Topics in Financial Risk Management and Fund Management

Department of Applied Finance and Actuarial Studies, Faculty of Business and Economics, Macquarie University Sydney, Australia

Haijie Weng

Thesis submitted to the Faculty of Business and Economics, Macquarie University, in fulfilment of requirement for the degree of Doctor in Philosophy

Singapore, October 2014

Declaration

I certify that except where due acknowledgement has been made, the work is that of the author alone; the work has not been submitted previously, in whole or in part, to qualify for any other academic award; the content of the thesis is the result of work which has been carried out since the official commencement date of the approved research program; any editorial work, paid or unpaid, carried out by a third party is acknowledged; and, ethics procedures and guidelines have been followed.

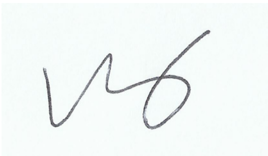

HAIJIE WENG

1 October 2014

Table of Contents

1. Abstract

This PhD thesis analyses three key problems in financial risk management and fund management. First, it tackles the problems of identifying risk contributors and performing Value-at-Risk analysis for Asian hedge funds. Second, it sheds light on agency problems affecting funds management in Australia. Last but not least, it discusses the problems of cleansing financial historical data essential for risk management. The thesis consists of three key chapters based on two published journal articles and one research paper.

Chapter 3, titled *Style Analysis and Value-at-Risk of Asia-focused Hedge Funds* has been published in the *Pacific-Basin Finance Journal, Volume 19 (2011)*. The chapter identifies risk factors and analyses Value-at-Risk (VaR) for Asia-focused hedge funds. Through a modified style analysis technique, we find that Asian hedge funds represented by Asian hedge fund indices show significant positive exposure to emerging equity markets. They also hold a significant portion of portfolio in cash and high credit rating bonds while taking short positions in world government and emerging market bonds. A rolling window style analysis is used to measure the time-varying risk exposure of Asian hedge funds. For both a static and rolling period style analysis, our model provides high explanatory power for returns on the considered hedge fund index. We further conduct a Value-at-Risk analysis using the results of a rolling window style analysis as inputs. Our results indicate that the accuracy of VaR models is dominated by their ability to capture the tail distribution of the hedge fund returns. Moreover, the distributional assumption seems to be more important than the chosen volatility model for the performance of the models in VaR prediction. Our findings further suggest that the considered parametric models outperform a simple historical simulation that is purely based on past return observations.

Chapter 4 is based on a journal article, titled *Agency Theory and Financial Planning*

Practice that has been published in the *Australian Economic Review, Volume 47 (2014).* The chapter extends an influential contribution to the literature on agency theory and then uses this extension, along with other theoretical contributions, to shed light on agency problems affecting funds management and financial planning in Australia. The case for pure fee for service in actively managed funds and plans turns out to be weak. The amount of money exposed to risk by an active manager should be less than the entire investible wealth of the client, especially in the case of investors on the cusp of retirement. Asset-based fees on actively managed funds should include a fulcrum component.

Chapter 5 titled *Backfilling Financial Data with an Iterative PCA-based Imputation* proposes an iterative PCA-based data imputation algorithm for handling missing values in financial time series. The designed backfilling algorithm generates satisfactory results for both simulated and empirical data, covering equity, rates and FX asset classes. Furthermore, our proposed model outperforms two of the most commonly used approaches for data imputation. Performance of our model depends on the fraction of the missing data and the noise of the data set. Our model serves as a robust tool for risk managers to backfill missing values in financial data, considering that complete data forms a prerequisite for generating a correct estimation of VaR and other performance measures.

2. Introduction

This PhD thesis deals with topics in financial risk management and fund management. The contribution of the thesis concentrates on three research areas:

- Risk management for Asia-focused hedge funds
- Fee structure in fund management
- Historical data for financial risk management

This introduction discusses the motivation and objectives for each of the research topics above and provides an overview of the thesis.

