lower cost of investment.

As well, investors' investment goals and tolerance for risk exposures may not align with the incentive structure of hedge fund and fund-of-fund managers. Therefore, a decomposition of a hedge fund's exposures to the risks arising from various asset classes is desired. Of particular interest is how hedge fund managers respond to various macroeconomic events. Using the model and estimation methodology presented here, we obtain a decomposition of the hedge fund industry's asset class risk exposures, which provide insight into their asset allocation process. Interestingly, we find large increases in exposure to municipal bonds during the Dot-com Bubble decline in 2000-2001 and the recent global financial crisis from 2007-2012.

Another important feature of the hedge fund investing world is that return performance is only reported on a monthly or even quarterly basis. Therefore, we can only observe how the fund has performed at certain discrete dates. Between those dates, we cannot observe the current state of an individual's invested capital. An investor could have doubled their money, or even lost half of their wealth overnight, but they will not know until the next reporting period. If an investor had access to the invested asset value weightings, then they could compute estimates for intraperiod return and volatility values. These values can have very important implications for current consumption choices, as well as risk management decisions.

This directly leads to a number of questions: Can we estimate what hedge funds are invested in, as well as how that asset allocation changes over time? Then, using these estimates of asset allocation, can we generate intra-reporting-period return and volatility estimates? That is, can we estimate how hedge funds are performing on a daily or even second-by-second granularity? Furthermore, can we replicate this hedge fund portfolio in order to produce a similar series of returns, but through investing in easily accessible assets?

This portfolio estimation setup suggests a state space estimation problem where the latent compositional weights are required to sum to 1. Due to this restriction, we venture beyond the classic Kalman Filter solution to estimate the weights (Kalman, 1960). The results from Chipman and Rao (1964) and Tintner (1952) on Constrained Least Squares (CLS) estimation allow for this restriction in static models. However, the CLS model does not allow for a dynamic compositional weight process. In response, Chia (1985) and Simon and Chia (2002) present a solution with this restriction for dynamic models. However, we show that these techniques do not perform well for this application. Other notable work in compositional time series models are presented in Grunwald, Raftery, and Guttorp (1993) and Cargnoni, Müller, and West (1997). These focus on multinomial observational models of pure proportions or multinomial counts. This work is in the spirit of those results, although we focus on univariate observations arising from a transformation of the latent compositional values. The form of the generative dynamic model is as follows:

$$w_t \sim Dir\left(\alpha \frac{w_{t-1} \circ (1 + r_{PA,t-1})}{\sum_{i=1}^{n} w_{t-1,i}(1 + r_{PA,t-1,i})}\right)$$

$$r_{HF,t} \sim t(w_t' r_{PA,t}, \sigma_\epsilon^2, \nu)$$

Our approach is to use the particle filtering methodology of Gordon, Salmond, and Smith (1993) to numerically solve the estimation problem on the portfolio weights w_t. Also, making use of the particle filtering methods allows us to venture outside the simple Gaussian observational error assumption, thereby giving more suitable estimation results.

The remainder of this paper is structured as follows: Section 2 describes the statistical setup and motivation of the basic dynamic model. Section 3 presents the fully specified Dirichlet Portfolio Model (DPM), the Sequential Monte Carlo approach for solving it, as well as an analytically solvable conditionally Normal approximation. Section 4 outlines previous approaches to solve this problem, their respective drawbacks, and proposed improvements. Section 5 compares the DPM to the other approaches under simulated and model hedge fund trading environments. Section 6 uses actual hedge fund return data to estimate latent investment weights. Section 7 outlines how these results can be used to estimate intraperiod hedge fund return and volatility values, as well as construct hedge fund replicating portfolios. Finally, section 9 concludes.

1.3 The Tracking & Filtering Problem

1.3.1 General Setup

First, consider the problem of estimating the latent weights on individual asset classes. That is, at each time period we desire to combine prior information on these weights with new information introduced through observed overall hedge fund performance given the contemporaneous performance on the asset classes of interest. This leads to defining a dynamic state space model of the following form:

$$r_{HF,t} = f(w_t, r_{PA,t})$$

$$w_t = g(\mathcal{F}_{t-1})$$

where $\mathcal{F}_t$ is the filtering of all information known at time t. Hence, this includes all previous hedge fund index returns r_{HF}, palette asset returns r_{PA}, and palette asset weights w up to and including time t. That is,

$$\mathcal{F}_t = \{r_{HF,1}, \ldots, r_{HF,t}, r_{PA,1}, \ldots, r_{PA,t}, w_1, \ldots, w_t\}.$$

Note that $w_t = (w_{t,1}, w_{t,2}, ..., w_{t,n})'$ is a $n \times 1$ vector of the weights on each asset at the beginning of time period t, $r_{PA,t} = (r_{PA,t,1}, r_{PA,t,2}, .., r_{PA,t,n})'$ is a $n \times 1$ vector of the palette asset returns over time period t, and $r_{HF,t}$ is a scalar value of the return on the hedge fund index over the same time period t. The chronology of the time period notation is illustrated below:

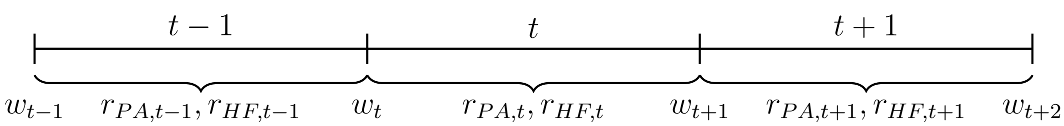

FIGURE 1.1: *Timeline Notation Illustration*

We can define a form for the observation function $f(\cdot)$. Since the aggregate return on a portfolio of assets in simply the sum of the value weighted returns on the assets, the observation equation can be written as follows:

$$r_{HF,t} = w_t' r_{PA,t} + \epsilon_t$$

where ϵ is a stochastic term to be given a distributional form later. This term is very important because it is unrealistic and potentially impossible to include all possible assets that a portfolio may be comprised of. Therefore, it is necessary to allow for a term to pick up the variation in the observed index returns which is orthogonal to the palette asset returns.

Determining the form of the transition function, $g(\cdot)$, is a bit more challenging. There may not necessarily be an exact science of how portfolio managers transition their asset weightings from period to period, but Amenc, Martellini, Meyfredi, and Ziemann (2010) suggest the following property:

$$E[w_t|\mathcal{F}_{t-1}] = w_{t-1}$$

This suggests that on average, portfolio managers keep the same asset value weightings from period to period. This is a reasonable assumption, however it does introduce a subtle problem. To illustrate this, take for example a portfolio of 2 assets where both are initially given equal value weighting (i.e. 50% each). Now, suppose that asset 1 yields a return of 0% and asset 2 yields a return of 100% over a given time period. Due to capital appreciation, assets 1 and 2 now have value weightings of 33.3% and 66.7%, respectively. If the above property were employed in creating a transitional distribution, the prior expectation of the asset weights would both be 50% (hence completely ignoring the idea of capital appreciation/depreciation). This property hereby causes an artificial "mean reversion" effect on the asset weights since assets with relatively high return performance will be forced to have a relatively low prior in the next period, and vice versa.

Since an estimation procedure is desired which does not favor a "mean reversion" effect over a "momentum" effect, it is much more intuitive to implement a true random-walk process for the weights, which is what Amenc et al likely intended. Since it is unlikely that the aggregate universe of portfolio managers consistently employs an asset allocation strategy which ignores capital appreciation, this is economically reasonable as well. In order to account for capital appreciation and depreciation, the previous period's weight estimates, w_{t-1} are updated by the relative increase in the observed period $t-1$ asset returns, $r_{PA,t-1}$. This gives the following property:

$$E[w_t|\mathcal{F}_{t-1}] = \frac{w_{t-1} \circ (1 + r_{PA,t-1})}{\sum_{i=1}^{n} w_{t-1,i}(1 + r_{PA,t-1,i})} \tag{1.1}$$

where $\circ$ is the Hadamard product.

It is important to note that the aggregate size of the hedge fund industry is about $2 trillion. Due to the very large nature of this aggregate hedge fund portfolio, it is unlikely that the entire industry could make major rebalancing shifts in the asset class weights from period to period. That is, it would be very unlikely, and incredibly difficult for the entire industry to consistently employ either a "mean reversion" or "momentum" style strategy. This is further, and potentially stronger support for the above property.

A stochastic component η is incorporated into the weight transition in order to allow for the period-to-period uncertainty about transitional changes in the weights. A distributional form will be imposed on this as well.

The general dynamic state space model for this problem is written:

$$r_{HF,t} = w_t' r_{PA,t} + \epsilon_t \tag{1.2}$$

$$w_t = \frac{w_{t-1} \circ (1 + r_{PA,t-1})}{\sum_{i=1}^{n} w_{t-1,i}(1 + r_{PA,t-1,i})} + \eta_t \tag{1.3}$$

$$s.t. \quad \sum_{i=1}^{n} w_t = 1$$

Note that in this model there is no non-negativity constraint on the weights. Negative values would imply a "short" weight on a palette asset.

1.3.2 Estimation

We are interested in solving for estimates of the asset weights conditional on all information available up to and including the current period.

First, because of the Markov property of the model, the true weights at time t can be written as conditionally independent of all earlier times given information in the previous time $t-1$:

$$p(w_t | w_0, ..., w_{t-1}, r_{PA,0}, ..., r_{PA,t-1}) = p(w_t | w_{t-1}, r_{PA,t-1})$$

As well, the observation model at time t is conditionally independent of all earlier times given information in the current time t:

$$p(r_{HF,t} | w_0, ..., w_t, r_{PA,0}, ..., r_{PA,t}) = p(r_{HF,t} | w_t, r_{PA,t})$$

Therefore, the probability distribution over all states in the model is:

$$p(w_0, ..., w_t, r_{HF,1}, ..., r_{HF,t}) = p(w_0) \prod_{\tau=1}^{t} p(r_{HF,\tau}|w_\tau, r_{PA,\tau}) p(w_\tau|w_{\tau-1}, r_{PA,\tau-1})$$

In order to estimate the weights w_t conditional on the information up to the current time t we simply need to marginalize out the previous time periods. Bayes rule gives the following expression:

$$p(w_t|r_{HF,t}, r_{PA,t}) \propto \underbrace{p(r_{HF,t}|w_t, r_{PA,t})}_{Likelihood} \underbrace{p(w_t|r_{HF,t-1}, r_{PA,t-1})}_{Prior}$$

This is the "update" step, where the prior on the weights at the current time period is given by:

$$p(w_t|r_{HF,t-1}, r_{PA,t-1}) = \int p(w_t|w_{t-1}, r_{PA,t-1}) p(w_{t-1}|\mathcal{F}_{t-1}) dw_{t-1}$$

The estimates of interest are obtained, $p(w_t|\mathcal{F}_t)$.

Under conditions of linearity and Normality this problem can be solved analytically with the Kalman filter. However, if either of those conditions are violated, then the above densities are intractable and therefore approximate inference must be employed via Sequential Monte Carlo Methods.

1.3.3 Prediction

At a given point in time t, the posterior predictive distribution is used as our next period forecast. This is given by:

$$p(w_{t+1}|w_t, r_{PA,t})$$

This is the same distribution as the prior for the next step ahead estimation problem. This prediction problem is of special interest since access is not available to the aggregate of “true” hedge fund industry asset allocations, and therefore if we are to believe the assumption that the aggregate hedge fund industry does not (and possibly cannot, due to its large size) change asset class weightings very quickly, then the “predictive” return accuracy will give insight into how accurate the estimation technique is when using real world data.

As well, by constructing an estimation method for the relative portfolio weights, then these estimated latent weights can be used to project what funds may be invested in at any point in the future. This gives the ability to estimate a distribution of potential latent intraperiod returns at any of these points:

$$p(r_{HF,t}|r_{PA,t}, \mathcal{F}_{t-1}) = p(w_t|r_{PA,t-1}, r_{HF,t-1})' r_{PA,t}$$

1.4 Sequential Monte Carlo Approach

Herein, the proposed general dynamic model for this state space problem will have distributional forms imposed on the stochastic nature of the expressions in order to develop a feasible model for estimation.

1.4.1 Model

First, the weight transition model (1.3) is considered. Recall the existence of the budgetary restriction $\sum_{i=1}^{n} w_t = 1$ on the relative portfolio weights. Although it may be easy for an individual hedge fund to take a short position on an asset class, it is very hard for the aggregate of all \$2+ trillion worth of hedge funds to take a net short position on some asset class. Therefore, we impose the restriction that asset class weightings may not take on a negative value, $w_{t,i} \geq 0$ for $i \in \{1, ..., n\}$. This suggests the use of a Dirichlet distribution for the weight transitions:

$$w_t \sim Dir\left(\alpha \frac{w_{t-1} \circ (1 + r_{PA,t-1})}{\sum_{i=1}^{n} w_{t-1,i}(1 + r_{PA,t-1,i})}\right) \tag{1.4}$$

where α is a scalar concentration parameter controlling how much the aggregate hedge fund industry changes its investment weightings each period. Notice how this satisfies the desired property suggested in (1.1):

$$E[w_t|\mathcal{F}_{t-1}] = \frac{\boldsymbol{\alpha}}{\boldsymbol{\alpha}_0} = \frac{w_{t-1} \circ (1 + r_{PA,t-1})}{\sum_{i=1}^{n} w_{t-1,i}(1 + r_{PA,t-1,i})}$$

where

$$\boldsymbol{\alpha} \equiv \alpha \frac{w_{t-1} \circ (1 + r_{PA,t-1})}{\sum_{i=1}^{n} w_{t-1,i}(1 + r_{PA,t-1,i})}$$

and

$$\boldsymbol{\alpha}_0 \equiv \sum_{j=1}^{n} \boldsymbol{\alpha}_j = \sum_{j=1}^{n} \alpha \frac{w_{t-1,j}(1 + r_{PA,t-1,j})}{\sum_{i=1}^{n} w_{t-1,i}(1 + r_{PA,t-1,i})} = \alpha$$

Consider the observation model (1.2). The purpose of the parameter ϵ_t is to pick up the variation in the hedge fund index returns $r_{HF,t}$ which is orthogonal to the palette asset returns $r_{PA,t}$, it is appropriate to consider a leptokurtic distribution due to the fat-tail property commonly exhibited by financial data first noted by Mandelbrot (1963). Therefore, the following scale-location Student-t model is used in our analysis:

$$r_{HF,t} \sim t(w_t' r_{PA,t}, \sigma_\epsilon^2, \nu) \tag{1.5}$$

By combining expressions (1.4) and (1.5), our dynamic model is completely defined to form the foundational Dirichlet Portfolio Model (DPM):

$$w_t \sim Dir\left(\alpha \frac{w_{t-1} \circ (1 + r_{PA,t-1})}{\sum_{i=1}^{n} w_{t-1,i}(1 + r_{PA,t-1,i})}\right)$$

$$r_{HF,t} \sim t(w_t' r_{PA,t}, \sigma_\epsilon^2, \nu)$$

This model can be used for estimation of latent asset weights in any portfolio where we are interested in the dynamics of weights changes due to active trading decisions.

1.4.2 Filtering Method

The use of a Dirichlet transition model, as well as the Student-t observation model has ruled out analytical solutions to this problem. Therefore, a Sequential Monte Carlo simulation technique is used to estimate the palette asset weights.

First, a prior distribution is placed over the initial palette asset weights:

$$w_0 \sim Dir\left(\alpha_0 \left(\frac{1}{n}, ..., \frac{1}{n}\right)'\right)$$

where α_0 is a scalar concentration parameter controlling initial uncertainty about the prior weight distribution. Each asset is given equal weight in expectation with the lack of better information. Simulating from this distribution gives a set of particles $\{w_0\}_p$, $p \in \{1, ..., P\}$, characterizing the approximation. Then, at each point in time we iteratively propagate and resample as defined in the Sequential Importance Resampling (SIR) algorithm of Rubin (1987) and Smith and Gelfand (1992). This process produces the weight distribution estimates of interest.

First, consider the propagation step. Let there exist a set of particles $\{w_{t-1}\}_p$ representing the distribution $p(w_{t-1}|r_{HF,t-1}, r_{PA,t-1})$ from a previous iteration. In order to find the prior distribution for the asset weights at time t, $p(w_t|w_{t-1}, r_{PA,t-1})$, draws from the transition model for each particle $\{w_{t-1}\}_p$ are made. This yields a set of particles $\{w_{t|t-1}\}_p$ approximating this prior distribution. Second, consider the resampling step. Given the set of particles $\{w_{t|t-1}\}_p$ approximating the prior distribution, they are resampled with respect to their relative likelihoods given by $\omega_t = p(r_{HF,t}|w_t, r_{PA,t})$. This set of resampled particles $\{w_t\}_p \equiv \{w_{t|t}\}_p$ will therefore approximate the desired distribution $p(w_t|r_{HF,t}, r_{PA,t})$.

Using these results, forms for all of the probability distributions in the SIR algorithm are fully specified. First, the 'Step 1' propagation step is defined:

$$p(w_t|w_{t-1}^{(p)}, r_{PA,t-1}) \equiv Dir\left(\alpha \frac{w_{t-1}^{(p)} \circ (1 + r_{PA,t-1})}{\sum_{i=1}^{n} w_{t-1,i}^{(p)}(1 + r_{PA,t-1,i})}\right)$$

Second, the 'Step 2' normalized importance weights are computed from the observation model $r_{HF,t} \sim t(w_t' r_{PA,t}, \sigma_\epsilon^2, \nu)$ where

$$p(r_{HF,t}|w_t, r_{PA,t}) = \frac{\Gamma(\frac{\nu+1}{2})}{\Gamma(\frac{\nu}{2})\sqrt{\pi\nu\sigma_\epsilon^2}}\left(1 + \frac{1}{\nu}\frac{(r_{HF,t} - w_t' r_{PA,t})^2}{\sigma_\epsilon^2}\right)^{-\frac{\nu+1}{2}}$$

Therefore, the importance weights are given by:

$$\omega_t^{(p)} \equiv \frac{p(r_{HF,t}|w_t^{(p)}, r_{PA,t})}{\sum_{\phi=1}^{P} p(r_{HF,t}|w_t^{(\phi)}, r_{PA,t})} = \frac{\left(\nu\sigma_\epsilon^2 + \left(r_{HF,t} - w_t^{(p)\prime} r_{PA,t}\right)^2\right)^{-\frac{\nu+1}{2}}}{\sum_{\phi=1}^{P}\left(\nu\sigma_\epsilon^2 + \left(r_{HF,t} - w_t^{(\phi)\prime} r_{PA,t}\right)^2\right)^{-\frac{\nu+1}{2}}}$$

Finally, the initial palette asset weight distribution is set:

$$p(w_0) \equiv Dir\left(\alpha_0\left(\frac{1}{n}, ..., \frac{1}{n}\right)'\right)$$

Note that if there exists better information about the distribution $p(w_0)$, the use of that will naturally lead to superior and more appropriate results.

Now, we substitute in these developed forms for the Dirichlet Portfolio Model to get the fully specified DPM Estimation Algorithm in Figure 1.2.

1.4.3 Conditionally Normal Approximation

Due to the non-Gaussian nature of this multivariate compositional model, the posterior distributions of the asset weights cannot be solved for in analytical closed form. Therefore, sequential Monte Carlo methods are employed to numerically approximate

DPM Estimation Algorithm

Initialize: Sample prior weights from

$$w_0^{(p)} \sim Dir\left(\alpha_0 \left(\frac{1}{n}, ..., \frac{1}{n}\right)'\right)$$

Iterate:

Step 1: Propagate new asset weights from

$$w_t^{(p)} \sim Dir\left(\alpha \frac{w_{t-1}^{(p)} \circ (1 + r_{PA,t-1})}{\sum_{i=1}^{n} w_{t-1,i}^{(p)}(1 + r_{PA,t-1,i})}\right) \text{ for } p = 1, ..., P$$

Step 2: Resample asset weights from

$$w_t^{(p)} \sim Mult_P\left(\left\{\omega_t^{(\phi)}, w_t^{(\phi)}\right\}_{\phi=1}^{P}\right)$$

where the normalized importance weights are given by

$$\omega_t^{(p)} = \frac{\left(\nu\sigma_\epsilon^2 + \left(r_{HF,t} - w_t^{(p)\prime} r_{PA,t}\right)^2\right)^{-\frac{\nu+1}{2}}}{\sum_{\phi=1}^{P}\left(\nu\sigma_\epsilon^2 + \left(r_{HF,t} - w_t^{(\phi)\prime} r_{PA,t}\right)^2\right)^{-\frac{\nu+1}{2}}}$$

FIGURE 1.2: *DPM Estimation Algorithm*

these distributions. Another approach to solving this problem is to approximate the Dirichlet errors by a multivariate Gaussian distribution replicating the first two moments at each step time. So, just as was done in the above solution, the error distributions of the transitions must be reparametrized at each time step based upon the estimation results from the previous step.

First, consider the transition model from the DPM in (1.4). It can be shown that the first two moments of w_t are:

$$\mu_t \equiv E[w_t|\mathcal{F}_{t-1}] = \frac{w_{t-1} \circ (1 + r_{PA,t-1})}{\sum_{i=1}^{n} w_{t-1,i}(1 + r_{PA,t-1,i})}$$

and

$$Cov[w_{t,i}, w_{t,j}|\mathcal{F}_{t-1}] = \frac{\boldsymbol{\alpha}_i(\boldsymbol{\alpha}_0 I_{i=j} - \boldsymbol{\alpha}_j)}{\boldsymbol{\alpha}_0^2(\boldsymbol{\alpha}_0 + 1)}$$

where

$$\boldsymbol{\alpha}_i \equiv \alpha \frac{w_{t-1,i}(1 + r_{PA,t-1,i})}{\sum_{i=1}^{n} w_{t-1,i}(1 + r_{PA,t-1,i})} \quad \text{and} \quad \boldsymbol{\alpha}_0 = \sum_i \boldsymbol{\alpha}_i = \alpha$$

from above. So, it can be shown that:

$$\Sigma_{t,i,j} \equiv Cov[w_{t,i}, w_{t,j}|\mathcal{F}_{t-1}] = \frac{w_{t-1,i}(1 + r_{PA,t-1,i})\left(\xi_{t-1} I_{i=j} - w_{t-1,j}(1 + r_{PA,t-1,j})\right)}{\xi_{t-1}(\alpha + 1)}$$

and

$$\xi_{t-1} = \sum_{k=1}^{n} w_{t-1,k}(1 + r_{PA,t-1,k})$$

Then, using this, the original DPM Model can be approximately rewritten into the Conditionally Normal Dirichlet Portfolio Model (CN-DPM) assuming a Gaussian observational distribution:

$$w_t \sim N\left(\mu_t, \Sigma_t\right)$$

$$\mu_t = \frac{w_{t-1} \circ (1 + r_{PA,t-1})}{\sum_{i=1}^{n} w_{t-1,i}(1 + r_{PA,t-1,i})}$$

$$\Sigma_t = \left[\frac{w_{t-1,i}(1 + r_{PA,t-1,i})\left(\xi_{t-1} I_{i=j} - w_{t-1,j}(1 + r_{PA,t-1,j})\right)}{\xi_{t-1}(\alpha + 1)}\right]_{i,j}$$

$$r_{HF,t} \sim N(w_t' r_{PA,t}, \sigma_\epsilon^2)$$

Note that although the observational distribution may not be best modeled by a Gaussian form, it is a necessary simplification to use the results from Kalman (1960) to solve the dynamic model analytically. Using the above form, the latent weights are solved for in the Conditionally Normal Dirichlet Portfolio Model in Figure 1.3.

CN-DPM Estimation Algorithm

Initialize: Set initial weight distribution

$$p(w_0) = N(\mu_0, \Sigma_0)$$

Iterate:

Step 1: Compute the prior weight distribution using the transition model

$$p(w_t|\mathcal{F}_{t-1}) = N\left(\mu_{t|t-1}, \Sigma_{t|t-1}\right)$$

$$\mu_{t|t-1} = \frac{w_{t-1} \circ (1 + r_{PA,t-1})}{\sum_{i=1}^{n} w_{t-1,i}(1 + r_{PA,t-1,i})}$$

$$\Sigma_{t|t-1} = \Sigma_{t-1|t-1} + \left[\frac{w_{t-1,i}(1 + r_{PA,t-1,i})\left(\xi_{t-1}I_{i=j} - w_{t-1,j}(1 + r_{PA,t-1,j})\right)}{\xi_{t-1}(\alpha + 1)}\right]_{i,j}$$

$$\xi_{t-1} = \sum_{k=1}^{n} w_{t-1,k}(1 + r_{PA,t-1,k})$$

Step 2: Compute the posterior weight distribution using the observation model

$$p(w_t|\mathcal{F}_t) = N\left(\mu_{t|t}, \Sigma_{t|t}\right)$$

$$\mu_{t|t} = \mu_{t|t-1} + K_t(r_{HF,t} - \mu'_{t|t-1}r_{PA,t})$$

$$\Sigma_{t|t} = \left(I - K_t r'_{PA,t}\right)\Sigma_{t|t-1}$$

where the optimal Kalman Gain value is given by

$$K_t = \left(\Sigma_{t|t-1}r_{PA,t}\right)\left(r'_{PA,t}\Sigma_{t|t-1}r_{PA,t} + \sigma_\epsilon^2\right)^{-1}$$

FIGURE 1.3: *CN-DPM Estimation Algorithm*

1.5 Alternative Approaches

We briefly review some alternative estimation techniques that either have been used, or could be used similarly to the Dirichlet Portfolio Model.

1.5.1 Rolling Window OLS

The general setup of these regression models is as follows:

$$r_{HF,\tau} = w_t' r_{PA,\tau} + \epsilon_\tau, \quad \epsilon \sim N(0, \sigma^2)$$

where $\tau \in \mathrm{T} \equiv \{t-k, t-k+1, ..., t\}$ for a k-sized window. The classic OLS estimator is given by $\hat{w}_t = \left(r_{PA,\mathrm{T}}' r_{PA,\mathrm{T}}\right)^{-1} r_{PA,\mathrm{T}}' r_{HF,\mathrm{T}}$. While this is a simplified first approach, it suffers from some major problems. First, this is a static model for the estimated asset weights, and therefore makes the incorrect assumption that the weights are constant over the estimation time period. Second, a window size k must be chosen, therefore having to deal with the trade-off of using more data to obtain better estimates but decreasing the relative importance of more recent observations. Lastly, the most apparent problem is the lack of the $\sum_{i=1}^{n} w_{t,i} = 1$ restriction. Nevertheless, a more appropriate estimator using this portfolio normalization constraint can be constructed. We have the following setup:

$$\bar{w}_t = \underset{w_t}{\operatorname{argmin}} \sum_{\tau \in \mathrm{T}} \left(r_{HF,\tau} - w_t' r_{PA,\tau}\right)^2 \quad \text{where} \quad \sum_{i=1}^{n} w_{t,i} = 1$$

This can be solved by Constrained Least Squares (CLS) from Chipman and Rao (1964) and Tintner (1952). In the context of this problem, estimates are obtained by:

$$\bar{w}_t = \hat{w}_t - \left(r_{PA,\mathrm{T}}' r_{PA,\mathrm{T}}\right)^{-1} \mathbf{1} \left(\mathbf{1}' \left(r_{PA,\mathrm{T}}' r_{PA,\mathrm{T}}\right)^{-1} \mathbf{1}\right)^{-1} \left(\mathbf{1}' \hat{w}_t - 1\right)$$

Note that although the above solution does place a normalizing restriction on the sum of the estimated weights, it still allows for individual weights to take any real value. That is, an estimated weight of 120 or -80 could be obtained, thereby implying an unrealistic 12,000% or -8,000% weight on that asset class. This explosive scaling effect happens widely in the presence of multicollinearity in the explanatory variables. Due to the prevalence of this in asset returns, the undesired scaling issue can be avoided by imposing the "no short selling" assumption of the DPM on the above CLS. That is, constrain the CLS with non-negativity: $w_{t,i} \geq 0, \quad \forall i \in \{1, \ldots, n\}$. Let us refer to this setup as the Inequality Constrained Least Squares (ICLS) method similar to Judge and Takayana (1966) and Liew (1976). This is easily solved via quadratic optimization.

Herein, the original OLS rolling regression approach will not be considered due to its gross misspecification for this problem. Instead, the CLS and ICLS approaches will be explored due to their increased suitability.

1.5.2 Naïve Kalman Filtering

Amenc et al outline an approach for using Bayesian inference to solve the dynamic state space model. The model is set up as follows:

$$r_{HF,t} = w_t' r_{PA,t} + \epsilon_t, \quad \epsilon_t \sim N(0, \sigma_\epsilon^2)$$

$$w_t = w_{t-1} + \eta_t, \quad \eta_t \sim N(0, Q)$$

This is a classic state space model which is analytically solvable via the Kalman Filter. Although this approach correctly identifies the problem as a dynamic model, it lacks the portfolio normalization constraint $\sum_{i=1}^n w_{t,i} = 1$. So, similar to the least-squares based techniques detailed above, we propose more suitable methods by adding in this constraint.

1. **Constrained Kalman Filtering via Restricted Covariance Structure**

   Consider the constraint $\sum_{i=1}^n w_{t,i} = 1$. It can be shown that

$$\sum_{i=1}^{n} Cov\left(w_{t,j}, w_{t,i}\right) = 0, \quad \forall j \in \{1, \ldots, n\}$$

Therefore, we can obtain a constrained estimation structure by choosing the initial weight uncertainty matrix Σ_0 and transition innovation matrix Q such that:

$$\Sigma_0\mathbf{1} = \mathbf{0},\, Q\mathbf{1} = \mathbf{0} \quad \text{and by symmetry} \quad \Sigma_0'\mathbf{1} = \mathbf{0},\, Q'\mathbf{1} = \mathbf{0} \tag{1.6}$$

Now, with the lack of better information, assign to $Var(w_{t,i})$ the average unconditional variance implied by the α parameter from the DPM for consistency purposes. Then, let

$$Cov(w_{t,i}, w_{t,j}) = -\frac{Var(w_{t,i})}{n-1}, \quad \forall j \neq i,$$

thereby satisfying the above restriction and creating a spherical covariance structure. Nevertheless, if there does exist information about a more suitable covariance structure, but it does not satisfy the above properties, we can project the original covariance matrices onto the space satisfying those constraints:

$$\Sigma_0^P = \left(I - \Sigma_0\mathbf{1}\left(\mathbf{1}'\Sigma_0\mathbf{1}\right)^{-1}\mathbf{1}'\right)\Sigma_0 \quad \text{and} \quad Q^P = \left(I - Q\mathbf{1}\left(\mathbf{1}'Q\mathbf{1}\right)^{-1}\mathbf{1}'\right)Q$$

2. **Constrained Kalman Filtering via State Projection**

Chia (1985) and Simon and Chia (2002) detail a method to first derive the unconstrained state estimate and then project it onto the constraint surface. This can easily be applied in the context of this application. When computing the posterior distribution of the weights, we can arrive at the projected distribution $N(\mu^P_{t|t}, \Sigma^P_{t|t})$ by first computing the unconstrained solution $N(\mu_{t|t}, \Sigma_{t|t})$ in the classic manner, and then projecting via:

$$\Upsilon_t = \Sigma_{t|t}\mathbf{1}\left(\mathbf{1}'\Sigma_{t|t}\mathbf{1}\right)^{-1} \qquad \mu^P_{t|t} = \mu_{t|t} - \Upsilon_t\left(\mathbf{1}'\mu_{t|t} - 1\right) \qquad \Sigma^P_{t|t} = \left(I - \Upsilon_t\mathbf{1}'\right)\Sigma_{t|t}$$

Conveniently, for calculating the prior distribution/forecasts for the weights, our transition function is already normalized with respect to the posterior weights from the previous period, thereby already projecting into the constrained space. Again, we assign to $Var(w_{t,i})$ the average unconditional variance implied by the α parameter from the DPM for consistency purposes.