2.1 Risk Management for Asia-focused hedge funds

In the past decade, significant growth rates in Asian financial markets have attracted global investors' strong interest for capital allocation in Asia focused hedge funds. A number of studies concerned with measuring the performance and risk of hedge funds have been conducted in the literature already. In many of these studies, the performance of hedge funds, as alternative investments, is compared to traditional funds or asset classes (Ackermann et al., 1999; Brown et al., 1999; Liang, 1999; Agarwal and Naik, 2004). Some of the results in these studies suggest that hedge funds can outperform equity markets due to the superior investment skills of hedge fund managers (Brown et al., 1999; Liang, 1999), while other studies cast doubt on the persistence of the superior performance of hedge funds (Ackermann et al., 1999; Agarwal and Naik, 2004). From a risk management perspective, hedge funds are exposed to market risk, liquidity risk and credit risk (Amenc et al., 2002).

The performance and risk analysis of hedge funds may also be underestimated due to the presence of various biases in hedge fund indices as pointed out by Fung and Hsieh (2000). There are several difficulties as it comes to investigating the performances and risks of the hedge fund industry. The short data history of many hedge funds makes it difficult to compare the returns with those of traditional asset classes. Also, dynamic and less transparent investment strategies applied by hedge fund managers make it difficult to capture the effective style components for this asset class. Finally, hedge fund returns usually exhibit nonlinearities when being regressed on returns of traditional asset classes.

In order to explore the risk exposures of hedge funds, many researchers have attempted to map the returns onto a set of external factors. While the conventional return-based Sharpe's (1992) style analysis is commonly used in mutual fund analysis, Agarwal and Naik (2000) conduct a generalised style analysis of various hedge fund strategies by allowing negative style weights and relaxing the constraint that the sum of the style weights has to be one. They examine the significance of style weights by employing a two-step procedure initially proposed by Lobosco and DiBartolomeo (1997). Similarly, Dor et al. (2003) modify Sharpe's return based style analysis in order to examine the effective style of hedge funds. Therefore, the return based style analysis using traditional asset classes is augmented by index options to more appropriately characterize the risk of the hedge funds. Fung and Hsieh (2004) propose an asset based style factor model that can explain up to 80% of the monthly variation in hedge fund portfolios. More recently, Teo (2009) suggests augmenting the factor model of Fung and Hsieh (2004) with broad Asian equity indexes to study Asian focused hedge funds.

This chapter contributes to the literate in several dimensions. First, we make use of the return based style analysis framework suggested by Agarwal and Naik (2000) and Dor et al. (2003) to identify the effective style factors for Asia-focused hedge funds. To our knowledge, next to Teo (2009) this is one of the first empirical studies to apply this technique to the Asian hedge fund industry. Our model also differs from Teo (2009) who follows an approach similar to an APT (arbitrage pricing theory) model. In contrast, our approach is based on Sharpe's (1992) return based style analysis, in which there is no intercept term and the sum of coefficients is equal to one. Further, instead of averaging individual hedge fund returns as in Teo (2009), we adopt the HFRI Emerging Markets: Asia ex-Japan Index to represent the universe of Asia-focused hedge funds. Another contribution to the literature of the chapter is the focus on back-testing the proposed models in an extensive out-of-sample forecasting and risk analysis. We apply both parametric and nonparametric models and apply a variety of performance measures using VaR and density forecasts in combination with loss functions to examine the ability of the models to appropriately quantify the risk for the considered hedge fund index.

2.2 Fee Structure in Fund Management

This chapter examines agency problems affecting funds management and financial planning in Australia. Many studies have been devoted to modelling the behaviour of fund managers and investors, particularly in examining a fund manager's fee structure and its impact on investment decisions. Grinblatt and Titman (1989) apply option pricing theory to study performance-based fee contracts and show that option-like bonus components in performance-based compensation contracts can induce fund managers to seek excessive risk. To mitigate this adverse risk incentive, they suggest that contracts should be designed with a cap and apply penalties for underperformance. Similarly, Starks (1987) employs agency theory to study the impact of compensation contracts on fund managers' investment decisions. The author shows that contracts with symmetric payoffs dominate those with a bonus component in situations where there is asymmetric information between investors and fund managers. As a result investors cannot observe managers' efforts in selecting a portfolio's risk level.

In contrast, Das and Sundaram (2002) show that if asymmetric incentive fees and leveraging are allowed, an incentive fee with a large performance component provides a higher utility to fund managers and a low level of equilibrium volatility and, thus, a better risk sharing between the two parties.

Dybvig et al. (2010) consider three optimisation problems corresponding to increasingly severe agency problems. In the first-best case, agency problems are absent. In the second-best case the manager reveals truthfully the observed signal to the investor but has private information about her effort level. In the third-best case the adverse-selection problem and the moral-hazard problem are both present. They find that in a second-best case, the optimal contract should include a bonus in proportion to the fund's excess return over a benchmark return to give the manager incentives to work hard.

Other models, such as Carpenter (2000) and Cuoco and Kaniel (2001), assume fund managers are compensated with an option written on the portfolio's value. Carpenter (2000) specifies fee structure as a fixed component plus a call option with the portfolio value as underlying and benchmark return as strike price. Their findings are mixed. In general the option-like payoffs will increase a manager's appetite for risk – the manager will typically increase the volatility of the portfolio for a payoff to be "away from the money". However, if a manager were to trade his own account, he would occasionally target lower asset volatility. Similarly, Cuoco and Kaniel (2001) specify a put option on the managed portfolio as a penalty term in contracts for managers' poor performance. They find that the existence of a penalty for underperforming the benchmark induces fund managers to invest in benchmark's constituent stocks as a hedge against portfolio performance deviating from benchmark.

The chapter introduces generalised log utility into the setup of Dybvig et al (2010). Generalised log utility has the realistic implication that relative risk aversion is a declining function of wealth, unlike its log, quadratic, power and exponential competitors. Next, we examine Australian industry practice. In addition to our extension to Dybvig's model, we draw informally on the results of Stoughton et al. (2011), Bateman et al. (2007) and Grossman and Stiglitz (1976), to shed light on the agency problems raised by intermediated investment management, multiple time periods, and general equilibrium.

2.3 Dealing with Missing and Incorrect Data

This chapter deals with the problem of missing values in financial time series. Specifically it provides a new imputation algorithm that can help to backfill missing or incorrect data in financial time series such that it can be applied for risk management purposes or back-testing analysis in fund management.

The value-at-risk (VaR) concept has emerged as one of the most prominent measures of downside market risk, where the VaR is defined as the lower end of a 99% confidence interval for a given horizon (typically a day or two weeks). In 1997, a Market Risk Amendment to the Basle Accord permitted banks to use VaR estimates for setting bank capital requirements related to trading activity. To calculate VaR, historical simulation has been adopted by most banks as the standard industry approach, see e.g., Jorion (2000) and Alexander (2001).

One of the major concerns for historical simulation relates to the quality of the available data. Historic data is usually sourced from various data vendors and it is not uncommon for downloaded data to be of poor quality. Problematic data can be missing, meaningless or unlikely. Missing data may occur from a market close, or insufficient contributor depth in the case of composite quotation, or even system failure leading to the loss of data. In the case of meaningless data, although the data is present, it may violate some condition. For example, negative FX spot rates, negative FX option volatility, or a normally liquid time series that suddenly shows complete staleness: all these examples of data are meaningless. Unlikely data are those that behave abnormally, for example, a sudden spike of FX volatility in a normal market environment.

Once problematic data has been detected, correction protocols have to be applied.

Note that given the large set of time series to be cleansed and maintained, such a correction tool should run on a fully automated basis. The default corrective action is simply to re-query the original source system from which data is snapped. Should that fail to resolve the problematic value, the second action is to impute a value based on relevant statistical information about the time series and its nearby neighbours. For example, an error in the 7 year interest swap rate can be corrected with reference to its 5 year and 10 years swap rates.

To the best of our knowledge, limited research has been conducted in applying advanced and automated techniques to backfill financial time series. Karelmo (2010) uses a basic PCA based algorithm to fill the missing observations in corporate bond time series. Mailleta and Merlinb (2009) propose a way that does not require any hypothesis and is totally data driven to complete the missing values in hedge fund monthly return time series. Minsky et al. (2010) backfill missing return data for a hedge fund by randomly selecting peers' returns. All the studies above are either based on a simple approach or developed for a particular asset class. The chapter, by adopting the regularised PCA algorithm proposed by Josse et al. (2011), we propose an advanced PCA-based backfill procedure for financial time series and test the imputation performance using various asset class time series. The main focus of this chapter is on backfilling missing observations, which result from the removal of problematic data (i.e., missing, meaningless and unlikely data).