Note that the estimation error covariance Σ_0 of this method is always going to be greater than or equal to that obtained by using the restricted covariance matrix method (Ko and Bitmead, 2007). This is because in the restricted covariance matrix method, the transition innovation covariance matrix Q is assumed to be the true process noise covariance, thereby resulting in the optimal state estimates for the system. However, in the projection method, the transition innovation covariance matrix Q may be inconsistent with the estimated transition innovation process. Nevertheless, if Q satisfies (1.6), then both these methods are equivalent.

Also, it is possible to incorporate inequality constraints into the Kalman Filtering approach as we did for the ICLS. Gupta and Hauser (2007) detail a method to do so using quadratic optimization. We do not implement this approach here since generally the Constrained Kalman solutions above produce estimation which is consistent with the desired non-negativity constraints, thereby negating the need to implement them in the estimation procedure.

Note that the CN-DPM presented above is analytically solvable via a modified Kalman filtering approach, however it incorporates an approximation for the Dirichlet compositional structure suggested by the DPM. This importantly requires redefined error distributions at each period. As well, note that the CN-DPM is a special case of the Constrained Kalman Filtering via Covariance Structure class of models due to it's compliance with the covariance restrictions. Furthermore, due to the period-by-period redefined error distributions, it also allows for the non-negativity constraint.

Again, due to the naïve Kalman Filter's misspecification for this problem, we only explore the Constrained Kalman via Covariance Structure (CKalCov) and State Projection (CKalProj) methods.

1.6 Simulated Portfolio Trading Comparison

We now compare the DPM and CN-DPM with the presented alternative estimation techniques on various simulated portfolio environments to motivate the effectiveness of the procedure.

1.6.1 Simulated Assets & Trading

First, sets of simulated monthly asset returns are developed under the following example model:

$$r_{i,t} \sim t(\mu_i, \sigma_i^2, \nu_i)$$

where

$$\mu_i \sim N(0.007, 0.003^2), \quad \sigma_i^2 \sim IG(2.5, 0.004),$$

$$\nu_i \sim IU(0, 0.5), \quad \Sigma_i \sim IW(\mathbf{I}, 8)$$

with contemporary correlation induced by a Gaussian copula having correlation implied by Σ_i. This example parametrization was motivated by Gelman and Hill (2006). Nevertheless, we will later demonstrate that the results also hold with real asset returns. For the following simulations, the case of 6 investable assets is considered.

Using these simulated assets, there are various ways to construct time series of portfolio weights. Let us first construct simulated portfolio asset weights using the previously motivated random-walk process from (1.4). Using these simulated weights, the simulated time series of hedge fund returns is constructed via $r_{HF,t} \sim$

$t(w_t' r_t, \sigma_\epsilon^2, \nu_{HF})$. A Student-t distribution is used here since it is reasonable to observe that returns which are orthogonal to the set of included explanatory assets can potentially be very "fat-tailed" in nature due to some trading activities such as market-making, high frequency trading, highly illiquid asset pricing, etc.

The objective is to track the weights on each of these assets, but due to the unobservable nature of the true portfolio weights, the accuracy of the weight predictions can never be observed in the real world setting. Therefore, to gain a proxy of how close these weights are being estimated, we can create the forecasted set of weights from the model and then determine how close we are to the one-step-ahead returns from the hedge fund index. That is, we want to minimize the following error:

$$\epsilon_t = r_{HF,t} - E[w_t|\mathcal{F}_{t-1}]' r_t$$

There are various measures of accuracy for this estimation. We explicitly define the following four measures for use throughout the remainder of the paper:

Measure Name	**Expression**
Forecasted Root Mean Squared Error (F-RMSE)	$\sqrt{\frac{1}{T}\sum_{t=1}^{T}(r_{HF,t} - E[w_t|\mathcal{F}_{t-1}]' r_t)^2}$
Forecasted Mean Absolute Error (F-MAE)	$\frac{1}{T}\sum_{t=1}^{T}|r_{HF,t} - E[w_t|\mathcal{F}_{t-1}]' r_t|$
Forecasted Pearson Correlation (F-Corr)	$corr\left(r_{HF,t}, E\left[w_t|\mathcal{F}_{t-1}\right]' r_t\right)$
Forecasted R^2 (F-R^2)	$1 - \frac{\sum_{t=1}^{T}\left(r_{HF,t} - E[w_t|\mathcal{F}_{t-1}]' r_t\right)^2}{\sum_{t=1}^{T}\left(r_{HF,t} - \bar{r}_{HF}\right)^2}$

Note that this F-RMSE value is exactly the same as the "tracking error" concept used commonly in portfolio management to describe how close the returns of a portfolio track to the returns of a given benchmark index. In our case, the benchmark index is simply the hedge fund index of interest.

We run 100 simulations and estimate the portfolio weights using the DPM, CN-DPM, CLS, ICLS, CKalCov, and CKalProj. First, we compare the simulation results using Forecasted Mean Absolute Error in Figure 1.4. The DPM and CN-DPM procedures produce smaller and more precise forecasted mean absolute deviation values, and therefore more accurate forecasted returns than the other methods. The constrained Kalman Filter methods, CKalCov and CKalProj, generally produce the next best best results, however, with a median forecasted MAE of 0.0082 and 0.0081 versus the DPM's value of 0.0057, the Kalman Filters perform about 43% worse. This is compared to using rolling window CLS and ICLS, which give median forecasted MAE values of 0.0085 and 0.0079, 49% and 39% worse than the DPM. As well, the CN-DPM has a median forecasted MAE of 0.0078, second to the DPM.

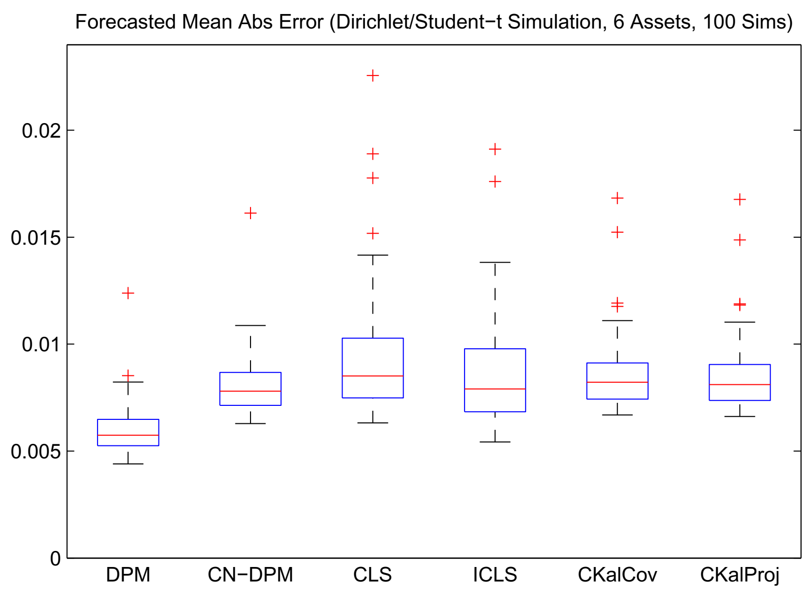

FIGURE 1.4: *Simulation Results – Forecasted Return MAE*

Another way to compare these estimation procedures is the Forecasted Coefficient of Determination, F-R^2. Naturally, this does not consider scale, as the MAE measure does, but it is useful to consider the proportion of variation in the hedge fund returns explained by the forecasted model. Examining the plot in Figure 1.5, the DPM produces much stronger forecasted R^2 values than the other methods, thus supporting its more accurate asset weight estimation. The DPM and CN-DPM produce median forecasted R^2 values of 0.979 and 0.965, respectively, as compared to rolling CLS and ICLS rolling with 0.938 and 0.944, and the CKalCov and CKalProj with 0.960 and 0.960. Similar plots can be created for the F-RMSE and F-Corr measures. These plots display similar results as those shown here.

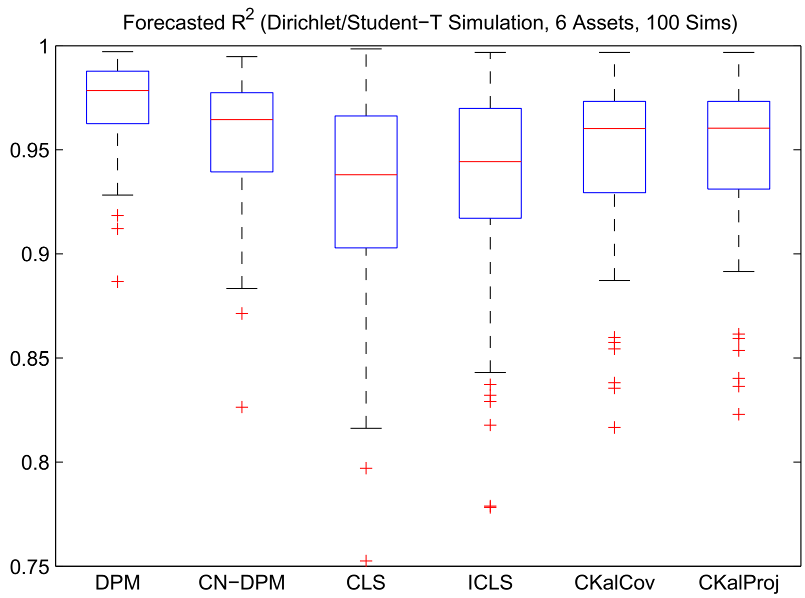

FIGURE 1.5: *Simulation Results – Forecasted Return* R^2

1.6.2 Increasing Number of Assets

Since the dimension of explanatory assets has the potential to grow large, we explore the effect of increasing the number of investable assets. Below, the same simulations are ran, but while increasing the number of investable assets in the simulated hedge fund index construction.

In Figure 1.6, the DPM procedure is used to estimate weights on the simulated assets, and then those weights are used to construct the forecasted return MAE values for each simulation. As expected, as the number of assets increases, the

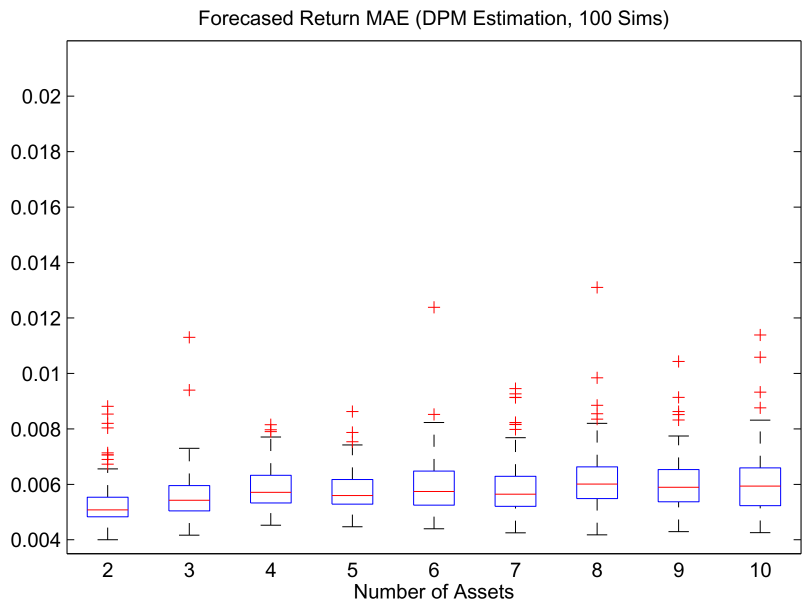

FIGURE 1.6: *DPM Forecasted Return MAE vs. Number of Assets*

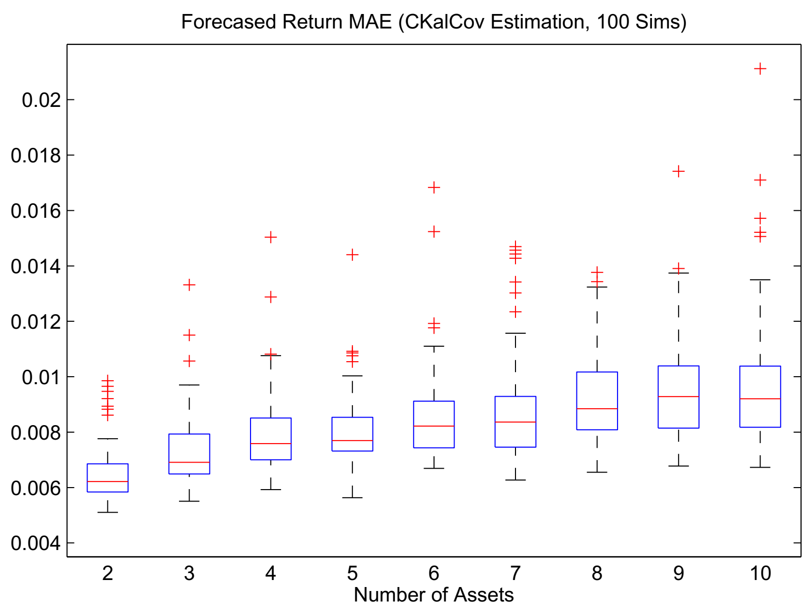

FIGURE 1.7: *CKalCov Forecasted Return MAE vs. Number of Assets*

forecasted return MAE increases. However, considering that the dimension of the estimation space is increasing, and therefore the potential estimation outcomes are exponentiating, the forecasted return MAE values do not degrade unreasonably. As well, although adding additional assets increases forecasted return MAE, as expected, we do observe that each new asset has a decreasing marginal effect on this accuracy measure.

For comparison, let us consider one of the constrained Kalman Filtering procedures from above, the CKalCov. Figure 1.7 uses this method to estimate weights on the simulated assets, and then constructs the forecasted return MAE values for each simulation. Here, the forecasted return MAE values degrade at a much faster rate than with the DPM. In fact, this relationship is unfortunately much more linear as the number of assets increase. Furthermore, at large numbers of assets, there are many more extreme values for inaccurate forecasted return MAE.

Although they are not exhibited here, similar plots for CLS and ICLS demonstrate even less accurate results. Therefore, the DPM is increasingly favored when considering a sizable potential space of investable assets. In the real world, we observe that hedge funds can invest in a vast number of potential assets, thus the DPM becomes an even more useful tool for estimating weights on that asset set.

1.6.3 Real Assets & Model Hedge Fund Example

To further motivate the suitability of the DPM, we construct a realistic hedge fund trading model similar to the one proposed in Khandani and Lo (2007). Then, the resulting asset weights are used to create a model hedge fund return series. We then use this return series, along with the portfolio asset returns, in the estimation procedures to obtain estimates for the asset weights. Finally, these estimated weights are compared to the true asset weights to infer accuracy. As the portfolio assets, we use the daily returns from the four largest sectors in S&P 500 Index (Technology, Financials, Health Care, and Energy) from January 4, 2010 to February 13, 2013.

Let us construct the model trading strategy as follows. Given a set of N equity sectors, consider a strategy where these are held proportional to their market capitalization $\tilde{w}$, but these weights are decreased or increased proportional to previous over or under-performance relative to their average. That is, the sectors which have previously over-performed are relatively under-weighted, while the sectors which have previously under-performed are relatively over-weighted. This is a form of "contrarian" strategy by under-weighting yesterday's winners and over-weighting yesterday's losers. For our example, we use the aggregated sum of daily returns over the last 30 days for each asset ($R_{t-1,i} \equiv \sum_{\tau=t-30}^{t-1} r_{\tau,i}$) when constructing our over/under-weight values. Specifically, the following asset weights are constructed:

$$w_{t,i} = \tilde{w}_{t,i} - (R_{t-1,i} - R_{t-1,m}), \quad R_{t,m} \equiv \frac{1}{N} \sum_{i=1}^{N} R_{t-1,i}$$

These are taken as our "true" weights used to construct the hedge fund return series. Now, to compare the accuracy of our procedures, we look at the estimation of our hedge fund's weight process. Figure 1.8 exhibits the resulting weight point-estimates and confidence intervals (or respective Bayesian credible intervals) for the S&P 500 Health Care Sector.

Visually, the DPM does the best job of tracking the true weight time series with an MAE of 1.16%, while the CN-DPM comes in a close second with 1.23%. For comparison, the OLS methods perform 43-75% worse, while the Kalman methods perform 12-21% worse, across the different component assets. As well, not only does the DPM accurately estimate the underlying weight process, but also it adjusts to large changes in the weights very quickly.

The rolling CLS and ICLS, as expected, take quite a few periods to adjust to large changes in the weights since the importance of each observation is given equal weight

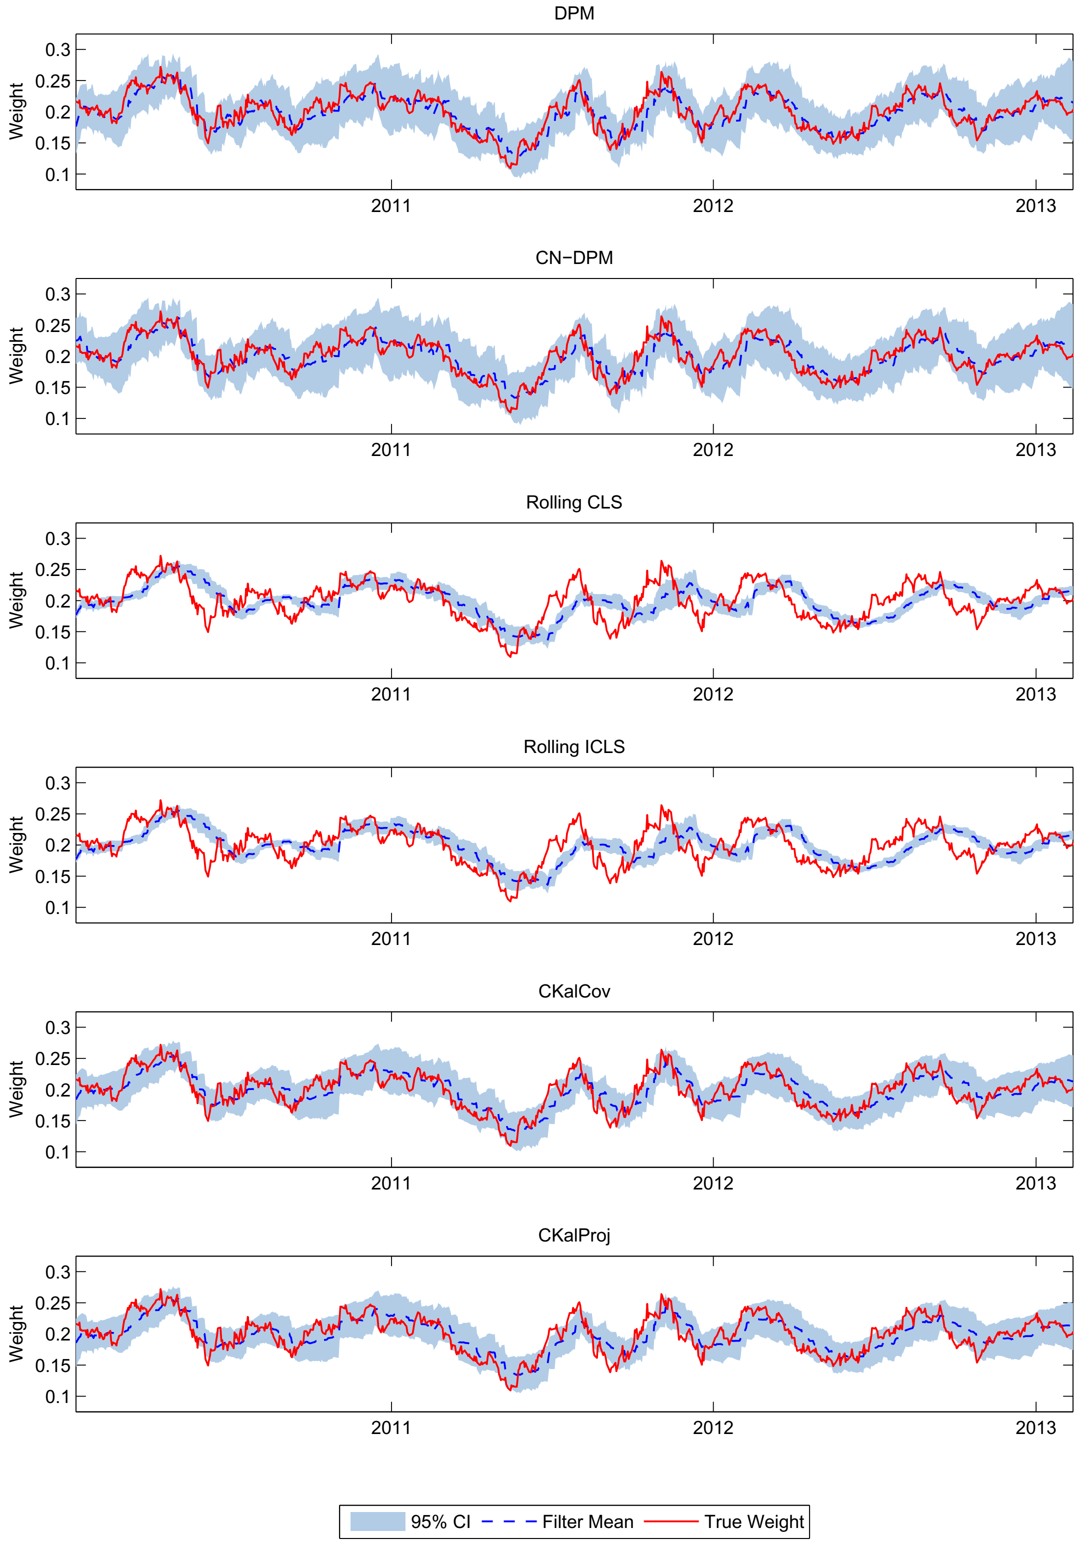

Figure 1.8: *Estimated Weight Accuracy Comparison*

in determining the resulting weight estimates. Therefore, a new portfolio return observation does not have a large impact on the weight estimation, especially if a large window size is used. However, if too small a window size is used, poor estimates and confidence intervals are obtained since the estimation sample is too small. Furthermore, since there is no intertemporal structure placed on weight transitions, we can obtain unreasonably large jumps in the estimated weights due to multicollinearity in the palette asset returns.

The constrained Kalman filtering methods do a better job, however they still do not track the true weight series as accurately as the DPM. This is especially evident when looking at quick changes in the weight values, however it is not as severe as compared to the rolling CLS and ICLS since the structure of the Kalman Filter allows for more appropriate updating of the weight estimation after obtaining a new portfolio return observation. Furthermore, similar to the results of the previous subsection, as the number of component assets increase, the DPM performs increasingly better, relative to the other methods.

1.7 Empirical Results & Comparison

1.7.1 Hedge Fund Data

We apply the DPM estimation methodology to monthly return data for the Hedge Fund Research Fund Weighted Composite Index (HFRIFWC) from January 1995 to October 2012. These return values are reported net of individual fund managers' fees. Since this index is constructed by compiling self reported hedge fund returns from individual managers, there are a few potential biases to identify in the data. First, there is no requirement that fund managers report their monthly returns, therefore only a subset of all funds report into these aggregated indices. Some hedge fund strategies have a maximum asset size which can be effectively invested, thereby creating a cap on the total fund size. In this case, some funds who have reached their maximum size may have no incentive to report their returns. This creates a downward bias in the self reported return indices if these missing funds are outperforming the average. On the other hand, there is a selection survivorship bias in the self reported returns since only the funds that are continuing to operate and therefore have not experienced large losses are reporting returns. This creates an upward bias in these self reported returns. Nevertheless, a large portion of fund managers do report returns, mainly for advertising purposes. Therefore, this aggregated index is the best proxy available for the whole hedge fund industry's returns to investors. So, we can reasonably use this return series in estimating the weights on our set of palette assets in the following sections.

1.7.2 Palette Assets & Parametrization

Since the universe of investable assets is quite large, it is very difficult to estimate weights on the complete set. However, since the goal is to estimate the invested weights on the assets that the value-weighted aggregate of hedge funds is invested in, this problem is simplified significantly. Instead of trying to estimate weights on each and every single investable asset, we can estimate the weights on portfolios of assets (or indices) representing broad classes of assets (e.g. US equity, emerging market equity, high yield bonds, etc.). Since the total hedge fund industry is so large, it is reasonable to make the assumption that the value weighted aggregate of hedge funds is invested in each of these broad asset classes in an approximate weighting scheme that is similar to the asset class's value weights. Therefore, instead of trying to estimate the weights on a potentially infinite set of individual assets, it is possible to estimate these exposures on a small subset of asset classes. Due to the smaller size of the number of asset classes, the dimension of this problem is significantly reduced,

Color	Asset Class Name
	Barclays Municipal Bond Index
	Barclays Short Term Treasury Bond Index
	Barclays Corporate High Yield Bond Index
	Deutsche Bank US Dollar Long Futures Index
	Dow Jones - UBS Commodity Index
	MSCI Emerging Markets Index
	MSCI EAFE Index (Europe, Australasia, Far East)
	MSCI US Equity Index

Table 1.1: *Asset Class "Palette"*

and therefore the resulting estimates are dramatically improved.

Herein, the following indices, similar to those used in Fung and Hsieh (1997), are used as a proxy for the asset classes that the hedge fund industry is investing in. Table 1.1 enumerates the asset class list and color key that will be used throughout the remainder of this paper. Note that it is not necessary to restrict ourselves to this specific set of asset classes. The methodology described above can be applied to any asset set of interest.

As well, the parametrization of the error distributions must be specified. For the transitional portion of the model, the Dirichlet errors are parametrized by the multivariate concentration parameter $\boldsymbol{\alpha}$. Intuitively, since the Dirichlet distribution is the conjugate prior of the multinomial distribution, this $\boldsymbol{\alpha}$ vector can be viewed as pseudo-counts for the prior distribution on the transitioned state of asset weights. In other words, it is the relative weight of the prior when updating with the return observation to obtain the posterior asset weight distribution at a given time period. Recall that $\sum_i \boldsymbol{\alpha}_i = \alpha$, therefore the relative weight on the prior is given solely by α, where the weight on a single observation is 1. This conveniently allows us to effectively quantify the influence that a single new observation has on each step in the estimation procedure as $1/(1+\alpha)$.

Nevertheless, it is convenient to perform Bayesian model comparison via Bayes Factors in this application. As essentially a likelihood ratio between competing model parametrizations, we simply can compute the marginal likelihoods for each α and choose the largest value. Figure 1.9 shows the log marginal likelihood values for the DPM procedure estimated at various values for α. The maximum value is achieved at $\alpha = 1600$. We note that, deviations around $\pm 70\%$ of this choice of α do not change the following results materially.

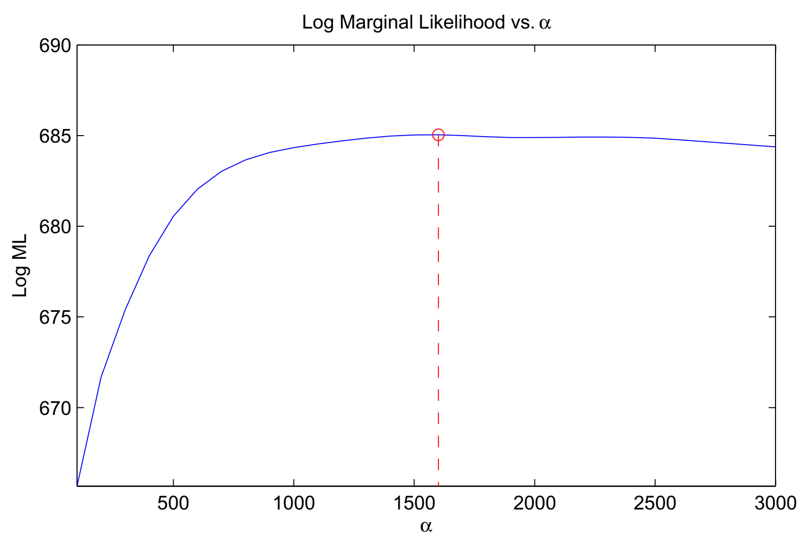

FIGURE 1.9: *Choice of Parameter α*

Lastly, the observational portion of the model needs to be specified. Therefore, the variance σ_ϵ^2 and the degrees of freedom ν of the Student-t distribution need to be chosen. Here, there is more flexibility in the choice of these parameters based upon the selection of palette assets and beliefs about the unexplained portion of returns. The better the palette assets represent the investable universe, the smaller the choice of σ_ϵ^2. As well, the more likely it is to observe extreme values in the unexplained returns, the smaller the value of ν is desired. Here, $\sigma_\epsilon^2 = 0.01$ and $\nu = 6$ are used since the above palette assets represent the investable universe well, but can allow for occasional extreme values in the unexplained component due to the large kurtosis commonly observed in financial return data (Mandelbrot, 1963).

Herein, we proceed with the above parametrization. Otherwise, we note that the Particle Learning work of Carvalho, Johannes, Lopes, and Polson (2010) could be further applied to the DPM estimation algorithm to allow estimation of the error parametrization at the same time as the estimation of the latent weight values.

In the following sections, our methods are used to estimate weights on the set of palette assets for the aggregate hedge fund return index. Then, using these weights, the model is used to forecast the one-step-ahead predicted weights to create a time series of "forecasted" index returns. These returns are plotted along with the observed index returns to compare how precise the estimation is in terms of forecasted return accuracy.

1.7.3 Rolling CLS & ICLS

From the previous sections, the ICLS has strictly produced more accurate results than the CLS without the positivity constraint. Hence for conciseness, only the ICLS results are exhibited. Figure 1.10 shows the forecasted returns and estimated weight plots for the Hedge Fund Research Fund Weighted Composite Index (HFRIFWC) using the rolling ICLS estimation method.

Notably, excessively large weights are placed on high yield bonds and short-term treasury securities. As well, we observe occasional periods with very large jumps in the portfolio weights. This effect is caused by multicollinearity between the palette asset returns, and therefore the static OLS procedure has a difficult time separating the ultimate return contribution of specific assets. Because of this, the resulting forecasted returns do no track the index well, thereby inferring poor asset class weight estimates.

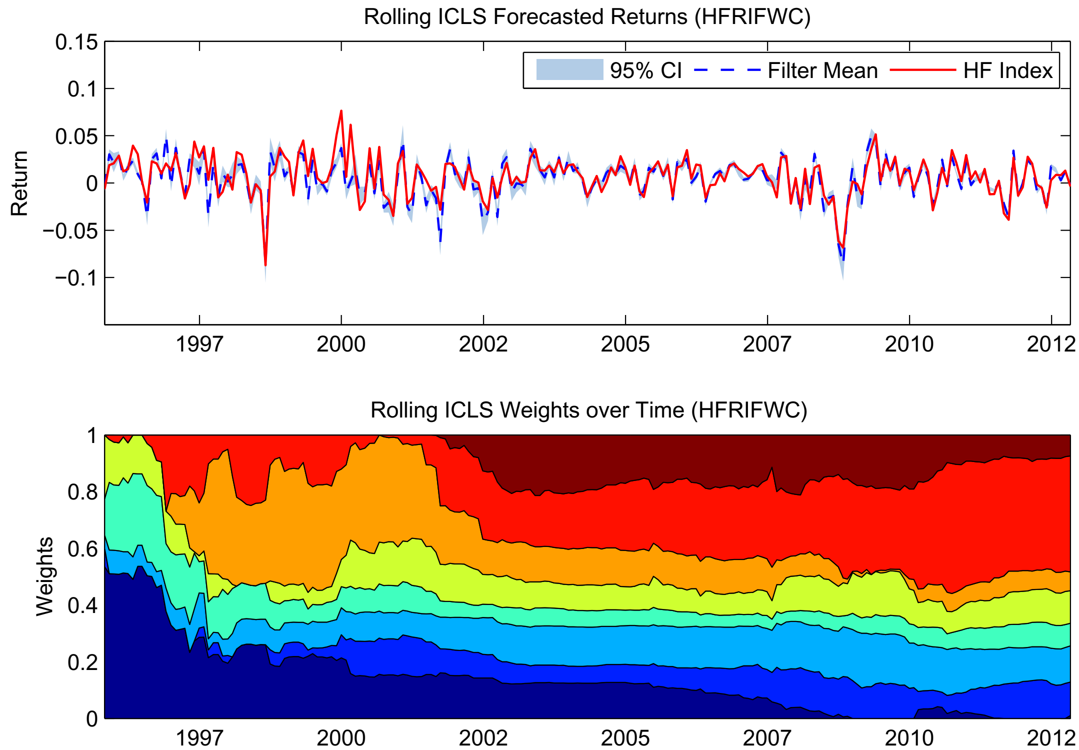

FIGURE 1.10: *Forecasted Returns & Estimated Weights - Rolling ICLS (HFRIFWC)*

Nevertheless, we see a large increase in the weight on short-term treasury bills around the recent economic downturn. High yield bonds have a very large weight until 2002. Equity investments seem to be originally focused in US stocks, with a general transition to emerging markets and Europe, Australasia, and the Far East (EAFE) investments over the sample period. Finally, the plot shows an extra large investment weight in municipal bonds, with a large spike starting around 2001.

This is used as a baseline to see how the three dynamic models compare to this static model's estimation approach.

1.7.4 *Constrained Kalman Filter*

Since the results for the CKalCov and CKalProj methods are very similar, we only exhibit CKalCov here. Figure 1.11 shows the forecasted returns and estimated weight plots for the Hedge Fund Research Fund Weighted Composite Index (HFRIFWC) using the CKalCov estimation method.

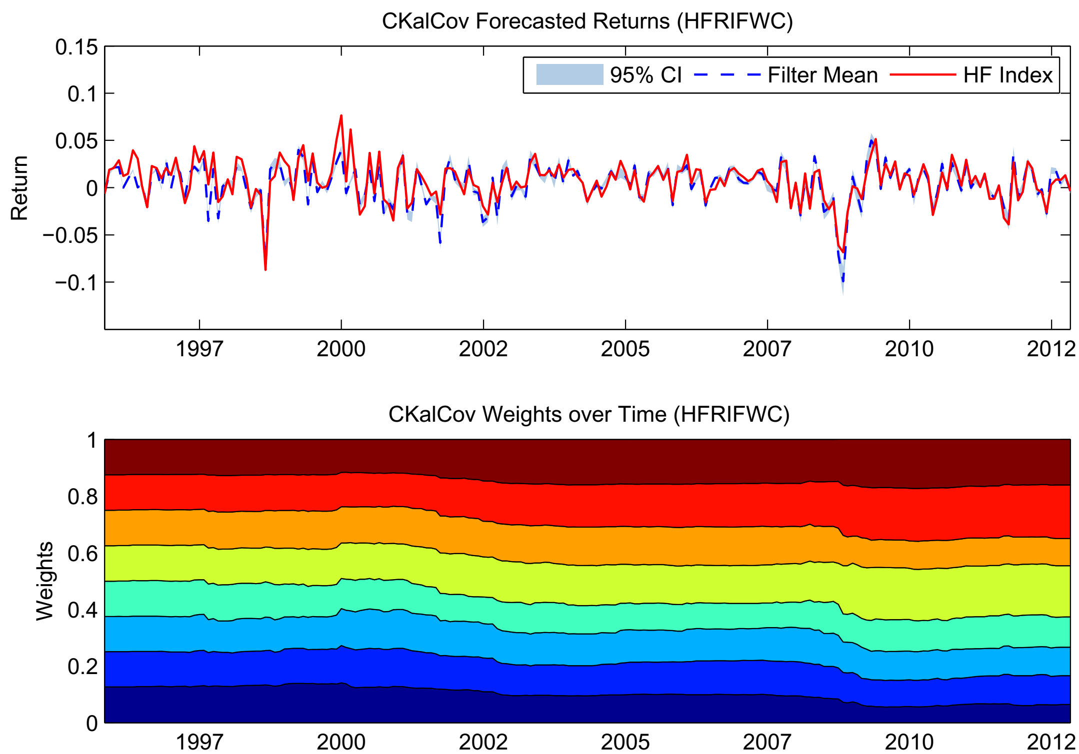

FIGURE 1.11: *Forecasted Returns & Estimated Weights - CKalCov (HFRIFWC)*

The Constrained Kalman Filter immediately looks like a superior model from the forecasted return comparison perspective. There are periods of time where it is inaccurate, but it is much better than the rolling OLS based technique. However, the estimated weights stay alarmingly even over the entire time horizon. This is caused by the spherical and time-invariant weight transition covariance structure combined with the contemporaneous correlation observed across financial asset returns. There are some values that increase or decrease, and not surprisingly, these are generally consistent with the direction of change in the weights from the rolling regressions. That is, there is larger weight placed on municipal bonds and short term treasuries, while a slowly decreasing weight is placed on US equities.

The most noteworthy attribute of the estimation is the small, but clear jumps in the weights around the dates of economic downturns, thereby indicating a shift in asset allocation occurring in the hedge fund industry in those periods.

1.7.5 Dirichlet Portfolio Model

Figure 1.12 shows the forecasted returns and estimated weight plots for the Hedge Fund Research Fund Weighted Composite Index (HFRIFWC) using the DPM estimation method.

Looking at the forecasted returns plot, the DPM does an excellent job of tracking the forecasted index returns. In general, we see fewer periods of poor estimation, as we did with the rolling regressions and the Kalman Filter methods.

As well, the weight estimation results are much more dynamic. Not only does the plot show clear jumps in the weights around economic downturns and also subsequent shifts to short-term treasuries, but also increases in municipal bond investments over the following years. Again, there is a noticeable shift from US stocks to increasing investments in emerging markets and EAFE seen over the recent years. Furthermore, a sizable investment in high yield corporate bonds is observed until 2000, when this weight shifted to investment grade corporate credit from 2000-2006, after the tech bubble.Finally, overall equity and fixed income exposures vary with economic cycles, which is consistent with beliefs about widespread portfolio allocation dynamics.

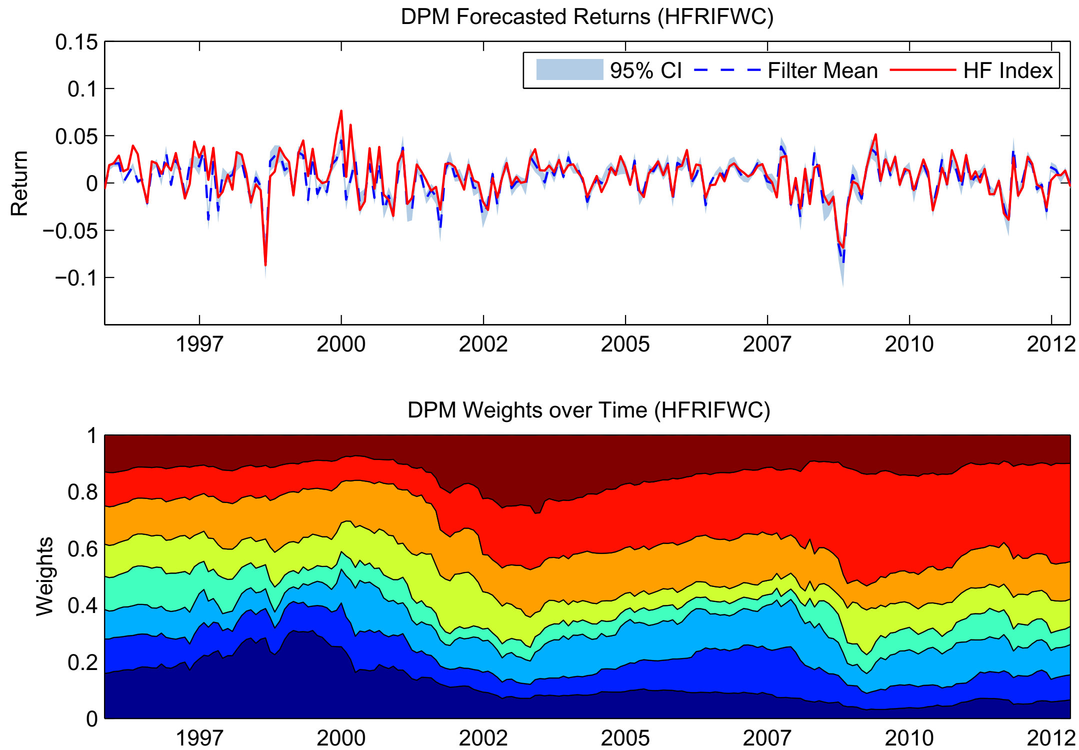

FIGURE 1.12: *Forecasted Returns & Estimated Weights - DPM (HFRIFWC)*

1.7.6 Conditionally Normal Dirichlet Portfolio Model

Figure 1.13 shows the forecasted returns and estimated weight plots for the Hedge Fund Research Fund Weighted Composite Index (HFRIFWC) using the CN-DPM estimation method.

The conditionally normal approximation of the DPM also provides accurate forecasted tracking of the hedge fund index. Generally the same changes in investment patterns are observed. However, weights remain more evenly distributed across all the asset classes than observed in the original DPM estimates. This effect is the result of the beta distribution's mode being closer to the bounds of its support. In other words, its skewed mode pulls values closer to either very large or very small val-

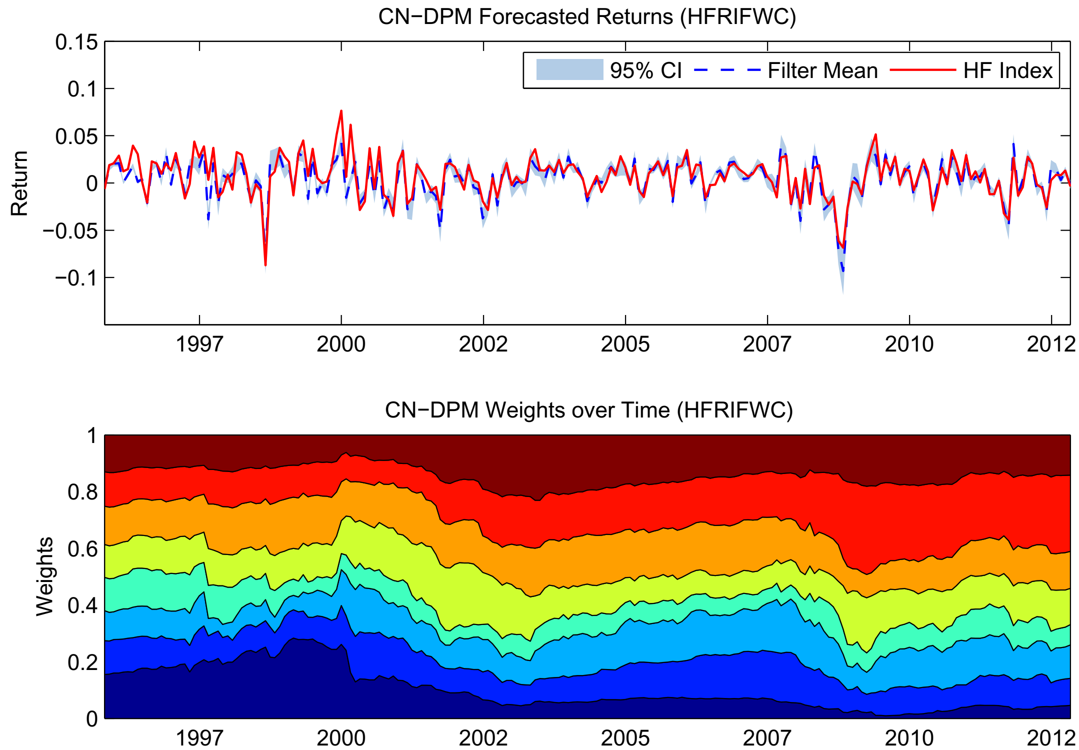

FIGURE 1.13: *Forecasted Returns & Estimated Weights - CN-DPM (HFRIFWC)*

ues, whereas a normal approximation distributes probability over the support with no skew. Lastly, note that this effect does not necessarily imply less or more accurate weight estimation since the forecasted return MAE value is very close to that found with the DPM estimation.

1.7.7 Comparison & Results

In the previous section, the DPM and CN-DPM demonstrated a superior job in constructing forecasted returns which were consistent with the hedge fund return index observed returns. Since the true invested hedge fund weights cannot be observed, the accuracy of the weight estimation cannot be directly assessed. Nevertheless, these forecasted returns give the best proxy for measuring accuracy using the information in this setup.

The DPM and CN-DPM produce asset weight shifts which are generally consistent the broad beliefs about how investment managers have shifted around their asset allocation weights over the last 15 year period. Examining the estimated changes in weights, perhaps the most interesting pattern is the systematic increase in investment in municipal bonds during recessions, and then a subsequent decrease during economic recovery. This is consistent with the perceived view that municipal bonds are a relatively safe investment, and therefore during times of high economic uncertainty they provide a reasonably safe investment vehicle. This is supported by Appleson, Parsons, and Haughwout (2012) who find that municipal defaults are less likely connected to economic downturns than defaults on corporate bonds. Therefore, observing hedge fund managers increasing investment flows to municipal bonds during these periods is easily rationalized. Nevertheless, the low risk nature of municipal bonds has come under much debate, especially since late 2007 to early 2008 when many municipal bond prices declined without seeing relative increases on similar duration swap contracts used to hedge interest rate risk (Deng and McCann, 2012).

Furthermore, a comparison of the sample statistics for the estimation methods is shown in Table 1.2. This table provides a comparison of the MAE, RMSE, R^2, correlation, mean, and standard deviation of the forecasted returns of the six estimation methods. Note that non-forecasted returns are constructed by applying the estimated weights to the same period returns, while the forecasted returns are the forecasted weights applied to the one-period-ahead returns. Here, the DPM and CN-DPM consistently outperform the other methods in terms of tracking error, mean absolute error, correlation, and R^2. As well, they generally do the best in replicating the respective mean and standard deviation values.

In all of the above estimation techniques, we point out that large weight is estimated on short-term treasury securities. As identified in Getmansky, Lo, and Makarov (2004), hedge fund managers commonly employ smoothing in their monthly self reported returns. Since these treasury assets are perceived to be risk free and have a very low return volatility as compared to the remaining assets, artificially low volatility in the hedge fund index can lead to a larger weight estimated on assets with low volatility themselves, like the short-term treasuries. From the hedge fund managers' perspective, this smoothing has the effect of improving their funds' observed risk-adjusted performance. Second, this self reported return smoothing can arise from the pricing of illiquid assets (Fisher et al., 2003; Kadlec and Patterson, 1999). Therefore, when using self reported hedge fund return data, it is common to estimate a desmoothing model on the return data. When implementing the model from Getmansky, Lo, and Makarov (2004) on the hedge fund return index, the resulting portfolio weight estimate on the short-term treasuries decreases significantly, while the weights on the remaining assets scale up, proportionately to each other.

	HFRIFWC						
Non-Forecasted	**Index**	**DPM**	**CN-DPM**	**CLS**	**ICLS**	**CKalCov**	**CKalProj**
Mean	0.00748	0.00457	0.00460	0.00516	0.00478	0.00430	0.00431
Standard Deviation	0.02100	0.02014	0.02003	0.02128	0.02100	0.02068	0.02078
RMSE	0	0.00921	0.00917	0.00970	0.00980	0.01021	0.01027
Mean Abs Error	0	0.00626	0.00644	0.00699	0.00670	0.00752	0.00756
Correlation	1	0.90893	0.90949	0.89930	0.89765	0.88967	0.88875
R^2	1	0.81062	0.81240	0.78994	0.78569	0.76718	0.76474

	HFRIFWC						
Forecasted	**Index**	**DPM**	**CN-DPM**	**CLS**	**ICLS**	**CKalCov**	**CKalProj**
Mean	0.00748	0.00414	0.00408	0.00477	0.00447	0.00412	0.00415
Standard Deviation	0.02100	0.02079	0.02084	0.02265	0.02181	0.02116	0.02125
RMSE	0	0.01049	0.01074	0.01181	0.01145	0.01076	0.01079
Mean Abs Error	0	0.00737	0.00760	0.00844	0.00806	0.00790	0.00791
Correlation	1	0.88439	0.87900	0.86160	0.86529	0.87986	0.87969
R^2	1	0.75509	0.74326	0.68966	0.70851	0.74224	0.74125

Table 1.2: *Replication Summary Statistics (Monthly)*

1.7.8 *Negative Portfolio Weights*

Throughout this paper, we have assumed and motivated the restriction of non-negativity on the latent asset class weights. Nevertheless, it can be mentioned that an estimation method which allows for negative weights can easily be constructed in the style of the DPM setup. The idea is to construct two portfolios, one for long (positive) positions, and another for short (negative) positions. This allows the long portfolio to capture the variation in the hedge fund index explainable by the positive asset returns, as well as the short portfolio to capture the variation explainable by the negative asset returns. Then, with estimated distributions for these sets of positive and negative weights, we can estimate a time varying combination factor used to obtain an overall portfolio weighting, thereby potentially increasing the overall explanatory power of the portfolio.

One way to do this is to estimate these separate long/short portfolios in each time period, with respective weights w_t^+ and w_t^-. A combined portfolio then can be constructed via $w_t = (1 + \gamma_t) w_t^+ - \gamma_t w_t^-$ where the time varying combination factor γ_t follows a Gaussian random walk model $\gamma_t \sim N(\gamma_{t-1}, \sigma_\gamma^2)$. This combination factor can also be estimated via a similar sequential Monte Carlo procedure. In our estimation problem, not surprisingly, this combination factor was generally found to be $\gamma_t \approx 1$, implying that the aggregate hedge fund industry portfolio does not have negative exposures to these asset classes.

1.8 Applications

1.8.1 *Replication of Investment Strategy*

The idea of replicating hedge fund investment strategies through low cost, liquid investments is not a new idea. Many large securities firms currently have products which seek to do exactly this. Many of these take a bottom-up approach which attempts to identify the types of trades and systematic patterns that funds employ to create their asset allocation, then implement these ideas in an algorithmic manner. We instead take the top-down approach which is much more statistically sound, as it is attempting to identify component exposures to candidate sets of asset classes, in order to best track the time series of returns.

In order to create a replicating portfolio in this manner, one simply needs a set of relative weights w^{Rep} on the asset set of interest. Ideally, these weights would be the same as the true hedge fund invested weights, but these are latent. Therefore, the DPM's expectation of the weights given all observable information up to that time can be used:

$$
\begin{aligned}
w_t^{Rep} &\equiv E[w_t|\mathcal{F}_{t-1}] \\
&= \frac{w_{t-1}^{Rep} \circ (1 + r_{PA,t-1})}{\sum_{i=1}^{n} w_{t-1,i}^{Rep}(1 + r_{PA,t-1,i})} \\
&= \frac{E[w_{t-1}|\mathcal{F}_{t-2}] \circ (1 + r_{PA,t-1})}{\sum_{i=1}^{n} E[w_{t-1,i}|\mathcal{F}_{t-2}](1 + r_{PA,t-1,i})} \\
&= \ldots
\end{aligned}
$$

Using these weights, one can invest in the assets of interest and construct the appropriate replicating portfolios. Naturally, the goal is to construct portfolios which have very similar returns to investors as the hedge fund indices being replicated. Since these indices are non-investable, access to these returns is usually obtained through investing in a fund-of-funds, which imposes their own aforementioned layer of fees. Therefore, the raw index returns are adjusted for these fund-of-fund fees for real comparison purposes. Nevertheless, we note that there exists upward bias in these index returns that remains unadjusted for. Figure 1.14 shows a plot of the cumulative return to investors for the adjusted HFRIFWC Index with an initial investment of $1.

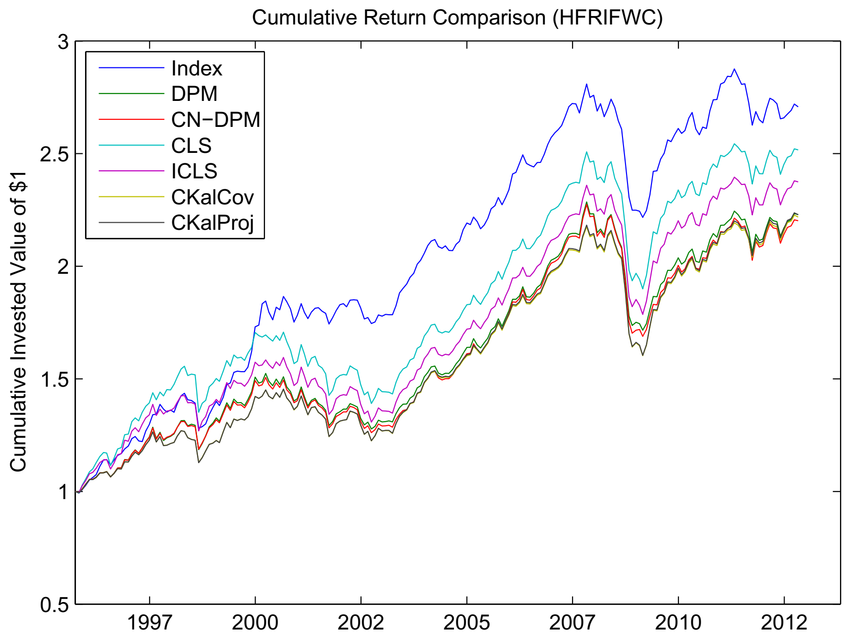

FIGURE 1.14: *Cumulative Return Comparison (HFRIFWC)*

The DPM and CN-DPM do a very good job in replicating the return series. Notably, the rolling CLS does the worst, due to its lack of positivity restriction on the weights, resulting in in-sample over-fitting. The CKalCov and CKalProj methods do a decent job, however their poor estimation of weights contributes to tracking inaccuracy and under-performance before 2000.

Finally, it is important to note that we do not take a stand on whether implementation of this replication is a good investment strategy. The answer to this lies in whether the hedge fund industry delivers superior risk adjusted returns. The answer to that is outside the scope of this paper.

1.8.2 Intraperiod Return & Volatility Approximation

Since funds only disclose their investment returns at discrete intervals, typically monthly or quarterly, it is difficult for investors to know how their invested capital is performing in the time between. They could have gained or lost a lot of wealth over the course of a few days, but they will not realize this information for potentially months later. This presents an informational problem, since the knowledge of this investment performance has clear implications for consumption decisions in the current period.

As well, it is very common to see portfolio managers for large pension funds, endowments, family offices, etc. not only invest in individual assets, but also other investment managers. Therefore, there is value in knowing how their less liquid and transparent hedge fund investments are performing in order to more appropriately manage risk in the rest of their overall portfolio. Hence, having an approximation of these intraperiod hedge fund return and volatility values is of great value.

With the DPM setup, we have an effective way of approximating these intraperiod returns. Consider the following setup where we want to approximate the intraperiod returns for τ units of time past reporting period time t:

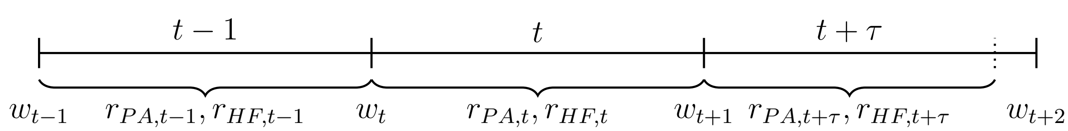

FIGURE 1.15: *Intraperiod Timeline Illustration*

We are interested in determining the value of the holding period return and volatility over time period $t+\tau$ given the observed information at time $t+\tau$. That is, we want to estimate $E[r_{HF,t+\tau}|\mathcal{F}_{t+\tau}]$ and $Vol[r_{HF,t+\tau}|\mathcal{F}_{t+\tau}]$. However, note that

$\mathcal{F}_{t+\tau} = \{\mathcal{F}_t, r_{PA,t+\tau}\}$ since the only new information observed since time t is the return on the palette assets. Since the true weights cannot be observed to use in this calculation, we can use our estimated weights from time t. Therefore, the estimate of intraperiod return, using the DPM model's estimates, simply becomes:

$$\begin{aligned} E\left[r_{HF,t+\tau}|\mathcal{F}_{t+\tau}\right] &= E\left[r_{HF,t+\tau}|\mathcal{F}_t, r_{PA,t+\tau}\right] \\ &= E\left[w_t|\mathcal{F}_t\right]' r_{PA,t+\tau} \end{aligned}$$

Similarly, an estimate of the intraperiod volatility can be computed:

$$\begin{aligned} Vol\left[r_{HF,t+\tau}|\mathcal{F}_{t+\tau}\right] &= Vol\left[r_{HF,t+\tau}|\mathcal{F}_t, r_{PA,t+\tau}\right] \\ &= \sqrt{E\left[w_t|\mathcal{F}_t\right]' \hat{\Sigma}_{r_{PA,t+\tau}} E\left[w_t|\mathcal{F}_t\right] + \sigma_\epsilon^2} \end{aligned}$$

where $\hat{\Sigma}_{r_{PA,t+\tau}}$ is an intraperiod covariance matrix for the palette assets. This matrix for the latent covariance structure can be constructed in various ways, including but not limited to Stochastic Volatility (Jacquier, Polson, and Rossi, 1994, 2004) and DCC-GARCH (Engle, 2002) approaches. Furthermore, the same expression above can be used to compute *forecasted* hedge fund volatility if the current period is taken to be time t and we want to forecast τ time into the future.

1.9 Conclusion & Discussion

This paper has presented a Bayesian dynamic model, the Dirichlet Portfolio Model (DPM), for the hedge fund industry weight transition process and aggregate return observations. We then exhibited a numerical solution to this model using sequential Monte Carlo methods, as well as a conditionally normal approximation (CN-DPM) which was solved analytically. In order to motivate the appropriateness of this dynamic model, other models and their respective solutions were compared. The simulated and model hedge fund results showed that both with simulated or real assets, as well as under simulated or model trading, the DPM produces more accurate estimates of the underlying weights as compared to the results produced by the other estimation methods. Overall, the DPM provided superior results across various measures of suitability. Interestingly, the estimation results on the the hedge fund industry aggregate return index identify a systematic increase in exposure to municipal bonds during economic downturns, and a subsequent decrease in exposure during economic recovery periods, which is consistent with the notion that defaults on municipal bonds are less connected to economic downturns than defaults on corporate bonds.

From the foundational DPM, there are many future extensions from this starting point. One of the challenges of the DPM estimation procedure is having to pre-specify the distributional error parameters, α, σ_ϵ^2, and ν. Just as there is value in obtaining the latent weight estimates, it would also be insightful to learn the magnitude of these tuning parameters during the estimation process. This can be achieved by applying the parameter learning concepts from Carvalho, Johannes, Lopes, and Polson (2010), Storvik (2002), Fearnhead (2002), or Liu and West (2001).

Furthermore, we identify applications of this methodology to both creation of hedge fund industry replicating portfolios, as well as intra-reporting-period return and volatility estimation. There is large value in being able to approximate these intra-reporting-period returns for both current consumption choices and various risk management decisions. As well, being able to create replicating portfolios from the asset class decomposition has the potential to construct more transparent portfolios with much lower cost structures. Therefore, the Dirichlet Portfolio Model is a convenient technique for decomposing unobservable portfolio compositions, allowing for future analysis on the dynamics of these weight processes.

2

Hedge Fund Portfolio Decomposition: Joint Estimation of Net Leverage and Asset Allocation Dynamics

2.1 Abstract

This paper extends a class of dynamic compositional state space models from Korsos (2013b) for estimating latent portfolio compositional weights when only observing the aggregate portfolio and underlying asset returns. Here, we augment the previous setup by demonstrating how to implement the joint estimation of net portfolio leverage dynamics. As well, by incorporating recent work for parameter learning in state space models, we detail how to not only sequentially estimate the time varying latent portfolio weight and leverage values, but also the distributional tuning parameters. We illustrate this technique on the estimation of asset class weights and leverage on a set of aggregate hedge fund indices from 1995 to 2012. Finally, using the resulting time series of hedge fund portfolio compositional values, we estimate the same-period price impact of portfolio allocation changes.

2.2 Introduction

In the field of portfolio risk attribution and analysis, extensive work has been done on discovering cross-sectional risk factors and estimating exposures to these various sources of risk. For years, this linear risk factor based approach has provided a convenient way to decompose investment portfolios into common components to explain variation in portfolio returns and then assess various economic questions such as pricing implications, managerial skill, as well as many others. Comparatively, little work has been done on compositional time series for decomposing and modeling portfolio weights directly on the simplex. Therefore, rather than estimating effects of risk factors on portfolio returns, we estimate relative portfolio weights on a set of investable assets.

This idea of decomposing portfolio weights onto investable assets originated with Sharpe (1992) where he estimates asset class weightings on individual mutual funds via constrained OLS estimation. In this seminal work, he estimates both time-invariant weights on the entire sample period, as well as uses a rolling window approach to form time-varying weight values. Fung and Hsieh (1997) extend this idea in the application of hedge fund portfolios. They identify that static asset class weights have much lower in-sample explanatory power than those for mutual funds, thereby indicating that dynamic asset allocation is a more significant component of hedge fund portfolios. With this in mind, Mamaysky, Spiegel, and Zhang (2007) introduce a Kalman Filter based approach which more appropriately identifies risk factor loadings as a dynamic process. Our approach acknowledges the benefits of this dynamic attribution process while directly modeling relative portfolio weights on the simplex, thereby developing a more suitably specified generative model.

Herein, we extend a class of compositional state space models based upon Dirichlet distribution latent value transitions, initially outlined by Korsos (2013b) for observing real valued portfolio return data and then modeling a latent set of time varying portfolio compositional weights on the simplex.

$$w_{t+1} \sim Dir\left(\alpha \frac{w_t \circ (1 + r_{A,t})}{\sum_{i=1}^{N} w_{t,i}(1 + r_{A,t,i})}\right)$$

$$r_{\Phi,t} \sim t\left(w_t' r_{A,t}, \sigma_\epsilon^2, \nu\right)$$

This approach allows a natural way to capture the variation in portfolio returns attributable to time-varying asset allocation, in addition to the variation in underlying asset returns. Since these models do not conform to the classic Gaussian-Linear form for Kalman filtering, solutions are based on the sequential Monte Carlo approach of Gordon, Salmond, and Smith (1993), Rubin (1987), and Smith and Gelfand (1992). Due to the sequential nature of this estimation procedure, this allows us to not only explain the cross-sectional variation in hedge fund returns, but also create out-of-sample forecasts at each period of compositional asset holdings in order to evaluate the forecasted tracking ability of the dynamic weight process.

Since many fund managers commonly employ the use of leverage to obtain their distribution of returns, we also wish to jointly identify this effect on portfolio returns. Similar to the relative portfolio weights, the dynamics of this net leverage amount is also unobservable. Therefore, in this paper we extend the idea of the Dirichlet Portfolio Model (DPM) by exhibiting how to incorporate net leverage into the estimation procedure to model both the normalized portfolio weights, as well as a time-varying net leverage scaling parameter. This is done via the addition of a capital appreciation adjusted net leverage transition process and a respective observational model modification:

$$\gamma_t \sim N\left(\frac{\gamma_{t-1}}{1 + r_{\Phi,t-1}(1 + \gamma_{t-1})}, \sigma_\gamma^2\right)$$

$$r_{\Phi,t} \sim t\left(\left(1+\gamma_t\right) w_t' r_{A,t}, \sigma_\epsilon^2, \nu\right)$$

As well, we show how to implement the parameter learning ideas of Lopes, Carvalho, Johannes, and Polson (2010) to perform online learning of the distributional tuning parameters of the various stochastic portions of the models.

In order to motivate the convenience and effectiveness of this class of compositional state space models, we estimate asset class level portfolio weights and leverage multiplier values on an index of aggregate hedge fund industry returns from 1995 to 2012 using a similar asset class set to that in Fung and Hsieh (1997). Then, to gain insight into the behavior of different types of hedge fund strategies, we examine how these time-varying leverage levels differ across various classifications of hedge funds.

Finally, with estimates for the time-varying compositions of the various asset classes, we analyze the dynamics of the weight processes and find that changes in asset class investment due to active trading are consistent with the expected market microstructure impact on prices for those assets (O'Hara, 1995; Madhavan, 2000; Hasbrouck, 2007). That is, as hedge funds actively increase holdings in a given asset class, a positive same-period return effect is observed for that asset, thereby bolstering the accuracy of the estimation results.

The remainder of this paper is structured as follows: Section 2 reviews the class of dynamic portfolio models used for estimation of time-varying portfolio weights. Section 3 extends these results to include net portfolio leverage estimation. Section 4 details the parameter learning procedures for online learning of distributional tuning parameters. Section 5 outlines the hedge fund and asset class data used for estimation. Section 6 presents the asset class portfolio weight estimation results for the hedge fund data. Section 7 compares net portfolio leverage estimation results across various hedge fund strategies. Section 8 estimates asset class investment effects implied by changes in these estimated hedge fund weights. Finally, section 9 concludes.

2.3 Portfolio Weight Estimation

We consider the portfolio weight estimation problem where both the portfolio returns and compositional asset returns are observable, however the relative investment weights on those assets are never completely observable to those outside the firm. As in Korsos (2013b), let there be an known investable set of N assets or asset classes with time-varying latent compositional weights $w_{t,i}$ for each asset i at time t which are required satisfy the budgetary restriction $\sum_{i=1}^{N} w_{t,i} = 1$. First, we briefly review the initial problem of estimating the latent weights on individual portfolio compositional assets in the absence of portfolio leverage. We will use this as a foundation for our net leverage estimation extension in the following section.

Let $\mathcal{F}_t$ represent the filtering of all information known at time t. Hence, this includes all previous portfolio returns r_Φ, compositional asset returns r_A, and compositional asset weights w up to and including time t. That is,

$$\mathcal{F}_t = \{r_{\Phi,1}, \ldots, r_{\Phi,t}, r_{A,1}, \ldots, r_{A,t}, w_1, \ldots, w_t\}.$$

Define $w_t = (w_{t,1}, w_{t,2}, ..., w_{t,N})'$ to be an $N \times 1$ vector of the weights on each asset at the *beginning* of time period t, $r_{A,t} = (r_{A,t,1}, r_{A,t,2}, .., r_{A,t,N})'$ to be an $N \times 1$ vector of the compositional asset returns over time period t, and $r_{\Phi,t}$ to be a scalar value of the return on the portfolio over the same time period t. The chronology of the time period notation is illustrated below:

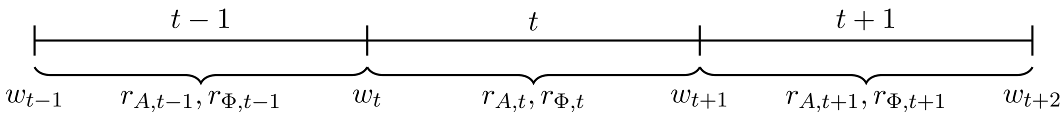

FIGURE 2.1: *Timeline Notation Illustration*

As proposed in Korsos (2013b), let us define a compositional weight transition model directly on the simplex $\mathbb{S}^{N-1} = \left\{ w \in \mathbb{R}_+^N : w'\mathbf{1} = 1 \right\}$ via a Dirichlet random walk model adjusted for capital appreciation of the compositional assets:

$$w_{t+1} \sim Dir\left(\alpha \frac{w_t \circ (1 + r_{A,t})}{\sum_{i=1}^N w_{t,i}(1 + r_{A,t,i})}\right) \tag{2.1}$$

The parameter α is a scalar controlling the dispersion width of the weight transition process. Therefore, the one-step-ahead expectation $E\left[w_{t+1}|\mathcal{F}_t\right]$ of this model is the unadjusted portfolio holdings given initial portfolio weights w_t and realized holding period asset returns $r_{A,t}$:

$$E\left[w_{t+1}|\mathcal{F}_t\right] = \frac{w_t \circ (1 + r_{A,t})}{\sum_{i=1}^N w_{t,i}(1 + r_{A,t,i})}$$

We note that although many asset managers make the decision to rebalance their portfolio holdings each period such that $E\left[w_{t+1}|\mathcal{F}_t\right] = w_t$, this is part of the investment process choice, and therefore this is not assumed for the broad range of portfolios, a priori. Nevertheless, if a manager does make the choice to rebalance assets based upon this notion, this effect will be picked up in the ultimate portfolio weight estimation results. As well, since different managers make different choices regarding the length of their rebalancing intervals, this is further motivation to not make a prior rebalancing assumption on the portfolio weight dynamics.

Next, a leptokurtic observational model is defined on the linear combination of the portfolio weights and the respective holding period returns on those assets:

$$r_{\Phi,t} \sim t\left(w_t' r_{A,t}, \sigma_\epsilon^2, \nu\right) \tag{2.2}$$

Since it may be impossible to fully capture the variation in portfolio returns by the set of assets of interest, observational error is controlled via the variance-scaled version of the Student-t distribution. It is reasonable that the set of chosen explanatory assets may not span the investment space of the portfolio of interest. Therefore, potential for error is introduced into the observed weighted portfolio return because of missing explanatory assets. Due to the presence of excess kurtosis in financial return data, as first documented in Mandelbrot (1963) and many others since then, the Student-t model allows for flexibility outside of the standard Gaussian model.

Combining the weight transition model in (2.1) with the portfolio observation model in (2.2) gives the foundational Dirichlet Portfolio Model (DPM) from Korsos (2013b):

$$w_{t+1} \sim Dir\left(\alpha \frac{w_t \circ (1 + r_{A,t})}{\sum_{i=1}^N w_{t,i}(1 + r_{A,t,i})}\right)$$

$$r_{\Phi,t} \sim t\left(w_t' r_{A,t}, \sigma_\epsilon^2, \nu\right)$$

Since neither the transition model nor the observation model is Gaussian, the latent portfolio weights cannot be solved for in closed form via Kalman (1960). Instead, the sequential Monte Carlo approach of Gordon, Salmond, and Smith (1993) is used to numerically solve for the latent portfolio compositional weights. This estimation algorithm, adapted to the DPM, is exhibited in Figure 2.2. Note that with the lack of better "time period 0" prior portfolio weight information, we assume $E[w_{0,i}] = 1/N, \ \forall i \in \{1, \ldots, N\}$, however if there is superior prior information available, appropriate incorporation into this starting distribution will yield more appropriate estimation results.

DPM Estimation Algorithm

Initialize: Sample prior weights from

$$w_0^{(p)} \sim Dir\left(\alpha_0\left(\frac{1}{N},...,\frac{1}{N}\right)'\right)$$

Iterate:

Step 1: Propagate new asset weights from

$$w_t^{(p)} \sim Dir\left(\alpha\frac{w_{t-1}^{(p)} \circ (1+r_{A,t-1})}{\sum_{i=1}^{N} w_{t-1,i}^{(p)}(1+r_{A,t-1,i})}\right) \text{ for } p = 1,...,P$$

Step 2: Resample asset weights from

$$w_t^{(p)} \sim Mult_P\left(\left\{\omega_t^{(\phi)}, w_t^{(\phi)}\right\}_{\phi=1}^{P}\right)$$

where the importance weights are given by

$$\omega_t^{(p)} \propto \left(1+\frac{1}{\nu}\frac{\left(r_{\Phi,t}-w_t^{(p)\prime}r_{A,t}\right)^2}{\sigma_\epsilon^2}\right)^{-\frac{\nu+1}{2}}$$

FIGURE 2.2: *DPM Estimation Algorithm*

2.4 Net Portfolio Leverage Estimation

There is growing evidence that many investment managers employ various amounts of leverage to obtain their portfolio returns. Similar to the compositional portfolio weights, this dynamic amount of net portfolio leverage is also unobserved. Herein, we extend the structure of the DPM from Korsos (2013b) to include the joint estimation of a time-varying net leverage value.

It is important to recognize the distinction between net leverage and gross leverage values. Consider a fund manager who believes that within a particular industry, a certain stock is over-priced relative to another stock. Wanting to place a bet that these two prices will converge, the fund does not want to be exposed to the common industry risk present in both of these assets. By placing a long bet on the relatively under-priced asset, and a corresponding short bet on the relatively over-priced asset they can achieve this goal. Their total net exposure of this position may be quite small due to the netting of the industry exposure, while the gross exposure can appear very large since it is a total of both long and short positions. This gross exposure, divided by contributed capital is the commonly quoted leverage multiplier value seen in industry reports and the popular media. Instead, we focus on leverage arising from non-netting exposures since this is the leverage which is ultimately meaningful in the distribution of portfolio returns.

Let us define a time-varying net leverage value γ_t which specifies a multiplier on the magnitude of portfolio's net holdings over-and-above the portfolio's contributed capital. That is, if $\gamma = 0.2$ and a manager has \$100 million in contributed capital, the fund has net borrowings of \$20 million in order to hold a total net value of \$120 million of financial assets. This implies the following modified observational model:

$$r_{\Phi,t} \sim t\left(\left(1+\gamma_t\right) w_t' r_{A,t}, \sigma_\epsilon^2, \nu\right) \tag{2.3}$$

If $\gamma = 0$ for all time periods, this model is equivalent to the one in the original DPM.

Next, a Gaussian random walk model, adjusted for capital appreciation, is specified for the dynamics of this leverage parameter:

$$\gamma_t \sim N\left(\frac{\gamma_{t-1}}{1+r_{\Phi,t-1}\left(1+\gamma_{t-1}\right)}, \sigma_\gamma^2\right) \tag{2.4}$$

Notice that as the a portfolio increases in value due to capital appreciation of the component assets, the effective leverage of the portfolio decreases. Just as we want to capture *active* trading decisions on the relative portfolio weights, we take a similar approach to the leverage values. That is, our prior at each period is that the manager has done nothing to change their portfolio since the previous period, thereby allowing the Bayes rule updates to reflect only the active trading decisions. Combining (2.3) and (2.4) with the original DPM gives the Leveraged Dirichlet Portfolio Model (L-DPM):

$$w_{t+1} \sim Dir\left(\alpha \frac{w_t \circ (1 + r_{A,t})}{\sum_{i=1}^{N} w_{t,i}(1 + r_{A,t,i})}\right)$$

$$\gamma_t \sim N\left(\frac{\gamma_{t-1}}{1 + r_{\Phi,t-1}\left(1 + \gamma_{t-1}\right)}, \sigma_\gamma^2\right)$$

$$r_{\Phi,t} \sim t\left(\left(1 + \gamma_t\right) w_t' r_{A,t}, \sigma_\epsilon^2, \nu\right)$$

Again, the sequential Monte Carlo estimation approach is applied to develop the estimation algorithm exhibited in Figure 2.3. We place a Beta distribution prior on the "time period 0" leverage values in order to initialize leverage multipliers in a historically reasonable range.

2.5 Parameter Learning

The main difficulty of these dynamic models is having to specify the values of the tuning parameters before filtering to estimate the latent states. One approach is to choose these tuning parameters via Bayes Factor optimization. This is done by selecting the parameter set $\mathcal{M}_i$ with the largest associated marginal predictive likelihood over the entire data sample:

$$p\left(r_{\Phi,t} | \mathcal{M}_i\right) = \prod_{\tau=1}^{t} p\left(r_{\Phi,\tau} | \mathcal{F}_{\tau-1}, \mathcal{M}_i\right)$$

Instead, we take the approach of simultaneously learning values for these parameters during the sequential latent weight estimation process. To date, a few approaches

L-DPM Estimation Algorithm

Initialize: Sample prior weights and leverage values from

$$w_0^{(p)} \sim Dir\left(\alpha_0\left(\frac{1}{N}, ..., \frac{1}{N}\right)'\right)$$

$$\gamma_0^{(p)} \sim Beta\,(a, b)$$

Iterate:

Step 1: Propagate new asset weights and leverage values from

$$w_t^{(p)} \sim Dir\left(\alpha \frac{w_{t-1}^{(p)} \circ (1 + r_{A,t-1})}{\sum_{i=1}^{N} w_{t-1,i}^{(p)}(1 + r_{A,t-1,i})}\right) \text{ for } p = 1, ..., P$$

$$\gamma_t^{(p)} \sim N\left(\frac{\gamma_{t-1}^{(p)}}{1 + r_{\Phi,t-1}\left(1 + \gamma_{t-1}^{(p)}\right)}, \sigma_\gamma^2\right) \text{ for } p = 1, ..., P$$

Step 2: Resample asset weights and leverage values from

$$\left\{w_t^{(p)}, \gamma_t^{(p)}\right\} \sim Mult_P\left(\left\{\omega_t^{(\phi)}, w_t^{(\phi)}, \gamma_t^{(p)}\right\}_{\phi=1}^{P}\right)$$

where the importance weights are given by

$$\omega_t^{(p)} \propto \left(1 + \frac{1}{\nu}\frac{\left(r_{\Phi,t} - \left(1 + \gamma_t^{(p)}\right) w_t^{(p)\prime} r_{A,t}\right)^2}{\sigma_\epsilon^2}\right)^{-\frac{\nu+1}{2}}$$

FIGURE 2.3: *L-DPM Estimation Algorithm*

have been proposed by Liu and West (2001), Storvik (2002), and Fearnhead (2002). Due to the superior results of Lopes, Carvalho, Johannes, and Polson (2010), we use their Particle Learning concept to learn the values of parameters α, σ_ϵ^2, ν, and σ_γ^2 in an online manner. In order to do so, prior sampling distributions and suitable sufficient statistics need to be identified.

2.5.1 *Weight Transition Model*

The transition model requires a parameter α which controls the dispersion of the weights when forming the prior for the next time period. In order to sample this α parameter, it is necessary to derive the conjugate prior of the Dirichlet distribution. In the generalized form, the Dirichlet distribution is written:

$$p(x|\boldsymbol{\alpha}) = \frac{1}{B(\boldsymbol{\alpha})}\prod_{i=1}^{N} x_i^{\boldsymbol{\alpha}_i - 1}, \quad \text{where} \quad B(\boldsymbol{\alpha}) = \frac{\prod_{i=1}^{N}\Gamma(\alpha_i)}{\Gamma\left(\sum_{i=1}^{N}\alpha_i\right)}$$

The Dirichlet distribution is an exponential family distribution, and therefore can be written in the form:

$$p(x|\boldsymbol{\eta}) = h(x)\, g(\boldsymbol{\eta}) \exp\left(\boldsymbol{\eta}'\mathbf{T}(x)\right)$$

where the base measure, natural parameter, partition, and sufficient statistic values are given, respectively:

$$h(x) = 1, \quad \boldsymbol{\eta} = \boldsymbol{\alpha} - 1, \quad g(\boldsymbol{\eta}) = \frac{1}{B(\boldsymbol{\eta}+1)}, \quad \mathbf{T}(x) = \ln x$$

For any exponential family distribution, there exists a conjugate prior of the form:

$$\begin{aligned} p(\boldsymbol{\eta}|\boldsymbol{\chi}, n) &= f(\boldsymbol{\chi}, n)\, g(\boldsymbol{\eta})^n \exp(\boldsymbol{\eta}'\boldsymbol{\chi}) \\ &\propto g(\boldsymbol{\eta})^n \exp(\boldsymbol{\eta}'\boldsymbol{\chi}) \end{aligned}$$

where $f(\boldsymbol{\chi}, n)$ is a normalization constant. Therefore, a conjugate prior to the Dirichlet distribution is:

$$p(\boldsymbol{\alpha}|\boldsymbol{\chi}, n) \propto \frac{1}{B(\boldsymbol{\alpha})^n} \exp\left((\boldsymbol{\alpha}-1)'\boldsymbol{\chi}\right)$$

Now, as in Grunwald, Raftery, and Guttorp (1993), let us reparametrize the concentration parameter into the decomposition $\boldsymbol{\alpha} = \alpha\boldsymbol{\theta}$ where $\alpha \in \mathbb{R}_+$ specifies the dispersion and $\boldsymbol{\theta} \in \mathbb{S}^{N-1}$ specifies the location. Therefore, we can write the conditional for the dispersion parameter α as

$$p(\alpha|\boldsymbol{\chi}, n, \boldsymbol{\theta}) \propto \frac{1}{B(\alpha\boldsymbol{\theta})^n} \exp\left((\alpha\boldsymbol{\theta} - 1)' \boldsymbol{\chi}\right)$$

and the conditional for the location parameter $\boldsymbol{\theta}$ as

$$p(\boldsymbol{\theta}|\boldsymbol{\chi}, n, \alpha) \propto \frac{1}{B(\alpha\boldsymbol{\theta})^n} \exp\left((\alpha\boldsymbol{\theta} - 1)' \boldsymbol{\chi}\right).$$

This conditional distribution for $\boldsymbol{\theta}$ is similar to the Dirichlet Conjugate distribution first given in Grunwald, Raftery, and Guttorp. We note that theirs is adapted to the setup of filtering location values from noisy compositional observations. In contrast, we are interested in the other part of the concentration parameter decomposition, that is, estimating values of the dispersion parameter α. Therefore, we instead concern ourselves with $p(\alpha|\boldsymbol{\chi}, n, \boldsymbol{\theta})$.

In the DPM, we specify the location parameter $\boldsymbol{\theta}$ to be:

$$\begin{aligned}
\boldsymbol{\theta}_t &= w_{t|t-1} \\
&= E\left[w_t|\mathcal{F}_{t-1}\right] \\
&= \frac{w_{t-1} \circ (1 + r_{A,t-1})}{\sum_{i=1}^{N} w_{t-1,i}(1 + r_{A,t-1,i})}
\end{aligned}$$

Therefore, we want to update the prior on α given new observations. Using the derived conjugate prior distribution with the exponential distribution properties, the posterior with a single new observation x is obtained:

$$\begin{aligned}
p(\alpha|x, \boldsymbol{\chi}, n, \boldsymbol{\theta}) &= p(\alpha|\boldsymbol{\chi} + \mathbf{T}(x), n+1) \\
&\propto \frac{1}{B(\alpha\boldsymbol{\theta})^{n+1}} \exp\left((\alpha\boldsymbol{\theta} - 1)' (\boldsymbol{\chi} + \mathbf{T}(x))\right) \\
&= \frac{1}{B(\alpha\boldsymbol{\theta})^{n+1}} \exp\left((\alpha\boldsymbol{\theta} - 1)' (\boldsymbol{\chi} + \ln x)\right)
\end{aligned}$$

So, for a given set of initial model parameters, n_0 and $\boldsymbol{\chi}_0$, and a new portfolio weight posterior $w_{t|t} \in \mathbb{S}^{N-1}$, updating the sample statistics at each iteration is done recursively via:

$$n_t = n_{t-1} + 1$$

$$\boldsymbol{\chi}_t = \boldsymbol{\chi}_{t-1} + \ln w_{t|t}$$

where the posterior distribution for sampling α is given by:

$$p\left(\alpha \middle| w_{t|t}, \boldsymbol{\chi}_t, n_t, w_{t|t-1}\right) \propto \frac{1}{B(\alpha w_{t|t-1})^{n_t}} \exp\left(\left(\alpha w_{t|t-1} - 1\right)' \boldsymbol{\chi}_t\right)$$

This is not a common distribution, and so sampling from this can be done via Slice Sampling (Neal, 2003).

2.5.2 Leverage Transition Model

Similarly, we can construct a particle learning setup for the leverage transition model parameter σ_γ^2. Let us first consider the Gaussian random walk model with prior transitional variance given by s_γ^2. The conjugate prior on this parameter is the inverse Gamma distribution:

$$\mathcal{IG}\left(n_t/2, n_t s_{\gamma,t}^2/2\right)$$

For a given set of initial prior parameters, n_0 and $s_{\gamma,0}^2$, updating the sample statistics at each iteration is done recursively via:

$$n_t = n_{t-1} + 1$$

$$n_t s_{\gamma,t}^2 = n_{t-1} s_{\gamma,t-1}^2 + \left(\gamma_t - \frac{\gamma_{t-1}}{1 + r_{\Phi,t-1}\left(1 + \gamma_{t-1}\right)}\right)^2$$

2.5.3 Gaussian Observation Model

Finally, we can construct a particle learning setup for the observational model parameters. Before we consider our Student-t model, let us first examine a strictly Gaussian observational model with variance given by σ_ϵ^2. The conjugate prior on this parameter is the inverse Gamma distribution:

$$\mathcal{IG}\left(n_t/2, n_t s_{\epsilon,t}^2/2\right) \tag{2.5}$$

Similar to the leverage results above, for a given set of initial prior parameters, n_0 and $s_{\epsilon,0}^2$, updating the sample statistics at each iteration is done recursively via:

$$n_t = n_{t-1} + 1$$

$$n_t s_{\epsilon,t}^2 = n_{t-1} s_{\epsilon,t-1}^2 + \left(r_{HF,t} - w_t' r_{PA,t}\right)^2$$

2.5.4 Student-t Observation Model

Let us consider our Student-t observational model with degrees-of-freedom parameter ν and scale parameter σ_ϵ^2. The method proposed in Lopes and Polson (2012) for Particle Learning of location-scale Student-t distributions can be implemented here to perform sequential learning of these parameters.

Using a scale mixture of normals representation, we write the Student-t errors $\eta \sim t_\nu\left(0, \sigma_\epsilon^2\right)$ as:

$$\eta = \sigma_\epsilon \sqrt{\lambda_i} \epsilon \quad \text{where} \quad \left(\lambda_i | \nu\right) \sim IG\left(\nu/2, \nu/2\right) \text{ and } \epsilon \sim N\left(0, 1\right)$$

First, assume the following Jeffreys prior from Fonseca, Ferreira, and Migon (2008) on ν:

$$p\left(\nu\right) \propto \frac{1}{\sigma}\left(\frac{\nu}{\nu+3}\right)\left\{\psi'\left(\frac{\nu}{2}\right) - \psi'\left(\frac{\nu+1}{2}\right) - \frac{2\left(\nu+3\right)}{\nu\left(\nu+1\right)^2}\right\}^{1/2}$$

where $\psi'(a) = d\{\psi(a)\}/da$ and $\psi(a) = d\{\log\Gamma(a)\}/da$ are the trigamma and digamma functions, respectively.

Now, using the scale mixture of normals representation from above, Bayes rule yields the sampling distribution for ν:

$$p(\nu|\lambda) \equiv p(\nu|S_{t1}, S_{t2}) \propto p(\nu) \left(\frac{\left(\frac{\nu}{2}\right)^{\frac{\nu}{2}}}{\Gamma\left(\frac{\nu}{2}\right)} \right)^t S_{t1}^{-(\nu/2+1)} \exp\{-\nu S_{t2}/2\}$$

Given the set of initial values $S_{01} = 1$, $S_{02} = 0$, updating the sample statistics at each iteration is done via:

$$S_{t1} = S_{t-1,1}\lambda \quad \text{and} \quad S_{t2} = S_{t-1,2} + 1/\lambda$$

Similar to the posterior sampling distribution for the dispersion parameter α, draws from this distribution can be obtained via Slice Sampling.

Next, assuming the same prior in (2.5) on the scale parameter σ_ϵ^2, the sampling distribution for the scale parameter is:

$$\sigma_\epsilon^2 \sim \mathcal{IG}(S_{t3}/2, S_{t4}/2)$$

Given the set of initial values $S_{03} = n_0$ and $S_{04} = n_0 s_{\epsilon,0}^2$, updating the sample statistics at each iteration is done via:

$$S_{t3} = S_{t-1,3} + 1 \quad \text{and} \quad S_{t4} = S_{t-1,4} + \eta_t^2/\lambda$$

As well, the predictive resampling distribution and latent state conditional posterior for propagation are:

$$\eta_{t+1} \sim t_{S_{t3}+2}\left(0, \frac{S_{t4}}{S_{t3}+2}\lambda\right)$$

$$\lambda \sim \mathcal{IG}\left(\frac{\nu+1}{2}, \frac{\nu + \eta_t^2/\sigma_\epsilon^2}{2}\right)$$

Combining all of these Particle Learning concepts with the L-DPM yields the complete algorithm for the Leveraged Dirichlet Portfolio Model with Parameter Learning detailed in Figures 2.4 and 2.5.

2.6 Data and Estimation

In order to illustrate our methodology, we evaluate the time-varying asset class composition and leverage on an index of the broad hedge fund industry returns. Similar to initial OLS based decomposition technique that Sharpe (1992) used on the static decomposition of mutual funds, and Fung and Hsieh (1997) used on the static decomposition of hedge funds, we update their results to allow for time-varying portfolio weights in a more appropriately specified dynamic model rather than their rolling-window OLS approach. As well, since leverage is an important component of hedge fund portfolio returns, we have introduced a technique to perform joint estimation of time-varying leverage multiplier values as part of the dynamic model. This gives us the ability to capture a larger portion of the variation in hedge fund portfolio returns than the previous normalized constrained rolling OLS approach which simply estimated relative portfolio weights with no leveraging effect.

We examine the monthly return data for the Hedge Fund Research Fund Weighted Composite Index (HFRIFWC) from January 1995 to October 2012. This index is formed from monthly self-reported returns complied from a source of over 2,200 hedge funds with over $50 million USD under management and a track record of greater than 12 months. These return values are reported net of individual fund managers' fees. We note that there can be some biases present in these self-reported aggregated hedge fund indices including illiquidity-induced serial correlation in fund returns (Getmansky, Lo, and Makarov, 2004), backfill and incubation bias (Fung and Hsieh, 2004), and self-selection bias (Fung and Hsieh, 2000, 2009). Nevertheless,

L-DPM with Parameter Learning Estimation Algorithm - (Part 1 of 2)

Initialize:

Sample prior weights and leverage values from

$$w_0^{(p)} \sim Dir\left(\alpha_0\left(\frac{1}{N},...,\frac{1}{N}\right)'\right) \quad \text{and} \quad \gamma_0^{(p)} \sim Beta\left(a,b\right)$$

Choose initial values for particle learning parameters

$$n_{\alpha,0},\ \boldsymbol{\chi}_t,\ n_{\gamma,0},\ s^2_{\gamma,0},\ S_{01}=1,\ S_{02}=0,\ S_{03}=n_{\epsilon,0},\ S_{04}=n_{\epsilon,0}s^2_{\epsilon,0}$$

Iterate:

Step 1a: Sample mixture parameter:

$$\lambda^{(p)} \sim \mathcal{IG}\left(\nu^{(p)}/2, \nu^{(p)}/2\right)$$

Step 1b: Resample particles with predictive distribution importance weights given by

$$\omega_t^{(p)} \propto \left(1+\frac{1}{S_{t3}+2}\frac{\left(r_{\Phi,t+1}-\left(1+\gamma^{(p)}_{t+1|t}\right)w^{(p)\prime}_{t+1|t}r_{A,t+1}\right)^2}{\lambda^{(p)}S^{(p)}_{t4}/\left(S^{(p)}_{t3}+2\right)}\right)^{-\frac{S_{t3}+3}{2}}$$

where the one-step-ahead weight and leverage predictions $w_{t+1|t}$ and $\gamma_{t+1|t}$ are:

$$w^{(p)}_{t+1|t}=\frac{w_t^{(p)}\circ(1+r_{A,t})}{\sum_{i=1}^N w^{(p)}_{t,i}(1+r_{A,t,i})} \quad \text{and} \quad \gamma^{(p)}_{t+1|t}=\frac{\gamma_t^{(p)}}{1+r_{\Phi,t}\left(1+\gamma_t^{(p)}\right)}$$

Step 2: Propagate asset weights and leverage values from

$$w^{(p)}_{t+1} \sim Dir\left(\alpha\frac{w_t^{(p)}\circ(1+r_{A,t})}{\sum_{i=1}^N w^{(p)}_{t,i}(1+r_{A,t,i})}\right)$$

$$\text{and} \quad \gamma^{(p)}_{t+1} \sim N\left(\frac{\gamma_t^{(p)}}{1+r_{\Phi,t}\left(1+\gamma_t^{(p)}\right)}, \sigma^2_\gamma\right)$$

FIGURE 2.4: *L-DPM with Parameter Learning - (Part 1 of 2)*

in the absence of better information, this data remains a reasonable and commonly used proxy for the performance of the aggregate hedge fund industry portfolio.

L-DPM with Parameter Learning Estimation Algorithm - (Part 2 of 2)

Step 3a: Sample Mixture Parameter

$$\lambda^{(p)} \sim \mathcal{IG}\left(\frac{\nu^{(p)}+1}{2}, \frac{\nu^{(p)} + \left(r_{\Phi,t+1} - \left(1+\gamma_{t+1}^{(p)}\right) w_{t+1}^{(p)\prime} r_{A,t+1}\right)^2 / \sigma_\epsilon^{(p)2}}{2}\right)$$

Step 3b: Update sample statistics

$$n_{\alpha,t+1} = n_{\alpha,t} + 1, \quad n_{\gamma,t+1} = n_{\gamma,t} + 1, \quad S_{t+1,3} = S_{t,3} + 1,$$

$$n_{\gamma,t+1} s_{\gamma,t+1}^{(p)2} = n_{\gamma,t} s_{\gamma,t}^{(p)2} + \left(\gamma_{t+1}^{(p)} - \frac{\gamma_t^{(p)}}{1 + r_{\Phi,t}\left(1+\gamma_t^{(p)}\right)}\right)^2,$$

$$\boldsymbol{\chi}_{t+1}^{(p)} = \boldsymbol{\chi}_t^{(p)} + \ln w_{t+1}^{(p)}, \quad S_{t+1,1}^{(p)} = S_{t,1}^{(p)} \lambda^{(p)}, \quad S_{t+1,2}^{(p)} = S_{t,2}^{(p)} + 1/\lambda^{(p)},$$

$$\text{and} \quad S_{t+1,4}^{(p)} = S_{t,4}^{(p)} + \left(r_{\Phi,t+1} - \left(1+\gamma_{t+1}^{(p)}\right) w_{t+1}^{(p)\prime} r_{A,t+1}\right)^2 / \lambda^{(p)}$$

Step 4: Sample parameters

$$\alpha^{(p)} \sim p\left(\alpha^{(p)}\right) \propto \frac{1}{B(\alpha w_{t+1|t}^{(p)})^{n_{\alpha,t+1}}} \exp\left(\left(\alpha w_{t+1|t}^{(p)} - 1\right)' \boldsymbol{\chi}_{t+1}^{(p)}\right),$$

$$\nu^{(p)} \sim p\left(\nu^{(p)}\right) \propto p(\nu) \left(\frac{\left(\frac{\nu}{2}\right)^{\frac{\nu}{2}}}{\Gamma\left(\frac{\nu}{2}\right)}\right)^t S_{t1}^{(p)-(\nu/2+1)} \exp\left\{-\nu S_{t2}^{(p)}/2\right\},$$

$$\sigma_\gamma^{(p)^2} \sim \mathcal{IG}\left(n_{\gamma,t+1}/2, n_{\gamma,t+1} s_{\gamma,t+1}^{(p)}{}^2/2\right),$$

$$\text{and} \quad \sigma_\epsilon^{(p)^2} \sim \mathcal{IG}\left(S_{t+1,3}/2, S_{t+1,4}^{(p)}/2\right)$$

FIGURE 2.5: *L-DPM with Parameter Learning - (Part 2 of 2)*

For the portfolio asset classes, we use an almost identical set to those used in Fung and Hsieh (1997). We, however, break fixed income securities into municipal, corporate high yield, and short-term treasuries, since they find that municipal and high yield bond funds have low correlation with their set of asset classes. Consistent with their findings, the inclusion of these asset classes in our estimation procedure

Color	Asset Class Name
	Barclays Municipal Bond Index
	Barclays Short Term Treasury Bond Index
	Barclays Corporate High Yield Bond Index
	Deutsche Bank US Dollar Long Futures Index
	Dow Jones - UBS Commodity Index
	MSCI Emerging Markets Index
	MSCI EAFE Index (Europe, Australasia, Far East)
	MSCI US Equity Index

Table 2.1: *Asset Class Color Code*

significantly increases the amount of variation in returns explainable in the index for the aggregate hedge fund industry. Therefore, the returns on both municipal bonds and corporate high yield bonds are an important component of the hedge fund industry return profile. As with any estimation technique, we emphasize the importance of appropriate explanatory variables, and in this case, an appropriate set of portfolio assets due to the potential for omitted variable bias arising from correlation across asset returns. Table 2.1 enumerates the asset class list and color key. Note that it is not necessary to restrict ourselves to this specific set of asset classes. The methodology described above can be applied to any asset set of interest.

Finally, it is important to acknowledge the widespread discussion of hedge funds taking large short positions on various assets. Notably, the structure of this Dirichlet Portfolio Model family does not allow for negative weight estimation since values are restricted to the simplex. Although it may be easy for an individual hedge fund to take conceivably large short positions on a particular asset, historical data suggests that it is unlikely that the aggregate of all hedge funds holds net short positions on any asset class. England's Financial Conduct Authority, formerly the Financial Services Authority, conducts periodic holdings surveys on a subset of hedge funds. These quarterly reports confirm that although many funds bolster large short positions, the net exposures inside all these asset classes are almost always net-long. In the very rare cases where exposures are net-short, they are almost indistinguishable from zero. Therefore, the assumed net-long estimation structure is largely appropriate.

2.7 Hedge Fund Results

We now use the Leveraged Dirichlet Portfolio Model with Parameter Learning to estimate portfolio weights, leverage multiplier values, and tuning parameter values on the aggregate hedge fund return data. Then, using these portfolio weights and leverage values, the model is used to forecast the one-step-ahead predicted weights to create a time series of "forecasted" returns. These returns are plotted along with the observed index returns to demonstrate the out-of-sample forecasted accuracy of the estimation procedure.

Figure 2.6 shows the forecasted returns and estimated weight plots. The forecasted returns are very close to the actual observed returns in the next period. We observe a cyclical pattern in the estimated weights which transitions from equities in periods of economic growth to fixed income securities during recessions. As well, a large weight is placed on short-term treasury securities beginning in 2001. As previously identified in Korsos (2013b), artificially low volatility in the hedge fund return data, due to the return smoothing and reporting biases, causes larger weight to be estimated on low volatility assets, such as short-term treasuries. Consistent with the results of Korsos, if the desmoothing model of Getmansky, Lo, and Makarov is implemented to desmooth the reported returns, weight on treasuries dramatically decreases while the other asset class weights scale up proportionately in order to appropriately capture the larger hedge fund return volatility.

Figure 2.7 shows the estimated values for the leverage multiplier γ_t. Interestingly, we see a rather constant, but slowly decreasing value of leverage, suggesting that

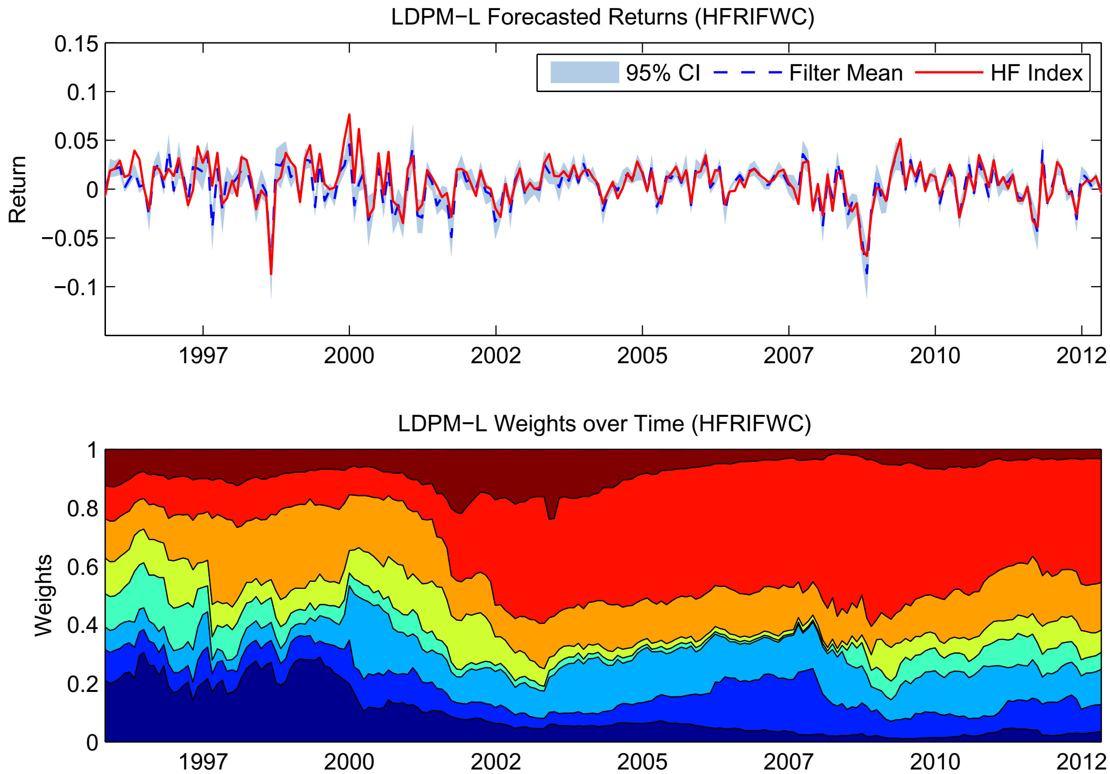

FIGURE 2.6: *Forecasted Returns & Estimated Weights*

although many funds may employ large amounts of gross leverage to obtain their distribution of returns, the net value of borrowing among this sample of the hedge fund industry is small.

At first, these small leverage values seem strikingly low due to the media imposed 'prior' manifested from the reported high gross leverage values. For example, in an April 2012 article in the *Financial Times*, Michael Hintze, chief executive of CQS, a \$9 billion London based hedge fund said, "hedge funds are presently leveraged one to three times; if they're mad, five times; if they're insane, 10 times." Recall that net and gross leverage values can differ dramatically. For example, suppose the hedge fund industry has around \$200 billion in short exposures and \$275 billion in long

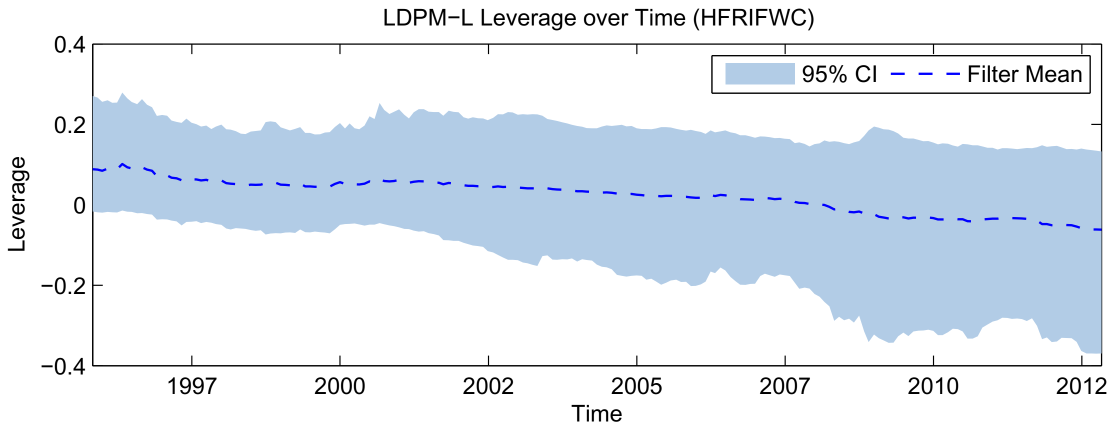

FIGURE 2.7: *Estimated Leverage Multipliers*

exposures to equities. Financing for much of the long positions can be offset via proceeds from the short positions. To keep this example simple, assuming negligibly small borrowing costs and margin requirements, \$200 billion of the long positions are financed with \$200 billion of proceeds from the shorts. This leaves \$75 billion remaining to finance in long equity positions. If this entire amount is financed via contributed capital, this yields a gross leverage multiple of $(200+200+75)/75 = 6.33$, whereas net leverage is simply $(-200 + 200 + 75)/75 = 1$, that is non-existent. Financing even \$10 billion of these equity positions via borrowing facilities still yields similar multiplier results of $(200 + 200 + 75)/65 = 7.31$ for gross leverage and only $(-200+200+75)/65 = 1.15$ for net leverage. That is, with \$10 billion of borrowing, gross leverage appears as 631% above contributed capital, when net leverage truly is only 15%.

The particle learning estimation for the tuning parameters is illustrated similarly. Figures 2.8, 2.9, 2.10, 2.11 show the progression of sequentially updating credible regions for each of the tuning parameters α, σ_γ, σ_ϵ, and ν.

Table 2.2 exhibits summary statistics for the monthly returns of the index and

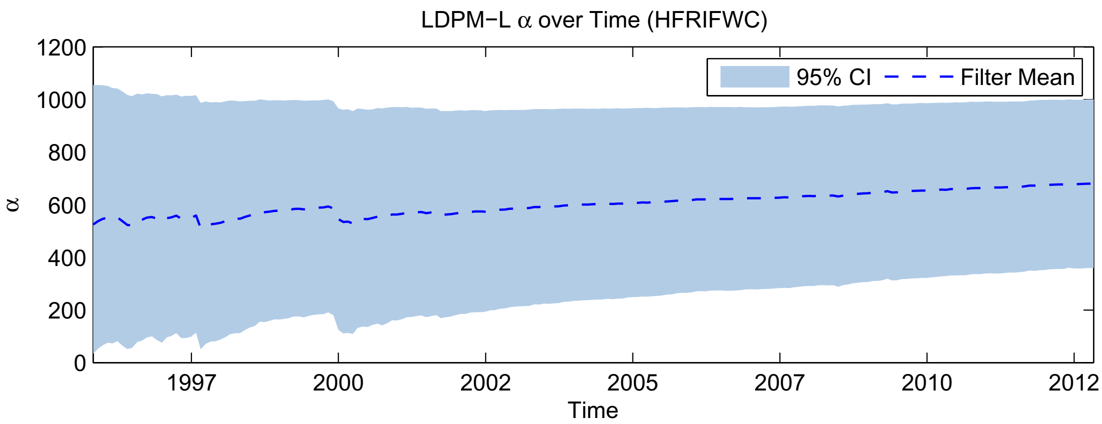

FIGURE 2.8: *Estimated α Value*

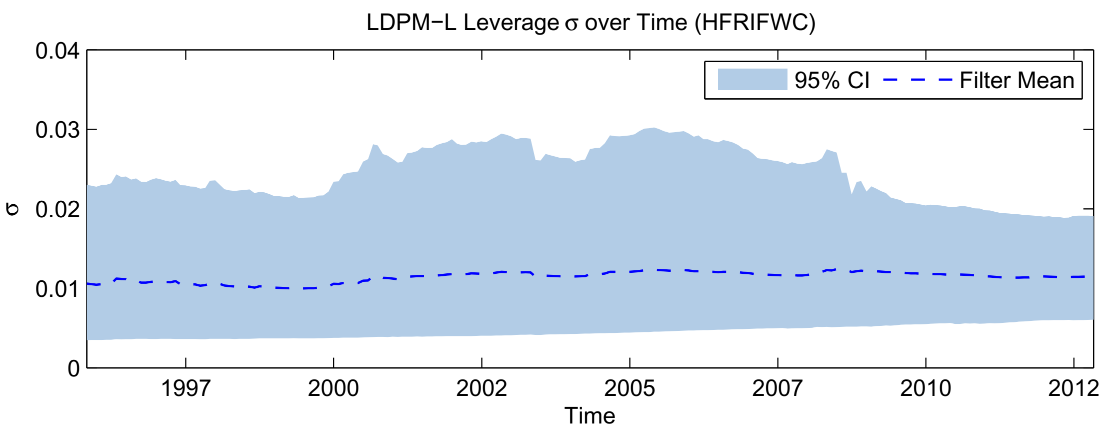

FIGURE 2.9: *Estimated σ_γ Value*

the returns implied by the weights and leverage of the Leveraged Dirichlet Portfolio Model with Parameter Learning (LDPM-PL) estimation. The 'Same-Period' return values are constructed with the sequential posterior estimates, thereby indicating explanatory power of the model given the compositional assets. The LDPM-PL's dynamic weight and leverage processes capture over 85% of the in sample variation in hedge fund returns. Notice that these sequential estimates do not include future period data, therefore not even using the entire sample period of data at each time

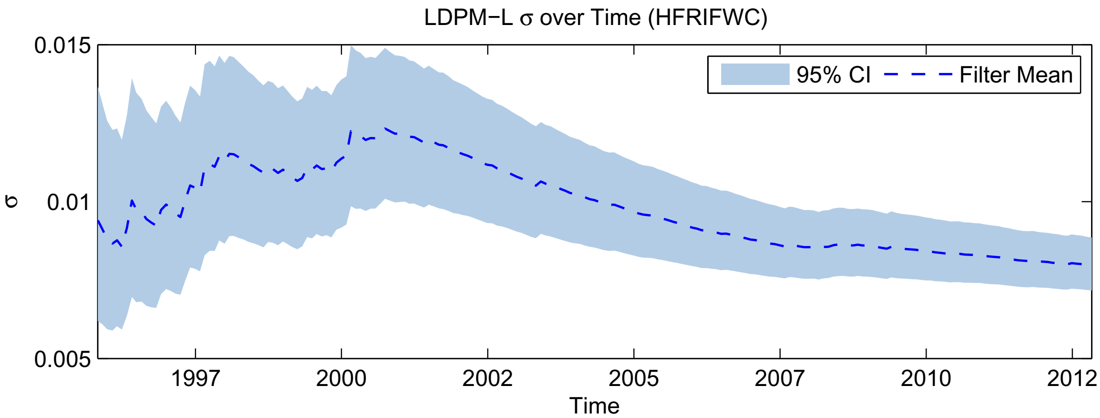

FIGURE 2.10: *Estimated σ_ϵ Value*

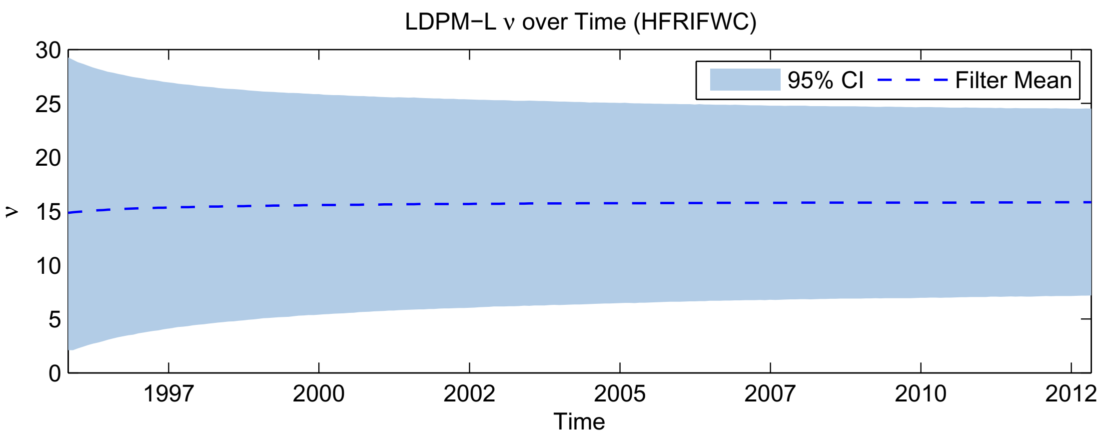

FIGURE 2.11: *Estimated ν Value*

period, as is done in many historical OLS based approaches. If the entire sample period data was to be used to refine the estimates at all periods, the fixed-interval particle smoothing technique of Carvalho, Johannes, Lopes, and Polson (2010) can be implemented here. However, note that since this algorithm costs $\mathcal{O}\left(P^2T\right)$ in computation time, where P is the number of particles used and T is the number of time periods, this can be prohibitively costly for commonly large values of P. The 'Forecasted' return values are constructed with the sequential prior estimates, thereby exhibiting the predictive power of the model at over 77% of future hedge fund return variation. Lastly, we point out that although the mean values of the estimated returns are noticeably below those of the original index, this effect arises from the aforementioned biases and return smoothing in the self-reported data. Applying a desmoothing model to the original return data, shifts the weights into higher volatility and higher returning assets, thereby decreasing, but not completely eliminating this gap. The remaining difference can be interpreted as the effect of reporting biases in the data. This is consistent with Malkiel and Saha (2005) and Jurek and Stafford (2012) who report the annualized effects of these biases from as low as 3% to potentially over 7% for various sources of hedge fund data. For comparison to our results, that would imply a monthly upward bias of between 0.25% to 0.583% in the mean return.

		LDPM-PL	
	HFRIFWC	**Same-Period**	**Forecasted**
Mean	0.00748	0.00490	0.00427
Standard Deviation	0.02100	0.01992	0.02057
RMSE	0	0.00837	0.01060
Mean Abs Error	0	0.00576	0.00750
Correlation	1	0.92424	0.87996
R^2	1	0.85423	0.77434

Table 2.2: *Estimation Summary Statistics (Monthly Returns)*

2.8 Net Portfolio Leverage by Fund Strategy

Since hedge funds and their respective managers can be very heterogeneous in their expertise and trading styles, they are commonly classified by the broad type of strategy they employ. These classifications regularly include Long/Short Equity, Quantitative, Event Driven, Macro, and Relative Value strategies. To gain insight into the behavior of different types of hedge funds, we examine how net portfolio leverage differs across the various types of these broad strategies.

In addition to the previously detailed 'Composite' index, we use corresponding return indices from Hedge Fund Research for each of the following hedge fund classifications: Equity Long/Short, Quantitative, Event Driven, Macro, and Relative Value. Similar to the composite index, these indices are formed from the same sample of 2,200 hedge funds with over $50 million USD under management. The Equity Long/Short index consists of funds maintaining both long and short positions primarily in equity and equity derivative securities. The Quantitative index consists of funds which use various forms of quantitative decision processes to select securities for purchase or sale. The Event Driven index consists of funds which hold securities of firms currently or prospectively involved in corporate transactions including mergers, restructurings, financial distress, tender offers, shareholder buybacks, debt exchanges, security issuance or other capital structure adjustments. The Macro index includes funds whose trading decisions are primarily influenced by movements in economic variables and their resulting effects on related securities. Finally, the Relative Value index consists of funds who attempt to profit off discrepancies in the relationship between multiple securities.

Running the L-DPM on each of these hedge fund return indices gives the following net portfolio leverage results in Figure 2.12. From this plot, while most funds tend to have very similar net portfolio leverage values to the composite, two styles stand out. The Relative Value funds exhibit a significant increase in leverage starting in 2000 and growing to almost double contributed capital in 2006. In 2007-2008 these funds experienced a sharp decline in leverage, but still remain higher than any of the other strategies. Since the majority of relative value strategies are concentrated in fixed income securities, it is very easy for hedge funds to obtain 'synthetic' leverage on

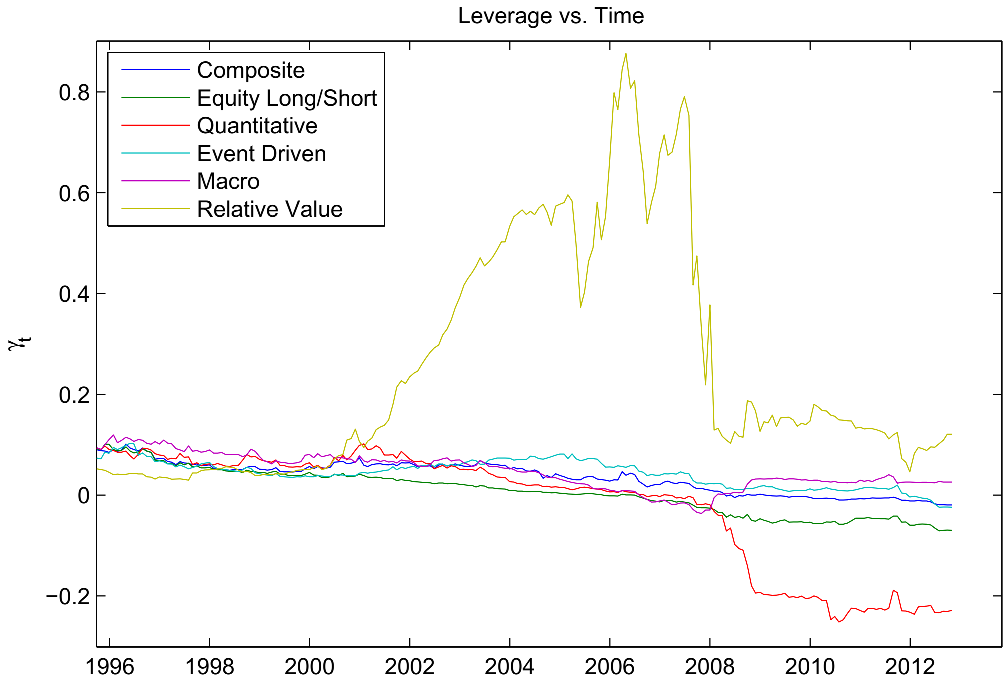

FIGURE 2.12: *Leverage Comparison by Hedge Fund Strategy*

these assets through the use of futures securities which have the effect of magnifying exposures to underlying assets without the need to borrow capital in the conventional ways. That is, the use of these fixed income futures securities allow funds to obtain greater notional exposure than their contributed capital without the need for repurchase agreements, prime brokerage borrowing facilities, unsecured borrowing, or traditional term financing, like the issuance of public debt. Consistent with this, the compositional weights are almost completely proportioned toward fixed income assets.

The other notable standout is the Quantitative style. These funds appear very similar to the composite until 2008, when their net portfolio leverage drops by about 20% and remains at that level. The estimated compositional portfolio weights also suggest net exposures of 70-80% in equities until 2008, when equities drop to about

60%, being replaced primarily by short-term treasuries. This deleveraging, combined with a significant increase in risk-free securities, suggests that quantitative hedge funds became more risk averse during the 2008 financial crisis, and that this cautious risk aversion still persists.

Finally, the remaining three fund classifications are very similar to the overall composite with regard to net portfolio leverage. The Equity Long/Short funds' compositional weights are proportionately similar to the composite, with a 20% absolute decrease in short-term treasuries. The Macro funds exhibit large investment swings between fixed income, equity, and currencies. The majority of these exposures are to the various fixed income securities except for high yield corporate bonds, since their shorter maturities cause lower sensitivity to macroeconomic information in exchange for increased sensitivity to firm-specific information. As expected, the Event Driven funds have a significantly larger weight placed on high yield corporate bonds, due to their sensitivity to corporate transactions and capital structure adjustments.

2.9 Estimation of Investment Effects

As previously detailed, due to the hidden nature of hedge fund investments, we may never observe true portfolio holdings. Small glimpses of long positions in equities and call/put options are available for some of the larger funds via 13D, 13G, and 13F SEC filings. Nevertheless, this is only a very small proportion of an individual fund's holdings, let alone those for the aggregate industry portfolio. Investment in fixed income, currencies, commodities, and even short positions do not have to be disclosed. Therefore, although it would be ideal to measure the accuracy of the portfolio estimation against the true portfolio compositions, this is not feasible. Instead, since changes in investment allocation through trading affects the demand for those securities, contemporaneous price effects can be observed due to the expected market microstructure impact on those asset prices (O'Hara, 1995; Madhavan, 2000; Hasbrouck, 2007). That is, if investment allocation is increased to a specific class of security, a positive expected price effect, over the same time period, would be consistent with this increase in asset demand.

In order to estimate the effect of changes in investment allocation on contemporaneous returns, we identify that portfolio weight changes arise from two causes: capital appreciation/depreciation of the compositional assets over a holding period and time varying asset allocation. Since the proportion of portfolio weight changes caused by capital appreciation/depreciation is not an investment decision controllable by a portfolio manager, we concern ourselves only with portfolio weight changes caused by active asset allocation trading decisions. In terms of our notation, the portfolio weight change over time t caused by these trading decisions is given by $w_{t|t} - w_{t|t-1}$, where $w_{t|t} \equiv w_t$ is the estimated portfolio weight at t and $w_{t|t-1}$ is the portfolio weight at t given a strict buy-and-hold transition from the previous period's estimated weight $w_{t-1|t-1}$. Therefore, this difference is the change in investment over period t due to an active decision to adjust portfolio allocation. In terms of the dynamic model, this difference is also thought of as the innovations on the weight transition process.

In the context of our portfolio estimation with leverage multiplier, this portfolio weight change is similarly given by $\left(1+\gamma_{t|t}\right) w_{t|t} - \left(1+\gamma_{t|t-1}\right) w_{t|t-1}$. We can now estimate the contemporaneous effect of asset allocation decisions on component asset i via:

$$r_{t,i} = \beta_{0,i} + \beta_{1,i}\left(\left(1+\gamma_{t|t}\right) w_{t,i|t} - \left(1+\gamma_{t|t-1}\right) w_{t,i|t-1}\right) + \xi_{t,i}$$

where

$$w_{t,i|t-1} = \frac{w_{t-1,i}(1+r_{A,t-1,i})}{\sum_{i=1}^{N} w_{t-1,i}(1+r_{A,t-1,i})} \quad \text{and} \quad \gamma_{t|t-1} = \frac{\gamma_{t-1}}{1+r_{\Phi,t-1}\left(1+\gamma_{t-1}\right)}$$

Table 2.3 exhibits OLS estimation details on each of the asset classes using the results from the L-DPM with Parameter Learning. We focus attention on the β_1 coefficients, which indicate the effects of active trading changes in the hedge fund industry portfolio on the same-period return in each asset class. Specifically, since actively increasing holdings in an asset class increases same-period demand for those assets, we would therefore expect a positive effect on prices and returns. We find that all of the estimated coefficients are positive, except short-term treasuries, with EAFE Equity, EM Equity, US Dollar Futures, and Municipal Bonds, significant at the 1% level, and Commodities significant at the 10% level. As well, we point out that although the coefficient on High Yield corporate bonds is not largely significant, these assets are commonly subject to illiquidity-induced return smoothing by fund managers (Fisher et al., 2003; Kadlec and Patterson, 1999). We note that if the desmoothing model proposed in Getmansky, Lo, and Makarov is applied to the original hedge fund data and our model is re-estimated, this coefficient on High Yield corporate bonds becomes significant at the 1% level with no major changes to those for the other assets.

2.10 Conclusion

This paper has extended a set of dynamic models and associated estimation procedures, based upon the foundational Dirichlet Portfolio Model (DPM) of Korsos (2013b). The Leveraged Dirichlet Portfolio Model (LDPM) and the Leveraged Dirichlet Portfolio Model with Parameter Learning (LDPM-PL) are valuable for estimation of both the unobserved time-varying portfolio compositions and also the net portfolio leverage of an aggregate portfolio return series. These techniques allow for a convenient decomposition of the major portfolio components in a different approach than the conventional rolling-window OLS approach on a set of risk factors.

Same-Period Return Effects of Active Trading

	US Equity	EAFE Equity	EM Equity	Commodities
β_0	0.00750**	0.00494	0.00812*	0.00308
	(0.00314)	(0.00332)	(0.00469)	(0.00321)
β_1	0.18587	0.92149***	1.45177***	0.65422*
	(0.19940)	(0.27040)	(0.36585)	(0.36460)
R^2	0.00410	0.05217	0.06945	0.01503

$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$

Same-Period Return Effects of Active Trading

	US Dollar	High Yield	Treasuries	Municipal
β_0	0.00206	0.00695***	0.00276***	0.00521***
	(0.00161)	(0.00187)	(0.00014)	(0.00079)
β_1	0.69524***	0.06179	−0.00035	0.36078***
	(0.14743)	(0.11827)	(0.00852)	(0.07834)
R^2	0.09535	0.00129	0.00001	0.09133

$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$

Table 2.3: *Same-Period Return Effects of Active Trading*

Importantly, these approaches not only allow for the identification of the sources of portfolio risk, but also time-varying estimates of the portfolio's investable holdings as they evolve through active trading decisions.

We then motivated this technique with the estimation of asset class weights on an index proxying the aggregate return of the hedge fund industry. We found that our dynamic weight and leverage processes captured over 85% of the variation in the hedge fund return index. As well, a discernible cyclical pattern is observed in the estimated weights which transitions from equities in periods of economic growth to fixed income securities during recessions. Comparing results across different classifications of hedge funds, we found that relative value hedge funds tend to have the highest net portfolio leverage levels due to the ease of obtaining magnified exposures using fixed income futures. Finally, we exhibited the accuracy of this technique by estimating asset class level regressions on the asset class return values against the same-period changes in the portfolio investment weights. The strong significance results are consistent with the expected market microstructure effects on prices for those assets, thereby bolstering the effectiveness of these portfolio decomposition techniques.

Finally, we note that this Dirichlet Portfolio Model family can be used in any environment where underlying compositions of portfolios may not be observed over time, but potentially noisy total portfolio returns are available. For example, consider the mutual fund industry. Exact portfolio holdings are available quarterly, however most changes in these holdings are unobservable until the next reporting period. Since daily return data on the portfolio is available, the DPM family can provide intra-reporting-period portfolio weight estimation for not only detection of when funds are changing their portfolio compositions, but also identification of where they are moving their assets.

3

Decomposing Hedge Fund Portfolios: Asset Allocation Choice & Its Role as a Volatility Buffer

3.1 Abstract

Unlike for mutual funds, detailed reports of hedge fund portfolio holdings are never accurately available via SEC filings except for long positions in equities and equity options for funds with over $100 million net in those assets. Since the complete picture of hedge fund holdings is not observable, this presents a significant analytical hurdle for more detailed analysis. To overcome this, we use a compositional state space model to estimate the dynamics of both asset class level portfolio holdings and leverage values on an index of hedge fund industry returns from 1995 to 2012. With estimates of these portfolio holdings, we find that net leverage levels in the hedge fund industry are smaller than popular belief due to netting both internally and across different funds. As well, using these estimates, we confirm previous findings that hedge funds do not contribute to herding behavior in most asset classes, and in

fact exhibit negative-feedback trading behavior in oil and municipal bonds.

3.2 Introduction

Although more transparent than they once were, hedge funds are a veritable "black box" of investing since outsiders, as well as current invested clients, may never observe the the entire true composition of a fund's portfolio. On the other hand, mutual funds present a much less opaque investment vehicle, where exact portfolio holdings can be accurately measured each quarter via SEC filings. This allows for straight-forward analysis to be performed using the compositions of these mutual fund portfolios. Unfortunately, since these compositions are not available for hedge fund portfolios, this manifests a significant analytical hole, and therefore it remains an econometric challenge to infer economic effects from the features of these changing portfolios.

Nevertheless, the hedge fund industry does provide a very narrow glimpse into their portfolio holdings via quarterly 13F SEC filings. These reports, required of investment managers with at least $100 million in equity assets under management, detail long positions in US equities, American Depository Receipts (ADRs), convertible notes, and call and put options. As well, 13D or 13G filings are required of any entity who obtains beneficial ownership of at least 5% of any class of publicly traded securities in a publicly company. We point out that these filings do not require reporting of any short positions in equities, or even any detail on fixed income, commodities, currencies, etc. Furthermore, since 13Fs are only required of funds with at least $100 million in *equity* assets under management, large hedge funds dealing primarily in the aforementioned, non-equity asset classes may not even have to report their share of equity holdings. Lastly, since there is mounting evidence that hedge funds employ large amounts of leverage to achieve their distribution of returns, funds can have equity holdings which are multiples greater than $100 million, while having less than \$100 million in equity assets under management, thereby avoiding the requirement to report holdings. Hence, these filings only provide a very small glimpse into the overall portfolio composition of hedge funds, thereby necessitating an alternative method to determine time-varying portfolio compositions.

The idea of decomposing investment fund compositions originated with Sharpe (1992), where he uses constrained OLS estimation to obtain asset class level portfolio weights on individual mutual funds. He estimates both time-invariant weights on the entire sample period, as well as demonstrates rolling window OLS regressions to obtain time varying values. Fung and Hsieh (1997) extend this "style analysis" of Sharpe to a different set of asset classes for both mutual funds and hedge funds. In doing so, they observe that static asset class weights on hedge fund portfolios have much lower in-sample explanatory power as compared to those for mutual funds. This indicates dynamic asset allocation/trading strategies play a larger role in hedge fund than mutual fund portfolios. In order to attribute the unexplained variation of these hedge fund portfolios, Fung and Hsieh (2001) construct a set of Primitive Trends Following Strategies (PTFS) which are used to explain common variation in returns in the cross-section of hedge funds. These PTFS risk factors are constructed to account for the dynamic nature of hedge fund weights over time due to the restrictive estimation procedure assuming static weight coefficients. To overcome the static weight hurdle, and to demonstrate the benefits of dynamic models on factor loading estimation, Mamaysky, Spiegel, and Zhang (2007) introduce a Kalman Filter based approach to obtain time varying risk factor coefficients on mutual funds. Our approach recognizes these benefits and extends this idea to directly estimate the time varying nature of portfolio weights in the spirit of Sharpe's rolling window OLS. However, we use advances in sequential Monte Carlo methods for dynamic model estimation, thereby also releasing us from the restrictions of the Kalman Filter, and allowing for a more appropriately specified model.

As Fung and Hsieh (1997) identify, hedge fund trading strategies are much more dynamic than those of mutual funds. Therefore, static buy-and-hold models largely cannot explain the variation in returns. Instead, this suggests the use of a model allowing for time-varying weight dynamics. Doing so enables us to capture the variation in hedge fund returns attributable to time-varying asset allocation, in addition to the variation in underlying asset returns. We implement the Leveraged Dirichlet Portfolio Model (L-DPM) from Korsos (2013a) to model both the normalized portfolio weights, as well as a time varying leverage scaling parameter. Due to the sequential nature of the estimation procedure, this allows us to not only explain the in-sample variation in hedge fund returns, but also create out-of-sample forecasts at each period of hedge fund asset class holdings in order to evaluate the forecasted tracking ability of the dynamic weight process.

First, we apply this estimation procedure to a set of actively managed diversified equity (AMDE) mutual funds, constructed similar to that in Kacperczyk, Sialm, and Zheng (2005), in order to obtain a "style analysis" across a set of equity industries. This is done by forming a value weighted return series on the set of AMDE mutual funds, and industry portfolio returns on each respective industry of interest. Then, time-varying industry style weights are estimated. Unlike previous work, we then compare these estimated style weights to the true time-varying industry weights implied by the set of AMDE mutual funds' holdings to evaluate the accuracy of the style analysis. We find that the dynamic model's allowance for time variation in industry weights produces explanatory power of 98.8% for these mutual funds on a set of 5 value-weighted industry portfolios.

Motivated by these results, asset class weights are estimated on the Hedge Fund Research (HFRI) Fund Weighted Composite Index, representing a proxy on the returns on the hedge fund industry. As expected, we find a discernible cyclical pattern in the estimated asset class weights with transitions from equities in periods of economic expansion, to fixed income securities during economic contraction. On the other hand, we find a lack of statistical support for large net leverage in the hedge fund industry, suggesting that although some funds may employ large amounts of gross leverage to achieve their distribution of returns, the net amount of leverage aggregated internally and across all funds is smaller than popular belief. However, this is consistent with various hedge fund surveys from England's Financial Conduct Authority and the European Central Bank.

Finally, we use these estimated asset class weights to assess whether hedge fund trading exhibits herding behavior in individual asset classes. That is, we pose the question: Do funds contribute to increased asset price volatility via selling as prices decrease or buying as prices increase? Since our estimation technique produces estimates of portfolio holdings, these can be used to form changes in portfolio weights for analysis of asset class level trading behavior. Consistent with previous findings by Kodres and Pritsker (1996) for futures transactions, we find that funds do not exhibit herding-like behavior across most asset classes. In fact, evidence of the opposite of herding, or negative-feedback trading, is found in oil and municipal bonds. By acting as a ready counterparty in directional markets, hedge funds increase liquidity and thereby contribute to decreased volatility in these assets (Morris and Shin, 1999; Persaud, 2000; Shiller, 1990).

The remainder of this paper is structured as follows: Section 2 describes the dynamic portfolio model and estimation procedure used for determining time-varying portfolio weights and leverage values. Section 3 outlines the mutual fund and hedge fund data used for estimation. Section 4 exhibits the resulting accuracy of the portfolio weight estimation on the mutual fund data. Section 5 presents the asset class portfolio weight estimation results for the hedge fund data. Section 6 analyzes these results as they relate to the various hedge fund surveys. Section 7 estimates the magnitude of herding behavior by hedge funds using the estimated portfolio compositions. Finally, section 8 concludes.

3.3 Dynamic Portfolio Model and Estimation

Consider the portfolio weight estimation problem where both the portfolio returns and compositional asset returns are observable, however the relative investment weights on those assets are never completely observable to those outside the firm. As in Korsos (2013a, 2013b), let there be an known investable set of N assets or asset classes with time-varying latent compositional weights $w_{t,i}$ for each asset i at time t which are required satisfy a normalizing budgetary restriction $\sum_{i=1}^{N} w_{t,i} = 1$. As Sharpe (1992) initially proposed, this restriction may be imposed directly on the least squares optimization. However, since OLS regression techniques directly assume a static model on the estimated weight coefficients, even rolling-window approaches to estimate time-varying values suffer from model misspecification. Since the entire aim of the estimation exercise is to infer how the weights are changing over time, we assume a dynamic model to be true *a priori*, therefore leading to a violation of the Gauss-Markov theorem. Hence, inference from these static OLS approaches implies an *a priori* bias and inefficiency in the resulting estimates. This is strong motivation for a more suitably specified dynamic model for estimation.

Mamaysky, Spiegel, and Zhang (2007) identify the estimation benefits of specifying a dynamic model and appropriate filtering methodology for time-varying risk factor identification. Since portfolio risk factor exposures can take on values in the reals $\mathbb{R}^N$, a classic Normal-Normal dynamic linear model, solvable via Kalman (1960) is an easily justifiable approach. In order to model relative portfolio weights which take values on the simplex $\mathbb{S}^{N-1} = \left\{ w \in \mathbb{R}_{+}^{N} : w'\mathbf{1} = 1 \right\}$, we need to constrain the estimation space. Korsos (2013b) details various ways that this constraint can be applied to dynamic models. First, this can be imposed directly on the period-by-period posterior weight optimization problem in a simple multivariate normal transition model. Second, since the restriction on the simplex implies $N-1$ degrees of freedom, where N is the number of component assets, the following restriction can be placed on the transitional covariance matrix: $\sum_{i=1}^{N} Cov\left(w_{t,j}, w_{t,i}\right) = 0, \forall j \in \{1, \ldots, N\}$. Korsos (2013b) shows that although these techniques seem convenient, they still require specification of $\mathcal{O}\left(N^2\right)$ covariance parameters, which for large numbers of explanatory compositional assets can be a challenging task. Instead, directly choosing a distribution with support on the simplex provides a natural way to model these proportions without needing to impose additional restrictions. As well, the choice of a Dirichlet distribution only requires the specification of a single univariate dispersion parameter, controlling the prior period-to-period variability in the portfolio weight transitions.

Since many fund managers commonly employ the use of leverage in their trading strategies, we also wish to jointly identify this effect on portfolio returns. Similar to the relative portfolio weights, the dynamics of this net leverage amount is also unobservable. We capture this net leverage by defining a time-varying value γ_t which specifies a multiplier on the magnitude of portfolio's net holdings over-and-above the portfolio's contributed capital. That is, if $\gamma = 0.2$ and a manager has \$100 million in contributed capital, the fund has net borrowings of \$20 million in order to hold a total net value of \$120 million of financial assets. Hence, $\gamma = 0$ indicates no leverage.

For convenience, let $\mathcal{F}_t$ represent the filtering of all information known at time t. This includes all previous portfolio returns r_Φ, compositional asset returns r_A, compositional asset weights w, and leverage scaling values γ up to and including time t. That is,

$$\mathcal{F}_t = \{r_{\Phi,1}, \ldots, r_{\Phi,t}, r_{A,1}, \ldots, r_{A,t}, w_1, \ldots, w_t, \gamma_1, \ldots, \gamma_t\}.$$

Define $w_t = (w_{t,1}, w_{t,2}, ..., w_{t,N})'$ to be an $N \times 1$ vector of the weights on each asset at the *beginning* of time period t, γ_t to be a scalar leverage multiplier also at the *beginning* of time period t, $r_{A,t} = (r_{A,t,1}, r_{A,t,2}, .., r_{A,t,N})'$ to be an $N \times 1$ vector of the compositional asset returns over time period t, and $r_{\Phi,t}$ to be a scalar value of the return on the portfolio over the same time period t. The chronology of the time period notation is illustrated below:

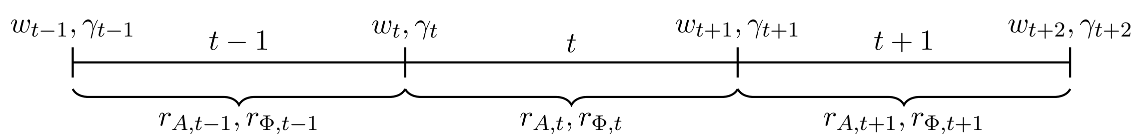

FIGURE 3.1: *Timeline Notation Illustration*

Following the Leveraged Dirichlet Portfolio Model (L-DPM) setup motivated in Korsos (2013a), we specify the following compositional weight transition model, leverage multiplier transition model, and observation model, respectively:

$$w_{t+1} \sim Dir\left(\alpha \frac{w_t \circ (1 + r_{A,t})}{\sum_{i=1}^{N} w_{t,i}(1 + r_{A,t,i})}\right) \tag{3.1}$$

$$\gamma_{t+1} \sim N\left(\frac{\gamma_t}{1 + r_{\Phi,t}\,(1 + \gamma_{t-1})}, \sigma_\gamma^2\right) \tag{3.2}$$

$$r_{\Phi,t} \sim t\left((1 + \gamma_t)\, w_t' r_{A,t}, \sigma_\epsilon^2, \nu\right) \tag{3.3}$$

First, the compositional weight transition model in (3.1) is a Dirichlet random walk model adjusted for capital appreciation of the invested assets. Plainly, this is a random walk model restricted to the simplex by decreasing directional variance as a respective boundary is approached. We note that since α is a scalar parameter controlling the prior dispersion of the period-to-period weight transitions, the one-step-ahead expectation is the unadjusted portfolio holdings given initial portfolio weights w_t and realized holding period asset returns $r_{A,t}$:

$$E\left[w_{t+1}|\mathcal{F}_t\right] = \frac{w_t \circ (1 + r_{A,t})}{\sum_{i=1}^{N} w_{t,i}(1 + r_{A,t,i})}$$

This provides the desired capital appreciation adjusted random walk property, while the scalar parameter $\alpha \in (0, \infty)$ can be thought of as the speed at which managers adjust their portfolio holdings. That is, smaller values of α represent faster changes in portfolio weights.

Although it is common for many asset managers to rebalance their portfolio holdings at predetermined periods such that $E\left[w_{t+1}|\mathcal{F}_t\right] = w_t$, this is part of the investment process choice. Therefore it is not assumed for the broad range of portfolios at every observable period. Nevertheless, if a manager does make the choice to rebalance asset weights, this effect will be picked up in the portfolio weight estimation results. As well, since different managers make different choices regarding their rebalancing intervals, this is further motivation to not make a prior rebalancing assumption. Ultimately, we are interested in estimating the *active* changes in portfolio weights due to *trading* activity, and since rebalancing a portfolio's investment weights involves active trading, these are exactly the effects we are trying to identify.

Second, the leverage multiplier transition model in (3.2) is a Gaussian random walk model, also adjusted for capital appreciation. As a portfolio increases in value due to capital appreciation of the component assets, the effective leverage of the portfolio decreases, and vice versa. Just as we want to capture active trading decisions on the relative portfolio weights, we take a similar approach to the leverage values. That is, our prior at each period is that the manager has done nothing to change their portfolio since the previous period, thereby allowing the Bayes rule filtering updates to reflect only the active trading decisions.

Lastly, the portfolio observation model in (3.3) is a location-scale Student-t model centered on the linear combination of the portfolio weights and the respective holding period returns for those assets, scaled by the leverage multiplier. It is reasonable that the set of chosen explanatory assets may not fully capture the variation in portfolio returns since they may not span the investment space of the portfolio of interest. Therefore, potential for leptokurtic error is introduced into the observed weighted portfolio return because of missing explanatory assets (Mandlebrot, 1963).

Estimation of the latent portfolio weights and leverage values is preformed via a sequential Monte Carlo technique. Even though the leverage transition model is Gaussian, the weight transition model and the observation model are not, therefore both the latent portfolio weights and leverage values cannot be solved for in closed form via Kalman (1960). Instead, the sequential Monte Carlo approach of Gordon, Salmond, and Smith (1993) is used to numerically solve for the latent portfolio compositional weights and leverage values. This Sampling Importance Resampling (SIR) estimation algorithm, adapted to the L-DPM, is exhibited in Figure 3.2.

3.4 Data

3.4.1 Mutual Fund Data

In order to compel the hedge fund asset class weight estimation results, we identify a similar portfolio tracking environment where true portfolio weights are available at discrete points in time. Mutual fund investments provide quarterly asset allocation data via the various required SEC forms 13F, 13D, 13G, N-1A, N-30D, N-CSR, and N-Q. Therefore, by first assuming that the mutual fund portfolio weights are completely unobservable, we construct weight estimates similarly to our hedge fund

L-DPM Estimation Algorithm

Initialize: Sample prior weights and leverage values from

$$w_0^{(p)} \sim Dir\left(\alpha_0 \left(\frac{1}{N}, ..., \frac{1}{N}\right)'\right) \quad \text{and} \quad \gamma_0^{(p)} \sim Beta\,(a, b)$$

Iterate:

Step 1: Propagate new asset weights and leverage values from

$$w_t^{(p)} \sim Dir\left(\alpha \frac{w_{t-1}^{(p)} \circ (1 + r_{A,t-1})}{\sum_{i=1}^{N} w_{t-1,i}^{(p)}(1 + r_{A,t-1,i})}\right) \text{ for } p = 1, ..., P$$

$$\gamma_t^{(p)} \sim N\left(\frac{\gamma_{t-1}^{(p)}}{1 + r_{\Phi,t-1}\,(1 + \gamma_{t-1})}, \sigma_\gamma^2\right) \text{ for } p = 1, ..., P$$

Step 2: Resample asset weights and leverage values from

$$\left\{w_t^{(p)}, \gamma_t^{(p)}\right\} \sim Mult_P\left(\left\{\omega_t^{(\phi)}, w_t^{(\phi)}, \gamma_t^{(p)}\right\}_{\phi=1}^{P}\right)$$

where the importance weights are given by

$$\omega_t^{(p)} \propto \left(1 + \frac{1}{\nu} \frac{\left(r_{\Phi,t} - \left(1 + \gamma_t^{(p)}\right) w_t^{(p)\prime} r_{A,t}\right)^2}{\sigma_\epsilon^2}\right)^{-\frac{\nu+1}{2}}$$

Figure 3.2: *L-DPM Estimation Algorithm*

estimation. Then, these estimates can be directly compared to the true weight values in order to motivate estimation accuracy.

We proceed by constructing an aggregated dataset of the actively managed diversified equity (AMDE) mutual funds. For this exercise, we focus attention on portfolios consisting of almost all equity securities since exact snapshots of these equity holdings are readily and accurately available. As similarly performed in Kacperczyk, Sialm, and Zheng (2005), we merge the CRSP Survivorship Bias Free Mutual

Fund Database with the Thomson-Reuters Mutual Fund Holdings Database, formerly known as the CDA Investment Technologies/Spectrum holding database, and the CRSP individual stock level data. The CRSP Mutual Fund Database includes information on various fund characteristics including total net assets, fund returns, and investment objectives/classifications. Recently, it has been updated to include stock level holdings, however the Thomson-Reuters Mutual Fund Holdings Database includes a much longer history of these holdings, and therefore provides a more extensive time series of data for analysis. We follow Wermers (2000) and use the Wharton Research Database Services (WRDS) Mutual Funds Links (MFLinks) Database to appropriately join the the CRSP and Thomson-Reuters mutual fund databases. Our set of actively managed diversified equity funds is then constructed by excluding balanced, bond, index, international, and sector funds. As well, we exclude all mutual fund observations where total net assets were less than $1 million in that quarter. Our final sample spans the period from December 1987 to June 2010.

Using the resulting mutual fund dataset, we construct an aggregate return index on the net asset value weighted set of actively managed diversified equity mutual funds from the fund returns and net asset values in the CRSP Mutual Fund Database. Then, with the merged holdings data for this set of mutual funds, we can obtain true relative portfolio weights on any mutually exclusive division of the underlying compositional assets. Since the estimation space on individual stocks can be on the order of 10^4 dimensional, we reduce the dimension of equities into industry groupings. For a straightforward exhibition of our methodology, we adopt the Ken French 5 industry groupings detailed in Table 3.1. The 'Consumer' group includes consumer durables, non-durables, wholesale, retail, and some services (laundries, repair shops). The 'Manufacturing' group includes manufacturing, energy, and utilities. The 'High Technology' group includes business equipment, telephone, and television transmission. The ‘Health’ group includes healthcare, medical equipment, and drugs. Finally, the ‘Other’ group includes mines, construction, building management, transportation, hotels, business services, entertainment, and finance.

The true industry weights for the set of mutual funds is constructed by matching individual stock SIC codes with the associated industry grouping. Then, total capital invested in each industry is summed over all mutual funds and normalized over total capital in all industries to obtain relative portfolio weight values for each time period. This establishes our *true* industry relative portfolio weights for the AMDE mutual fund set. Consistent with this, the Ken French 5 industry portfolio returns are used for the explanatory compositional asset returns. These returns are constructed by the value weighted average return on the complete set of firms in the given industry, rebalanced yearly at the end of June. That is, it represents the return on holding the portfolio of all firms in an industry in proportion to their market capitalization. If the mutual funds hold the set of firms in a different proportion, or if the funds perform large amounts of active trading between reporting periods, these industry portfolio returns may become less suitable. Nevertheless, although these value-weighted industry returns may not constitute the exact composition of the mutual funds’ investment choices and resulting returns on each industry, they do represent a decent proxy as the size of the chosen mutual fund set grows large.

3.4.2 Hedge Fund Data

In the hedge fund portfolio estimation problem, we consider the monthly return data for the Hedge Fund Research Fund Weighted Composite Index (HFRIFWC) from January 1995 to October 2012. This index is formed from monthly self-reported returns complied from a source of over 2,200 hedge funds with over $50 million USD under management and a track record of greater than 12 months. These returns are

Industry	SIC Codes
Consumer	0100-0999
	2000-2399
	2700-2749
	2770-2799
	3100-3199
	3940-3989
	2500-2519
	2590-2599
	3630-3659
	3710-3711
	3714-3714
	3716-3716
	3750-3751
	3792-3792
	3900-3939
	3990-3999
	5000-5999
	7200-7299
	7600-7699
Manufacturing	2520-2589
	2600-2699
	2750-2769
	2800-2829
	2840-2899
	3000-3099
	3200-3569
	3580-3621
Manufacturing (cont.)	3623-3629
	3700-3709
	3712-3713
	3715-3715
	3717-3749
	3752-3791
	3793-3799
	3860-3899
	1200-1399
	2900-2999
	4900-4949
High Tech	3570-3579
	3622-3622
	3660-3692
	3694-3699
	3810-3839
	7370-7379
	7391-7391
	8730-8734
	4800-4899
Health	2830-2839
	3693-3693
	3840-3859
	8000-8099
Other	All Remaining

Table 3.1: *Ken French 5 Industry Groupings*

reported net of individual fund managers' fees. As extensively detailed in historical literature, we note that there can be some biases present in these self-reported aggregated hedge fund indices including illiquidity-induced serial correlation in fund returns (Getmansky, Lo, and Makarov, 2004), backfill and incubation bias (Fung and Hsieh, 2004), and self-selection bias (Fung and Hsieh, 2000, 2009). Nevertheless, in the absence of better information, this data remains a reasonable and commonly used proxy for the performance of the aggregate hedge fund industry portfolio.

For the portfolio asset classes, we use a similar set to those used in Fung and Hsieh (1997). Table 3.2 enumerates the asset class list and color key. Equities are divided into US, non-US (Europe, Australasia, Far East), and emerging markets. The prevalent Morgan Stanley Capital International (MSCI) indices are used are used for each of these respective locales. Commodity returns are represented by spot returns on Gold and WTI Cushing Crude Oil. Returns on holding long US Dollar futures against the Euro, Japanese Yen, British Pound, Canadian Dollar, Swedish Krona and Swiss Franc are represented by the Deutsche Bank US Dollar Long Futures Index. In contrast, we break fixed income securities into municipal, corporate high yield, and mortgage backed securities (MBS), since Fung and Hsieh (1997) find that municipal and high yield bond funds have low correlation with their asset class set. We use the respective Barclays Capital Indices for each of these fixed income classifications. Consistent with their findings, the inclusion of these asset classes in our estimation procedure significantly increases the amount of variation in returns explainable in the index for the aggregate hedge fund industry. This supports the intuition that both municipal bonds and corporate high yield bonds are important components of the hedge fund industry portfolio. We also include an index of MBS returns since we find that their inclusion significantly increases explainable index variation. As with any estimation technique, we emphasize the importance of appropriate explanatory variables, and in this case, an appropriate set of portfolio assets due to the potential for omitted variable bias arising from correlation across asset returns.

Additionally, we would like to recognize the popular hypothesis of hedge funds taking large short positions on various assets. Notably, the structure of this Dirichlet Portfolio Model family does not allow for negative weight estimation since values are restricted to the simplex. Although it may be easy for an individual hedge fund to

Color	Asset Class Name
	MSCI US Equity Index
	MSCI EAFE Index (Europe, Australasia, Far East)
	MSCI Emerging Markets Index
	Gold Spot
	WTI Cushing Crude Oil Spot
	Deutsche Bank US Dollar Long Futures Index
	Barclays Corporate High Yield Bond Index
	Barclays Mortgage Backed Securities Index
	Barclays Municipal Bond Index

Table 3.2: *Asset Class Color Code*

take conceivably large short positions on a particular asset, historical data suggests that it is unlikely that the aggregate of all hedge funds holds net short positions on any asset class. England's Financial Conduct Authority, formerly the Financial Services Authority, conducts periodic holdings surveys on a subset of approximately 50 hedge fund managers across over 100 funds. These quarterly reports confirm that although some funds bolster large short positions, the net exposures inside all these asset classes are almost always net-long. In the very rare cases where exposures are net-short, they are almost indistinguishable from zero. Therefore, the assumed net-long estimation structure is largely appropriate.

3.5 Mutual Fund Results

In this section, we exhibit the estimation results on our constructed index of actively managed diversified equity mutual fund performance. Figure 3.3 shows the monthly series of estimated industry relative portfolio weights for each of the five industry portfolios. Overlaid on these plots are the quarterly constructed series of "true" industry weights. We note that while these portfolio weights are not *precisely* weights on the value weighted industry portfolios, they closely proxy for these as the net asset value size of the overall index portfolio, and therefore the aggregate holdings size in each industry, grows large.

Examining the estimation results, we find that the L-DPM procedure strongly captures the dynamics of the portfolio weight allocation process. Particularly, the fitted weight series tends to be most accurate for the Technology, Health, and Other groupings due to their lower correlation with each other as well as the remaining industry returns. For our hedge fund example in the next section, this suggests potential for even more precise estimation results since contemporary correlations across various asset classes are significantly less than those across equity industries. We also point out that since the constructed mutual fund index surely holds its particular selection of firms within an industry in slightly different proportions than the corresponding value weighted industry portfolio, minor estimated weight differences are expected due to marginally different industry correlation structures. Nevertheless, since we are especially interested in *changes* in asset allocation, we observe that even in the assets where the estimated weights do not track precisely with the true weights, their first-differences do track very closely.

It is also important to mention that although these industry weights appear constant over some periods of time, there are a number of periods where weights change rather significantly. To test the null hypothesis that these portfolio weights are constant over the sample period, and that simple OLS would be suitable, we employ the Generalized Fluctuation test of Kuan and Hornik (1995) to test for time varying regression coefficients. The result of this test gives a p-value of 1.749e-9, indicating a strong rejection of the constant weight hypothesis in favor of a model with time-varying weights. The test indicates that the majority of the rejection strength arises from the High Tech, Manufacturing, and Other industries.

Figure 3.4 illustrates the estimated means of these weights, stacked in an area plot

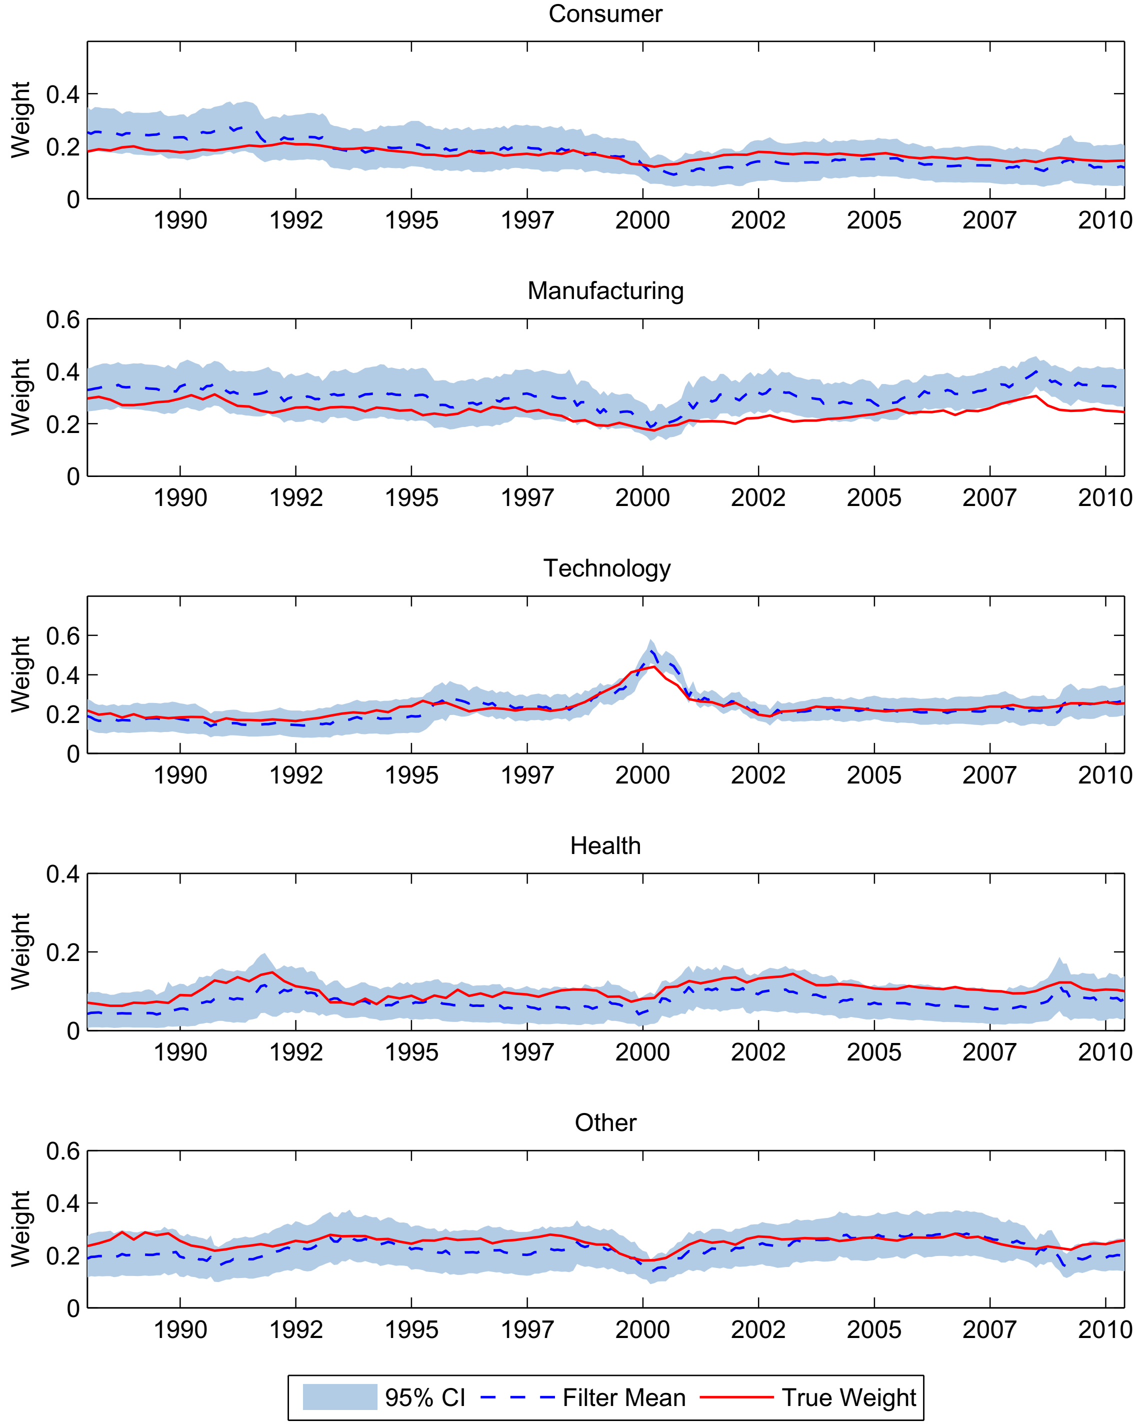

FIGURE 3.3: *Estimated Weight Accuracy Comparison*

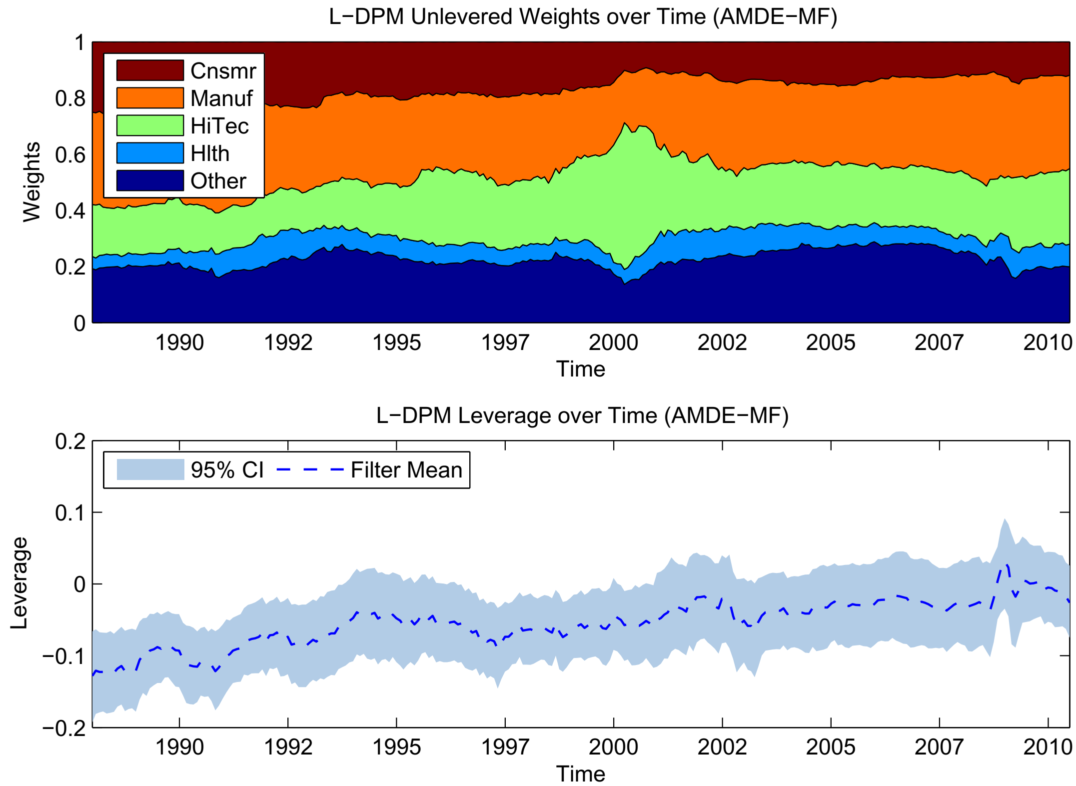

FIGURE 3.4: *Aggregated Estimated Weights & Leverage*

to give perspective on how they evolve dynamically with respect to each other. Along with this, the joint estimation results for the leverage scaling value is plotted below. As expected, this value is close to, and sometimes below zero indicating a net absence of leverage across the set of actively managed diversified equity mutual funds. This net de-leveraged effect can be indicative of cash holdings, underperformance due to management fees, and even underperformance in firm selection within each relative industry set. Notably, there is a discernible upward trend in these leverage values, becoming not statistically different from zero in the most recent period. This suggests a diminishing effect of these aforementioned causes.

Figure 3.5 exhibits the cumulative return series for the individual industry port-

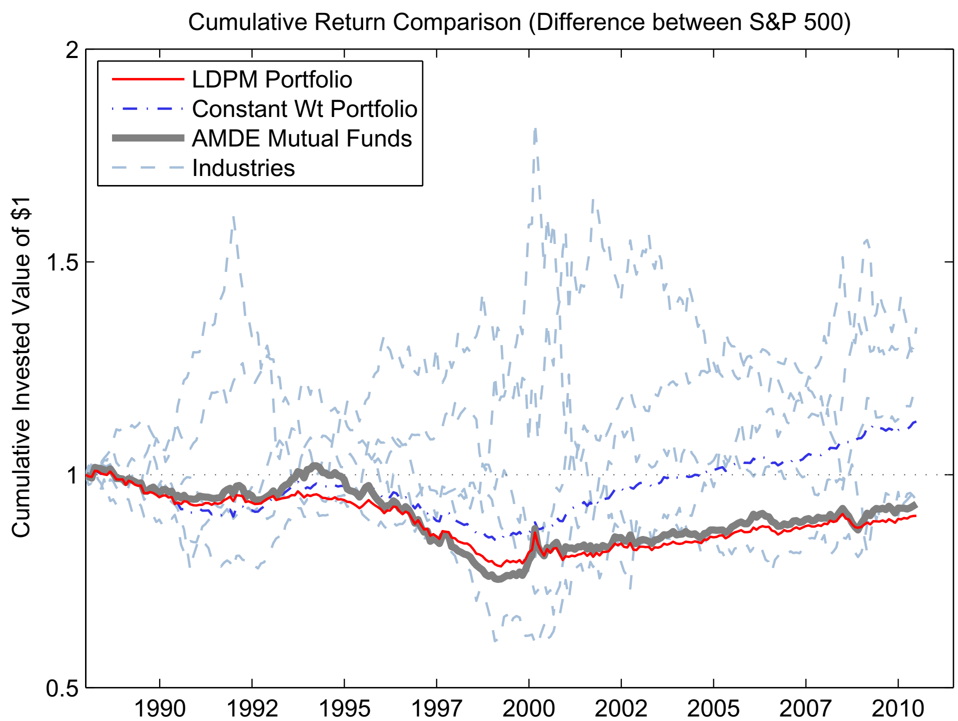

FIGURE 3.5: *Cumulative Return Comparison*

folios, the constructed actively managed diversified equity mutual fund index, and the replicating portfolio constructed by the L-DPM's estimation on the 5 industry portfolios, less the return on the S&P 500. We subtract the return on the S&P 500 to remove broad market variation due to equity mutual funds' high correlation with this market proxy, and therefore focus on the differences between our series of interest. For comparison, we also include a return series on a time-invariant, constant weight portfolio estimated on the entire sample via OLS. We see that the L-DPM procedure does a very accurate job in explaining the variation in the index of mutual fund returns by the combination of component industry groups and the estimated dynamic asset allocation process. Table 3.3 shows monthly sample statistics associated with these return series. While individual industry correlations with the constructed mutual fund index range from 65.3% to 89.4%, the L-DPM portfolio has a 99.4% correlation with the index. As well, the dynamic asset allocation process of the L-DPM is able to explain 98.8% of the variation in the index returns from the individual industry portfolios' explanatory power of 42.6% to 79.9%. This is consistent with the theoretical unexplainable variation implied by the constructed true industry weights of around 2%. For comparison, the results using constant OLS weights gives smaller correlation and R^2 values of 98.6% and 97.3%, respectively. Finally, we point out that the mean returns over the sample period for all of the industry portfolios are greater than that of the constructed mutual fund index, thereby motivating the previously mentioned de-leveraged effect.

Monthly Mutual Fund Summary Statistics

Monthly Rets	AMDE	L-DPM	Cnsmr	Manuf	HiTec	Hlth	Other
Mean	0.00786	0.00772	0.00879	0.00922	0.00865	0.00988	0.00809
Std Dev	0.04197	0.04186	0.04124	0.04096	0.06466	0.04578	0.05217
RMSE	0	0.00456	0.02301	0.02276	0.03293	0.03670	0.02616
Mean Abs Err	0	0.00319	0.01667	0.01593	0.02257	0.02776	0.01865
Correlation	1	0.99408	0.84692	0.84960	0.89438	0.65301	0.86682
R^2	1	0.98819	0.71727	0.72182	0.79992	0.42642	0.75138

Table 3.3: *Explanatory Summary Statistics (Monthly Returns)*

3.6 Hedge Fund Results

In this section, we detail the estimation results for the decomposition of the aggregate hedge fund industry portfolio. Similarly to the mutual fund example, we use the Leveraged Dirichlet Portfolio Model to estimate relative portfolio weights and leverage multiplier values on the aggregate hedge fund return data. We then use these values to compute implied portfolio returns to evaluate explanatory power of the dynamic portfolio weight allocation process and leverage values. As well, since true values of the hedge fund portfolio weights are never available, we construct one-step-ahead predicted weights and leverage values to produce a time series of "forecasted" returns. These are compared with the observed index returns to demonstrate the out-of-sample forecasting accuracy of the estimation procedure.

Figure 3.6 shows the estimation results for both the relative portfolio weight process and the leverage multiplier values. The corresponding color key is previously exhibited in Table 3.2. A clear cyclical pattern is discernible in the estimated weights, with transitions from equities in periods of economic expansion to fixed income securities during economic contraction. Revisiting the previously mentioned Financial Conduct Authority's sample of 50 surveyed hedge fund managers, we can check if these estimated weights are generally consistent with their reported weights, even for such a small sample of the hedge fund industry. We note that their chosen asset classes represent different mutually exclusive divisions than ours. For example, they divide equities into 'listed' and 'un-listed', whereas we divide them into the more descriptive US, EAFE (Europe, Australasia, and Far East), and emerging markets categories. Therefore, although we cannot appropriately compare results on a very focused level, we still can compare across broad asset classes. That is, in March 2012, they show relative net equity, fixed income, and other exposures of around 35%, 57%, and 8%, where our methodology produces estimates of 36.6%, 52.9% and, 10.5%, respectively. Even though differences in these values can surely manifest due to differences between the groups of hedge funds included in the data sets, these values are remarkably similar.

The figure also shows the estimated values for the leverage multiplier γ_t. Similar to Korsos (2013a), we observe a rather constant, but slowly decreasing leverage value. Notably, the estimated value is never statistically different from zero. This suggests that although some funds may employ large amounts of gross leverage to obtain their individual portfolio returns, the net value of borrowing among this entire sample of the hedge fund industry is in fact, negligible. We will examine and discuss detailed reasoning for these results in section 6.

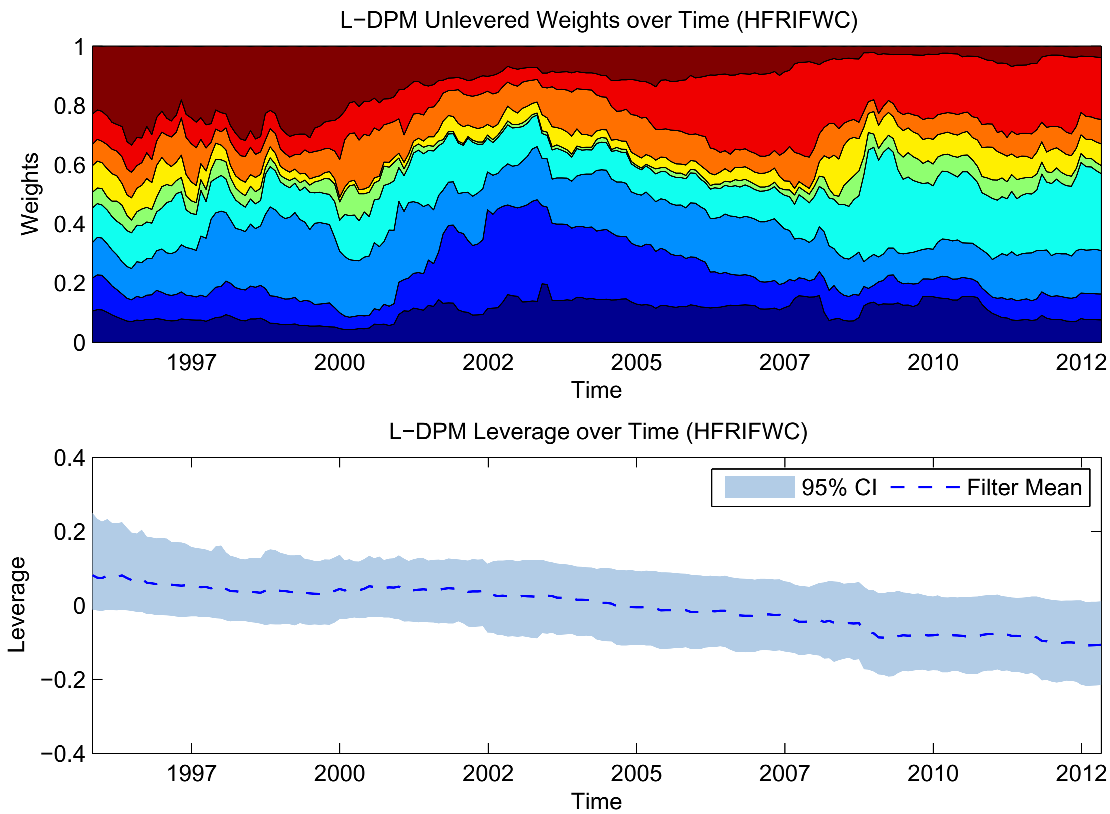

FIGURE 3.6: *Estimated Relative Weights & Leverage Values*

Figure 3.7 compares the portfolio returns constructed from using both the same-period estimated, as well as one-step-ahead forecasted weights and leverage values. We observe that the portfolio returns produced by using the same-period estimated portfolio weights and leverage are exceptionally close to the true hedge fund returns, indicating strong in-sample explanatory power of the L-DPM. Even more compelling is the series of portfolio returns produced with the one-step-ahead forecasted portfolio weights and leverage. The forecasted returns are also remarkably close to the actual observed returns in the next period, thereby strongly supporting the procedure's forecasting ability and thus estimation accuracy. The most noticeable differences in the forecasting accuracy occur around periods of high volatility in the hedge fund returns. We identify 3 major reasons for this effect. First, since funds are likely to attempt to strategically adjust their asset allocation, it is difficult to capture quick, intra-month changes in portfolio weights due to the monthly data granularity. Second, since the correlation of asset returns commonly increases during volatile periods (Jacquier and Marcus, 2001; Engle, 2002), it becomes more difficult to identify and decompose the sources of portfolio risk. Third, since the realization of high volatility decreases value to investors, funds are more likely to employ smoothing of illiquid asset returns in a highly volatile environment (Kadlec and Patterson, 1999; Fisher, Gatzlaff, Geltner, and Haurin, 2003). Despite these potential issues, during periods of high asset price volatility, although the *absolute* amount of error increases, the absolute amount of correctly captured variation increases as well. Hence, as the underlying asset returns become more volatile, the *relative* amount of error does not significantly increase as a percentage of total volatility, thereby implying a rather constant R^2 value over the entire sample period.

Table 3.4 details summary statistics for the monthly returns of the index, the returns implied by the weights and leverage of the Leveraged Dirichlet Portfolio Model estimation, and the returns on the individual compositional asset classes. As in the previous return comparison plots, the 'explanatory' return values are constructed with the same-period sequential estimates, thereby indicating explanatory power of the model given the compositional asset classes. While the individual asset class correlations with the hedge fund index range from -22.3% to 83.5%, the L-DPM portfolio has a 92.3% correlation with the index. As well, the dynamic asset allocation process of the L-DPM is able to explain 85.2% of the variation in the index returns from the individual asset classes' explanatory power of 0.3% to 69.7%. As well, it

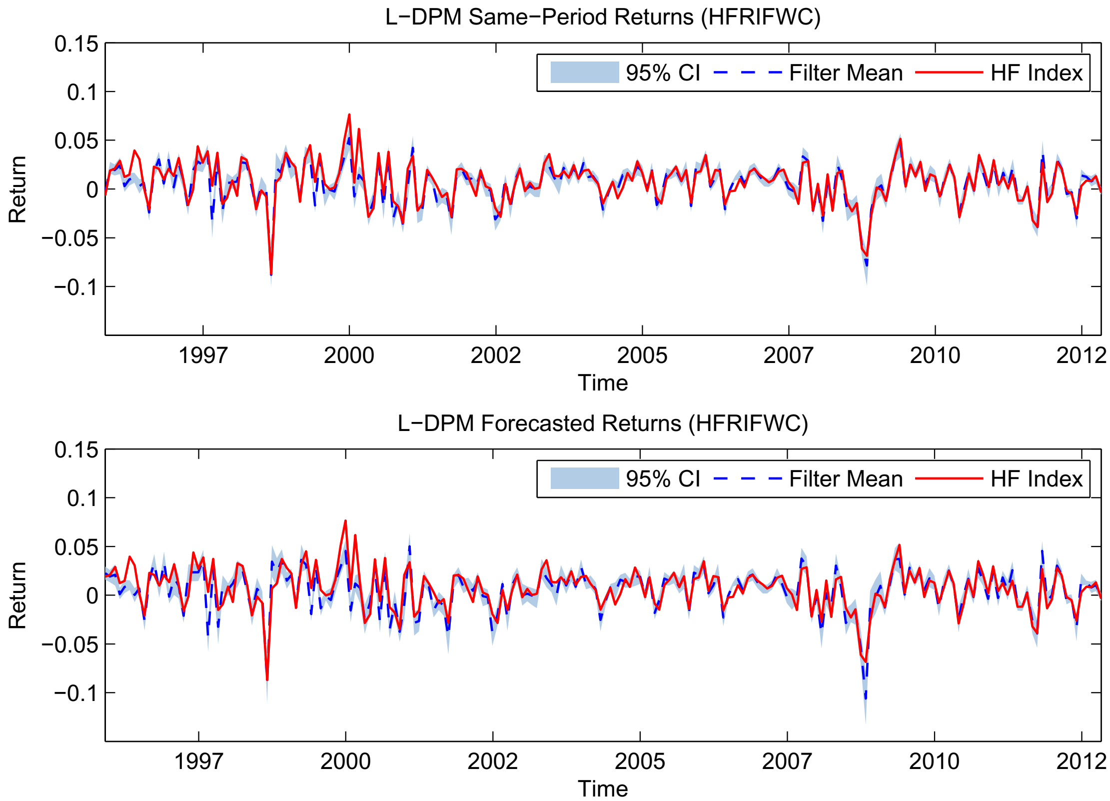

FIGURE 3.7: *Same-Period & Forecasted Returns*

is important to identify that these sequential estimates do not even include future period data in the current period estimates. Therefore, they do not use the entire sample period of data at each time period, as is commonly done in the many historical OLS based approaches. Thus, these estimates may be improved by implementing the fixed-interval particle smoothing technique of Carvalho, Johannes, Lopes, and Polson (2010) in order to refine the portfolio weight and leverage estimates at all periods. Nevertheless, due to the sequential nature of the estimation procedure, the L-DPM does have a stark advantage of being able to produce 'forecasted' portfolio return values constructed with the prior weights and leverage implied by the dynamic transition model. Significantly, these forecasted estimates exhibit predictive power of 75.7% of the one-month-ahead hedge fund index return variation.

Although the mean values of the returns implied by the L-DPM estimated weights and leverage are below those of the original index, this effect arises from the aforementioned biases and return smoothing in the self-reported data. As previously suggested, applying a desmoothing model to the original return data, shifts the weights into higher volatility and higher returning assets. This decreases, but does not completely eliminate the gap. The remaining difference can be interpreted as the effect of reporting biases in the data. The size of this underperformance gap is consistent with Malkiel and Saha (2005) and Jurek and Stafford (2013) who report the annualized effects of these biases from as low as 3% to potentially over 7% for various sources of hedge fund data. For comparison to our results, that would imply a monthly upward bias of between 0.25% to 0.583% in the mean return.

Figure 3.8 illustrates a cumulative return comparison, similar to that in the mutual fund example. Plotted are the individual compositional asset class returns, the aggregate hedge fund industry index, and the replicating portfolio constructed by the L-DPM's estimation on the 9 asset classes. We see that the L-DPM procedure is able to construct a cumulative return series which is shaped very similarly to that of the hedge fund index, especially when considering the differing variability in the individual compositional assets. Similar to the numerical results above, the underperformance size, relative to the index, is consistent with the data bias sizes suggested in the aforementioned literature.

Strikingly, this graphical perspective identifies that most of the "underperformance" does not occur evenly throughout the entire sample period, but instead mostly around late 1999 to early 2000. In fact, aligning the true index and the L-DPM replicating portfolio as of 2001 produces an almost identical fit. As previously mentioned, data issues surely contribute to this difference especially because

Monthly Hedge Fund Summary Statistics

Monthly Returns	**HFRIFWC Index**	**Explanatory L-DPM**	**Forecasted L-DPM**	**US Equity**
Mean	0.00707	0.00545	0.00480	0.00652
Standard Deviation	0.02122	0.02063	0.02190	0.04643
RMSE	0	0.00837	0.01121	0.03337
Mean Absolute Error	0	0.00552	0.00784	0.02484
Correlation	1	0.92302	0.86998	0.75456
R^2	1	0.85196	0.75687	0.56935
Monthly Returns	**EAFE Equity**	**Emerging Mkt Equity**	**Gold Spot**	**Crude Oil Spot**
Mean	0.00459	0.00767	0.00846	0.01233
Standard Deviation	0.05016	0.07193	0.04768	0.09517
RMSE	0.03612	0.05529	0.04808	0.08921
Mean Absolute Error	0.02671	0.04214	0.03789	0.07079
Correlation	0.77959	0.83514	0.19936	0.38071
R^2	0.60777	0.69746	0.03975	0.14494
Monthly Returns	**Long US Dollar**	**High Yield Corporates**	**Mortgage Backed Sec**	**Municipal Bonds**
Mean	0.00203	0.00657	0.00504	0.00476
Standard Deviation	0.02463	0.02761	0.00763	0.01203
RMSE	0.03619	0.02190	0.02325	0.02391
Mean Absolute Error	0.02731	0.01598	0.01788	0.01863
Correlation	-0.22282	0.62389	-0.09114	0.05295
R^2	0.04965	0.38924	0.00831	0.00280

Table 3.4: *Explanatory & Forecasted Summary Statistics (Monthly Returns)*

the early 2000's saw a dramatic increase in the number of hedge funds, and since the index is comprised of funds which have existed for at least 12 months, backfill and incubation biases are an especially significant consideration during this time period.

3.7 Net Leverage & Asset Class Exposures

Examining the estimated hedge fund results, it is initially striking to see net leverage values which are in the range from 0% to 10%. From widespread reports in the popular media, it is not uncommon to hear about hedge funds taking leveraged

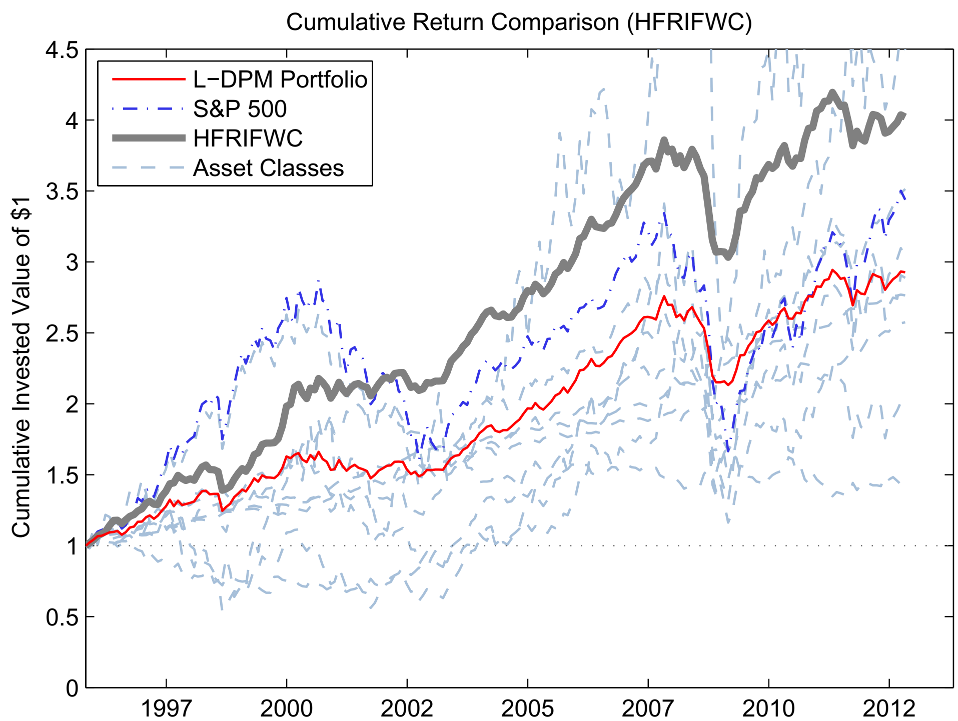

FIGURE 3.8: *Cumulative Return Comparison*

positions of over 10 times their contributed capital. For example, in an April 2012 article in the *Financial Times*, Michael Hintze, chief executive of CQS, a $9 billion London based hedge fund said, "hedge funds are presently leveraged one to three times; if they're mad, five times; if they're insane, 10 times." Given this media imposed 'prior', it seems that we should expect estimated leverage multiplier values in our model of around 1,000%. So this certainly provokes the question: Where is all the leverage in the hedge fund return data?

First, let us define the idea of leverage. That is, leverage is the means by which an investor captures an exposure level to a particular asset which is greater than what could normally be captured by only using contributed capital and directly investing in the notional amount of that asset.

Now, in order to answer the question above, we first detail the various sources of hedge fund leverage. First, it must be noted that not all forms of leverage directly involve borrowing to fund asset purchases. This usually is referred to as 'synthetic' or 'embedded' leverage. For example, futures contracts provide profit or loss on a notional amount of an asset for only a posted margin of around 5 to 15% of that asset value. Therefore, performance indices based on these futures contracts, like our commodities and US Dollar futures indices above, already have leverage built into their returns. Options contracts are a similar concept where investor can gain exposure to underlying asset price movements for a fraction of the notional value of those assets. Structured products like collateralized debt obligations (CDOs) also exhibit leverage-like qualities for junior payment tranches. Outside of exchange traded or over-the-counter (OTC) securities, funds can structure swap contracts with investment banks to exchange return differences between two reference assets. Since the hedge funds only must post a fractional margin amount and do not directly purchase the full notional value of these assets, these swaps exhibit leveraging qualities as well. Accompanying these products, funds may choose to directly borrow in order to fund other desired exposures. Most commonly, this is achieved via repurchase agreements (REPOs), but can also be done through prime brokerage borrowing facilities, and to a lesser extent unsecured borrowing or even traditional term financing, like the issuance of public debt.

Investment managers employ these different types of leverage for a variety of reasons. The most familiar use is to obtain general return enhancement when a manager has strong conviction to increase directional exposure. On the other hand, leverage can be used as part of a risk reduction strategy if a manager places short positions against their long positions to potentially decrease exposure to a cross-sectional risk factor. In this case, gross exposures will surely exceed the amount of contributed capital in the strategy, but the net exposure will be below this value. As well, leverage is commonly used for the magnification of low risk returns in spread strategies where directional market speculation is not necessarily made, but instead where arbitrage may exist or nearly complete risk hedging is postulated for a particular strategy. Finally, leverage can occur when a manager is concerned with liquidity and transaction costs. For example, in the cases of commodities and corporate bonds, futures and other derivatives on the reference assets are commonly more liquid or carry lower transaction costs than the underlying assets themselves. Therefore, fund managers will commonly invest in these products as a more efficient means to gain exposure to the referenced assets' price movements.

Returning to our originally posed question, let us now consider a hedge fund who believes that within a particular industry, a certain stock is slightly over-priced relative to another stock. Wanting to place a bet that these two prices will converge, the fund does not want to be exposed to the common industry risk present in both of these assets. By placing a long bet on the relatively under-priced asset, and a corresponding short bet on the relatively over-priced asset they can achieve this goal. Although their total exposure of this position may be quite small due to the netting of the industry exposure, the gross exposure appears quite large since it is a total of both long and short positions. This gross exposure, divided by contributed capital is the commonly quoted leverage multiplier value seen in industry reports and the popular media. If we instead consider the case where these funds invest in the corresponding options contracts rather than the reference equities, this difference between the net and gross exposures can quickly grow rather large and obfuscate our view of aggregate exposure levels.

Furthermore, individual funds, who naturally will have varying directions of net exposures, will additionally net out to an even smaller magnitude aggregate net exposure. Concurrently, each of their gross exposures will invariably sum together to create an even larger aggregate gross exposure. This effect creates a compounding amplification of aggregate gross exposure levels while aggregate net exposures can remain relatively small. Although not explicit as to this effect, evidence of this is found in England's Financial Conduct Authority's (FCA) periodic reports on hedge fund risks. For example, in their subset of surveyed funds they find around \$200 billion in short exposures and \$275 billion in long exposures to listed equities, thereby giving \$475 billion in gross exposures and only \$75 billion in net exposures. Nevertheless, the FCA's July 2011 report does identify that net exposures are "generally positive but low across most asset classes [and that] this low 'net long' exposure is a common characteristic of hedge funds and differentiates them from other asset managers."

Interwoven in this concept of net versus gross asset exposures is net versus gross leverage. In the listed equities example above, financing for much of the long positions can be offset via proceeds from the short positions. To keep this example simple, assuming negligibly small borrowing costs and margin requirements, \$200 billion of the long positions are financed with \$200 billion of proceeds from the shorts. This leaves \$75 billion remaining to finance in long equity positions. If this entire amount is financed via contributed capital, this yields a gross leverage multiple of $(200 + 200 + 75)/75 = 6.33$, whereas net leverage is simply $(-200 + 200 + 75)/75 = 1$, that is non-existent. Financing even \$10 billion of these equity positions via borrowing facilities still yields similar multiplier results of $(200+200+75)/65 = 7.31$ for gross leverage and only $(-200 + 200 + 75)/65 = 1.15$ for net leverage. That is, with \$10 billion of borrowing, gross leverage appears as 631% above contributed capital, when net leverage truly is only 15%. As well, these differences in leverage are even more pronounced in fixed income securities since very large leveraged positions can be achieved via futures, swaps, and structured products contracts with no need for direct borrowing to fund these positions.

The 2008 financial crisis, the near-failure of Bear Stearns, and the bankruptcy of Lehman Brothers have surely highlighted the importance to hedge funds of the seemingly more administrative parts of their operations according to the European Central Bank's (ECB) periodic *Financial Stability Review*. This includes simple functions like monitoring both cash and securities balances with their prime brokerages and even their counterparty exposures. As well, banks have become more cautious with their transactions with these prime brokerage clients. This precaution has manifested into higher margin requirements and more conservative collateral valuation. As well, investors in hedge funds have demanded more liquidity and higher cash balances for redemption waves in order to avoid fire-sales on existing portfolio assets. All of these have a strong diminishing effect on the amount of leverage used by hedge funds.

According to the ECB, this deleveraging shift is "intensified by mounting and expected redemption requests from investors." They suggest that this has caused many hedge funds to reallocate investments to less risky assets such as cash and equivalents. As well, they report that the number of hedge funds reporting leverage multiples of less than one reached record high levels in September and October 2008. Our estimated results confirm this suspicion. Furthermore, the aforementioned FCA reports find that ratios of unencumbered cash to total borrowings from 2009 to 2012 have averaged around 75%, suggesting that even when funds use borrowing in the construction of their portfolio, most of it remains in cash as a protection against redemption waves and margin calls.

To answer the question of where the leverage is in the data, the simple remark is: It's there, but it's hiding. Due to the netting of opposing gross positions, both intra- and inter-fund, the aggregate net leverage exposure remains quite small. Since the net exposures and leverage of the industry of hedge funds is very small compared to the gross exposures and leverage, this suggests that these funds are largely just taking sizable bets against each other.

3.8 Implications for Herding Behavior

In the aftermath of financial crises, the idea of 'herding' behavior, or positive-feedback trading, commonly arises in examinatory discussions of the financial markets. Herding is a pejorative term which describes the behavior of groups of investors who sell when prices are falling and buy when prices are rising. Morris and Shin (1999), Persaud (2000), and Shiller (1990) express concern and detail why this herding behavior can result in increased volatility as well as potentially destabilized markets. In the 2008 financial crisis, widespread criticism was brought upon hedge funds due to the belief that their market participation exacerbated asset volatility. In one of the seminal works on herding, Kodres and Pritsker (1996) study a sample of hedge funds and their trading behavior in futures contracts from 1992 to 1994 and find that funds, in fact, exhibit not positive, but negative-feedback trading behavior. That is, they tend to buy when prices are decreasing and sell when prices are increasing. Therefore, since they provide a ready counterparty for trading, they not only increase liquidity, but also decrease volatility by reducing price pressures in the presence of directional market forces.

In more recent events, the idea of potential herding behavior by hedge funds has once again become widespread. As well, since complete and detailed time series of hedge fund holdings are not publicly available, this has left the fund industry open to mass criticism since evidence is not readily available to either support or refute the claims of the popular media. Conveniently, the portfolio decomposition of the Dirichlet Portfolio Model family allows for the construction of active portfolio weight changes. Analyzing these asset class weight changes, as they relate to previous period's returns allows us to assess not only if hedge funds still exhibit negative-feedback trading behavior, but also if that behavior is limited to particular asset classes.

As Korsos (2013a) previously identified, there exists a contemporary relationship between active changes in hedge fund portfolio asset allocation and corresponding asset class returns, attributable to classic demand variation on market prices. These effects were identified by examining the relationship of active portfolio weight changes and same-period asset class returns. Contrastingly, in order to evaluate potential asset class feedback behavior, we are interested in assessing effects of historical asset class returns on current period active changes in portfolio allocation. In order to estimate these potential relationships, we similarly identify that portfolio weight changes arise from two causes: capital appreciation/depreciation and time varying asset allocation. The proportion of portfolio weight changes caused by holding period capital appreciation/depreciation is not an investment decision directly controllable by a portfolio manager. Therefore we concern ourselves only with portfolio weight changes caused by active asset allocation trading decisions. In terms of our notation, the portfolio weight change over time t caused by these active trading decisions is given by $w_{t|t} - w_{t|t-1}$, where $w_{t|t}$ is the estimated portfolio weight at t and $w_{t|t-1}$ is the portfolio weight at t given a strict buy-and-hold transition from the previous period's estimated weight $w_{t-1|t-1}$. This difference is the change in investment over period t due to an active decision to adjust portfolio allocation. In terms of the dynamic model, this difference is also thought of as the innovations on the weight transition process. In the context of the L-DPM portfolio estimation with leverage multiplier, this portfolio weight change is similarly given by $\left(1+\gamma_{t|t}\right) w_{t|t} - \left(1+\gamma_{t|t-1}\right) w_{t|t-1}$. We can now estimate the effect of historical returns on current active asset allocation decisions on component asset i via:

$$\left(1+\gamma_{t+1|t+1}\right) w_{t+1,i|t+1} - \left(1+\gamma_{t+1|t}\right) w_{t+1,i|t} = \beta_{0,i} + \beta_{1,i} r_{t-1,i} + \xi_{t,i}$$

where

$$w_{t+1,i|t} = \frac{w_{t,i}(1+r_{A,t,i})}{\sum_{i=1}^{N} w_{t,i}(1+r_{A,t,i})} \quad \text{and} \quad \gamma_{t+1|t} = \frac{\gamma_t}{1+r_{\Phi,t}\left(1+\gamma_t\right)}$$

The constant coefficient β_0 represents the average unconditional active change in portfolio weights each month. The coefficient β_1 represents the average effect on the portfolio-size-normalized active change in asset class investment allocation over month t given asset class return over month $t-1$. That is, β_1 is the average active portfolio weight change effect explainable by variation in the previous month's asset class return.

Table 3.5 exhibits the OLS estimation details for each of the asset classes using the results from the L-DPM. Not surprisingly, the constant coefficients β_0 are all indistinguishable from zero since the portfolio weight series are all stationary. We focus attention on the β_1 coefficients, which indicate the effects of last month's lagged asset class returns on the current month's active changes in portfolio compositions. Although not all statistically different from zero, 7 out of 9 of the asset classes have negative coefficients, suggesting that when prices drop, hedge funds increase exposure to those assets in the following month. This effect is strongest in oil and municipal bonds. On the other hand, the two asset classes with positive coefficients, mortgage backed securities and emerging market equity, are both statistically significant. Mortgage backed securities is most significant, suggesting that when MBS prices drop, hedge funds have sold these assets. That is, although hedge funds exhibit negative-feedback trading in most asset classes, and therefore contribute to lower asset price volatility in those respective assets, they do exhibit positive-feedback trading behavior in MBS products and EM equity, thereby potentially contributing to increased MBS and EM equity price volatility through herding.

Feedback Trading Effects

	US Equity	EAFE Equity	EM Equity
β_0	−0.00105	0.00055	−0.00016
	(0.00095)	(0.00076)	(0.00073)
β_1	−0.02830	−0.01442	0.05925*
	(0.01831)	(0.02920)	(0.03388)
R^2	0.00799	0.00131	0.00755

$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$

Feedback Trading Effects

	Gold	Oil	US Dollar
β_0	−0.00041	−0.00054	0.00091
	(0.00059)	(0.00049)	(0.00081)
β_1	−0.02089	−0.16486*	−0.00701
	(0.03455)	(0.08941)	(0.00860)
R^2	0.00125	0.01939	0.00161

$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$

Feedback Trading Effects

	High Yield	MBS	Municipal
β_0	0.00014	−0.00104	−0.00005
	(0.00109)	(0.00082)	(0.00072)
β_1	−0.02471	0.00768**	−0.01375*
	(0.01655)	(0.00286)	(0.00847)
R^2	0.01709	0.01204	0.01351

$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$
HAC standard errors are reported due to potential for autocorrelation & heteroskedasticity in the errors terms

Table 3.5: *Active Feedback Trading Effects*

However, we note that although there is evidence that when prices decrease in mortgage backed securities, hedge funds subsequently decrease relative holdings in these assets, there is another possible explanation outside of the postulated herding behavior which contributes to higher asset price volatility. In particular, mortgage backed securities followed a consistent downward trend during the sub-prime mortgage crisis of 2008. As prices declined, hedge funds decreased exposures to these assets. Assuming that initially, prices under-reacted to declining housing prices and increasing mortgage default rates, hedge fund market participation instead assisted in the market's price discovery process. Hence, their prudent selling allowed them to avoid further losses rather than purely contributing to excessive volatility. A similar effect, but in the opposite direction, is observed for the other asset class with an estimated positive coefficient, emerging market equities. Since emerging markets experienced largely consistent positive returns with relatively low volatility through 2007, hedge funds who increased their exposures to these assets assisted in the price discovery process and thereby realized profits from contributing appropriate information which implied higher prices.

Pooled Feedback Trading Effects

	All Asset Classes	**Without MBS & EM**
β_0	-0.00011	-0.00009
	(0.00027)	(0.00031)
β_1	-0.00861	-0.01836^{*}
	(0.00957)	(0.01073)
R^2	0.00042	0.00197

$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$

Table 3.6: *Pooled Active Feedback Trading Effects*

In Table 3.6, we form a pooled OLS regression model using all the asset classes. That is, we estimate the following:

$$\left(1 + \gamma_{t+1|t+1}\right) w_{t+1,i|t+1} - \left(1 + \gamma_{t+1|t}\right) w_{t+1,i|t} = \beta_0 + \beta_1 r_{t-1,i} + \xi_{t,i}$$

where the coefficients β_0 and β_1 are assumed to be the same across all asset classes

i. Focusing on the fitted β_1 values, we find that although the pooled herding effect coefficient is still negative, the positive relationships from mortgage backed securities and emerging market equities hinder statistical significance. However, when estimating this pooled effect without those asset classes, we do find a statistically significant negative coefficient, implying that across the remaining asset classes, hedge funds exhibit negative-feedback trading behavior.

Finally, we note that although hedge funds commonly employ smoothing of their reported returns for various reasons, this does not pose a problem for the herding estimation results. Take for example a hedge fund who experiences a large loss/gain due to falling/rising prices on a particular asset over time period $t-1$. In order to smooth returns, the fund reports a higher/lower return than the true return, thereby implying a lower (invariably) relative weight on the depreciated asset at the beginning of time period $t-1$. Assuming the hedge fund amends this missing 'smoothed difference' in the next period, this piece is appropriately added onto the return over time period t. Since it is assumed that asset returns are independent over time, this missing difference is independent of returns on the given asset in that period. Therefore, although the exposure to this asset was effectively under-reported for $t-1$, and therefore under estimated, this added missing piece does not directly affect the weight estimation at time t or future periods since estimated weight levels will return to true levels. Hence, return smoothing should not create an artificial herding effect in our estimated results. Therefore, when the desmoothing procedure of Getmansky, Lo, and Makarov (2004) is performed on the initial hedge fund return series, these results do not change materially.

3.9 Conclusion

In this paper we presented a new analytical tool, the Leveraged Dirichlet Portfolio Model (L-DPM) for decomposing portfolio returns into their unobservable relative portfolio weights and leverage multiplier values. This compositional state space modeling technique allows us to overcome the significant analytical hurdle that the detailed set of hedge fund holdings is never completely observable for research. To compel our hedge fund estimation results, we constructed a dataset of the actively managed diversified equity subset of the mutual fund universe. Using this data, we formed an aggregate return index, as well as "true" time-varying portfolio weights on a division of equity industries. Using corresponding value-weighted industry portfolio returns, we were able to closely estimate the these true industry weights via the the L-DPM estimation procedure. Motivated by these empirical accuracy results for mutual funds, we estimated the dynamics of both asset class level portfolio holdings and leverage values on an index of hedge fund industry returns from 1995 to 2012.

With resulting estimates of these portfolio holdings, we discovered that net leverage levels in the hedge fund industry are smaller than popular belief due to netting both internally and across different funds. As well, using these estimates, we confirmed previous findings that hedge funds do not contribute to herding behavior in most asset classes, and in fact exhibit negative-feedback trading behavior in oil and municipal bonds.

Bibliography

[1] Aitchison, J. (1986). *The Statistical Analysis of Compositional Data*. Chapman & Hall, London.

[2] Amenc, N., L. Martellini, J. Meyfredi, and V. Ziemann (2010). Passive Hedge Fund Replication - Beyond the Linear Case. *European Financial Management*, 16(2), 191-210.

[3] Ang, A., M. Rhodes-Kropf, and R. Zhao (2008). Do Funds-of-Funds Deserve their Fees-on-Fees?. *Journal of Investment Management*, 6(4), 34-58.

[4] Appleson, J., E. Parsons, and A. Haughwout (2012). The Untold Story of Municipal Bond Defaults. *Liberty Street Economics*, Federal Reserve Bank of New York, August 15, 2012.

[5] Cargnoni, C., P. Müller, and M. West (1997). Bayesian Forecasting of Multinomial Time Series Through Conditionally Gaussian Dynamic Models. *Journal of the American Statistical Association*, 92(438), 640-647.

[6] Carlin, B., N. G. Polson, and D. Stoffer (1992). A Monte Carlo approach to nonnormal and nonlinear state-space modeling. *Journal of the American Statistical Association*, 87, 493-500.

[7] Carvalho, C., M. Johannes, H. Lopes, N. G. Polson (2010). Particle Learning and Smoothing. *Statistical Science*, 25, 88-106.

[8] Carvalho, C., M. Johannes, H. Lopes, and N. G. Polson (2011). Particle Learning for Sequential Bayesian Computation. *Bayesian Statistics*, 9, 317-360.

[9] Cassar, G. and J. Gerakos (2011). Hedge Funds: Pricing Controls and the Smoothing of Self-Reported Returns. *Review of Financial Studies*, 24(5), 1689-1734.

[10] Chipman, J. S. and M. M. Rao (1964). The Treatment of Linear Restrictions in Regression Analysis. *Econometrica*, 32(12), 198-209.

[11] Deng, G. and C. McCann (2012). Leveraged Municipal Bond Arbitrage: What Went Wrong?. *The Journal of Alternative Investments*, 14(4), 69-78.

[12] Engle, R. F. (2002). Dynamic Conditional Correlation: A Simple Class of Multivariate Generalized Autoregressive Conditional Heteroskedasticity Models. *Journal of Business and Economic Statistics*, 20, 339-350.

[13] European Central Bank (2008). *Financial Stability Review*. Frankfurt, Germany, December 2008.

[14] Fama, E. F. and K. R. French (1993). Common Risk Factor in the Returns on Stocks and Bonds. *Journal of Financial Economics*, 33(1), 3-56.

[15] Fama, E. F. and K. R. French (1992). The Cross-Section of Expected Stock Returns. *Journal of Finance*, 47(2), 427-465.

[16] Fearnhead, P. (2002). Markov Chain Monte Carlo, Sufficient Statistics, and Particle Filters. *Journal of Computational and Graphical Statistics*, 11, 848-862.

[17] Financial Services Authority (2011). Assessing the Possible Sources of Systemic Risk from Hedge Funds: A Report on the Findings of the Hedge Fund Survey and Hedge Fund as Counterparty Survey. Survey Report, London, July 2011.

[18] Financial Services Authority (2012). Assessing the Possible Sources of Systemic Risk from Hedge Funds: A Report on the Findings of the Hedge Fund Survey and Hedge Fund as Counterparty Survey. Survey Report, London, February 2012.

[19] Fisher, J., D. Gatzlaff, D. Geltner, and D. Haurin (2003). Controlling for the Impact of Variable Liquidity in Commercial Real Estate Price Indices. *Real Estate Economics*, 31, 269-304.

[20] Fonseca, T., M.A.R. Ferreira, and H.S. Migon (2008). Objective Bayesian Analysis for the Student-t Regression Model. *Biometrika*, 95, 325-333.

[21] Fung, W. and D. A. Hsieh (1997). Empirical Characteristics of Dynamic Trading Strategies: The Case of Hedge Funds. *Review of Financial Studies*, 10(2), 275-302.

[22] Fung, W. and D. A. Hsieh (2000). Performance Characteristics of Hedge Funds and CTA Funds: Natural vs. Spurious Biases. *Journal of Financial and Quantitative Analysis*, 35, 291-307.

[23] Fung, W. and D. A. Hsieh (2001). The Risk in Hedge Fund Strategies: Theory and Evidence from Trend Followers. *Review of Financial Studies*, 14, 313-341.

[24] Fung, W. and D. A. Hsieh (2004). Hedge Fund Benchmarks: A Risk Based Approach. *Financial Analysts Journal*, 60, 65-80.

[25] Fung, W. and D. A. Hsieh (2009). Measurement Biases in Hedge Fund Performance Data: An Update. *Financial Analysts Journal*, 65, 36-38.

[26] Gamerman, D. and H. F. Lopes (2006). *Markov Chain Monte Carlo: Stochastic Simulation for Bayesian Inference*. Chapman & Hall/CRC Press, Boca Raton, FL.

[27] Gelman, A. and J. Hill (2006). *Data Analysis Using Regression and Multilevel/Hierarchical Models*. Cambridge University Press, New York, NY.

[28] Getmansky, M., A. Lo, and I. Makarov (2004). An Econometric Model of Serial Correlation and Illiquidity in Hedge Fund Returns. *Journal of Financial Econometrics*, 74, 529-609.

[29] Gordon, N., D. Salmond, and A. F. M. Smith (1993). Novel Approach to Nonlinear/Non-Gaussian Bayesian State Estimation. *IEEE Proceedings*, F-140, 107-113.

[30] Grunwald, G. K., A. E. Raftery, and P. Guttorp (1993). Time Series of Continuous Proportions. *Journal of the Royal Statistical Society B*, 55(1), 103-116.

[31] Gupta, N. and R. Hauser (2007). Kalman Filtering with Equality and Inequality State Constraints. Technical Report, Oxford University Computing Laboratory.

[32] Hasanhodzic, J. and A. Lo (2007). Can Hedge-Fund Returns Be Replicated?: The Linear Case. *Journal of Investment Management*, 5, 5-45.

[33] Hasbrouck, J. (2007). *Empirical Market Microstructure*. Oxford Press, Oxford.

[34] Jacquier, E. and A. J. Marcus (2001). Asset Allocation and Market Volatility. *Financial Analysts Journal*, 57(2), 16-30.

[35] Jacquier, E., N. G. Polson, and P. Rossi (1994). Bayesian Analysis of Stochastic Volatility Models. *Journal of Business and Economic Statistics*, 12, 371-389.

[36] Jacquier, E., N. G. Polson, and P. Rossi (2004). Bayesian Analysis of Stochastic Volatility Models with Fat-Tails and Correlated Errors. *Journal of Econometrics*, 122, 185-212.

[37] Jegadeesh, N. and S. Titman (1993). Returns to buying winners and selling losers: Implications for stock market efficiency. *Journal of Finance*, 48, 65-91.

[38] Jones, S. (2012). Hedge funds keep a lid on leverage. *Financial Times*, April, 9, 2012. Retrieved May 29, 2013, from http://www.ft.com.

[39] Judge, G. G. and T. Takayana (1966). Inequality Restrictions in Regression Analysis. *Journal of the American Statistical Association*, 61, 166-81.

[40] Jurek, J. W. and E. Stafford (2012). The Cost of Capital for Alternative Investments. Working Paper, Bendheim Center for Finance, Princeton University.

[41] Kacperczyk, M., C. Sialm, and L. Zheng (2005). On the Industry Concentration of Actively Managed Equity Mutual Funds. *Journal of Finance*, 60(4), 1983-2011.

[42] Kadlec, G. and D. Patterson (1999). A Transactions Data Analysis of Nonsynchronous Trading. *Review of Financial Studies*, 12, 609-630.

[43] Kalman, R. E. (1960). A new approach to linear filtering and prediction problems. *Journal of Basic Engineering*, 82(1), 35-45.

[44] Khandani, A. E. and A. W. Lo (2007), What Happened to the Quants in August 2007?. *Journal of Investment Management*, 5, 5-54.

[45] Ko, S. and R. Bitmead (2007), State Estimation for Linear Systems with State Equality Constraints. *Automatica*, 43(8), 1363-1368.

[46] Kodres, L. and M. Pritsker (1996). Directionally-Similar Position Taking and Herding by Large Futures Market Participants. Unpublished Manuscript, International Monetary Fund, Washington.

[47] Korsos, L. F. (2013a). Hedge Fund Portfolio Decomposition: Joint Estimation of Net Leverage and Asset Allocation Dynamics. Working Paper, The University of Chicago Booth School of Business.

[48] Korsos, L. F. (2013b). The Dirichlet Portfolio Model: Uncovering the Hidden Composition of Hedge Fund Investments. Working Paper, The University of Chicago Booth School of Business.

[49] Kuan C.-M. and Hornik K. (1995), The Generalized Fluctuation Test: A Unifying View, *Econometric Reviews*, 14(2), 135-161.

[50] Liew, C. K. (1976). Inequality Constrained Least-Squares Estimation. *Journal of the American Statistical Association*, 71(355), 746-751.

[51] Liu, J. and M. West (2001). Combined Parameters and State Estimation in Simulation-Based Filtering. In *Sequential Monte Carlo Methods in Practice* (A. Doucet, N. de Freitas, and N. Gordon, eds.). Springer, New York.

[52] Lopes, H. F., C. Carvalho, M. Johannes, and N. G. Polson (2010). Particle learning for sequential Bayesian computation. In *Bayesian Statistics 9* (J. M. Bernardo, M. J. Bayarri, J. O. Berger, A. P. Dawid, D. Heckerman, A. F. M. Smith, and M. West, eds.). Oxford University Press, Oxford.

[53] Lopes, H. F. and N. G. Polson (2012). Particle Learning for Fat-tailed Distributions. Working Paper, The University of Chicago Booth School of Business.

[54] Lopes, H. F. and R. S. Tsay (2011). Particle filters and Bayesian Inference in financial econometrics. *Journal of Forecasting*, 30, 169-209.

[55] Madhavan, A. (2000). Market Microstructure: A Survey. *Journal of Financial Markets*, 3, 205-258.

[56] Malkiel, B. G. and A. Saha (2005). *Financial Analysts Journal*, 61(6), 80-88.

[57] Mamaysky, H., M. Spiegel, and H. Zhang (2007). Improved Forecasting of Mutual Fund Alphas and Betas. *Review of Finance*, 11, 359400.

[58] Mandelbrot, B. (1963). New Methods in Statistical Economics. *Journal of Political Economy*, 71, 421-440.

[59] Markowitz, H. M. (1952). Portfolio Selection. *The Journal of Finance*, 7(1), 77-91.

[60] Markowitz, H. M. (1959). *Portfolio Selection: Efficient Diversification of Investments*. John Wiley & Sons, New York, NY.

[61] Morris, S. and H.S. Shin (1999). Risk Management with Interdependent Choice. *Oxford Review of Economic Policy*, 15(3), 52-62.

[62] Neal, Radford M. (2003). Slice Sampling. *Annals of Statistics*, 31(3), 705-767.

[63] O'Hara, M. (1995). *Market Microstructure Theory*. Blackwell, Oxford.

[64] Persaud, A. (2000). Sending the Herd Off the Cliff Edge: The Disturbing Interaction Between Herding and Market-Sensitive Risk Management Practices. in Jacques de Larosiere Essays on Global Finance, Institute of International Finance, Washington.

[65] Polson, N. G., J. Stroud, and P. Müller (2008). Practical filtering with sequential parameter learning. *Journal of the Royal Statistical Society B*, 70, 413-428.

[66] Rubin, D. B. (1987). A noniterative sampling/importance resampling alternative to the data augmentation algorithm for creating a few imputations when fractions of missing information are modest: the SIR algorithm, comment to a

paper by Tanner and Wong. *Journal of the American Statistical Association*, 82, 543-6.

[67] Sharpe, W. F. (1992). Asset Allocation: Management Style and Performance Measurement. *Journal of Portfolio Management*, 18, 7-19.

[68] Shiller, R. J. (2000). *Irrational Exuberance*. Princeton University Press, Princeton, NJ.

[69] Smith, A. F. M. and A. E. Gelfand (1992). Bayesian statistics without tears. *American Statistician 46*, 84-88.

[70] Spellman, J. (2009). Hedge Funds: How They Serve Investors in US and Global Markets. *HedgeFundFacts.org*. Coalition of Private Investment Companies. August 2009.

[71] Storvik, G. (2002). Particle Filters in State Space Models with the Presence of Unknown Static Parameters. *IEEE Transactions of Signal Processing*, 50, 281-289.

[72] Tintner, G. (1952). *Econometrics*. John Wiley & Sons, New York, NY.

[73] Wermers, R. (2000). Mutual Fund Performance: An Empirical Decomposition into Stock-Picking Talent, Style, Transaction Costs, and Expenses. *Journal of Finance*, 55, 1655-1703.