2.4 Structure of the Thesis

This PhD thesis consists of three main chapters, which can be assigned to the research areas discussed above in the following way:

- Risk management for Asia-focused hedge funds

  Style Analysis and Value-at-Risk of Asia-focused Hedge Funds, co-authored with Stefan Trück, Pacific-Basin Finance Journal, Volume 19 (2011) p.491–510.

- Fee structure in fund management

Agency Theory and Financial Planning Practice, co-authored with Geoffrey Kingston, Australian Economic Review, Volume 47.

- Historical data for financial risk management

Backfilling Financial Data with an Iterative PCA-based Imputation, co-authored with Stefan Trück.

3. Style Analysis and Value-at-Risk of Asia-Focused Hedge Funds

Haijie Weng (contribution: 80%) and Stefan Trück (contribution: 20%)

This paper has been presented at:

- 22nd Australasian Finance and Banking Conference, Sydney, August 2009
- Higher Degree Research Expo, Macquarie University, Sydney, November 2010
- European Financial Management Symposium, Beijing, March 2011

This paper has been published in *Pacific-Basin Finance Journal 19 (2011) p.491–510.*

Abstract

In this paper we identify risk factors for Asia-focused hedge funds through a modified style analysis technique. Using an Asian hedge fund index, we find that Asian hedge funds show significant positive exposures to emerging equity markets and also hold significant portion of portfolio in cash and high credit rating bonds while they take short positions in world government and emerging market bonds. A rolling window style analysis is further employed to analyse the time-varying risk exposure of Asian hedge funds. For both a static and rolling period style analysis, our model provides a high explanatory power for returns of the considered hedge fund index. We further conduct a Value-at-Risk analysis using the results of a rolling window style analysis as inputs. Our results indicate that the accuracy of VaR models is dominated by their ability to capture the tail distribution of the hedge fund returns. Moreover, the distributional assumption seems to be more important than the chosen volatility model for the performance of the models in VaR prediction. Our findings further suggest that the considered parametric models outperform a simple historical simulation that is purely based on past return observations.

3.1 Introduction

In the past decade, significant growth rates in Asian financial markets have attracted global investors' strong interest for capital allocation in Asia focused hedge funds. The expansion of the sector results in over 1,000 hedge funds focusing on Asian markets, representing over 15 percent of the total number of funds in the global industry. Although Asia-focused funds are characteristically smaller, accounting only for 4.9 percent of total industry assets, the significant growth rates of the Asia-focused hedge fund industry over the past years has also drawn the attention of the research community.

A number of studies concerned with measuring the performance and risk of hedge funds have been conducted in the literature already. In many of these studies, the performance of hedge funds, as alternative investments, is compared to traditional funds or asset classes (Ackermann et al., 1999; Brown et al., 1999; Liang, 1999; Agarwal and Naik, 2004). Some of the results suggest that hedge funds can outperform equity markets due to superior investment skills of hedge fund managers (Brown et al., 1999; Liang, 1999), while other studies cast doubt on the persistence of the superior performance of hedge funds (Ackermann et al., 1999; Agarwal and Naik, 2004). From a risk management perspective, hedge funds are exposed to market risk, liquidity risk and credit risk (Amenc et al., 2002). The performance and risk analysis of hedge funds may also be underestimated due to the presence of various biases in hedge fund indices as pointed out by Fung and Hsieh (2000). There are several difficulties as it comes to investigating the performances and risks of the hedge fund industry. The short data history of many hedge funds makes it difficult to compare the returns with those of traditional asset classes. Also dynamic and less transparent investment strategies applied by hedge fund managers make it difficult to capture the effective style components for this asset class. Finally, hedge fund returns usually exhibit nonlinearities when being regressed on returns of traditional asset classes.

In order to explore the risk exposures of hedge funds, many researchers have attempted to map the returns onto a set of external factors. While the conventional return-based Sharpe's (1992) style analysis is commonly used in mutual fund analysis, Agarwal and Naik (2000) firstly conduct a generalised style analysis of various hedge fund strategies by allowing negative style weights and relaxing the constraint that the sum of the style weights has to be one. They examine the significance of style weights by employing a two-step procedure initially proposed by Lobosco and DiBartolomeo (1997). Similarly, Dor et al. (2003) modify Sharpe's return based style analysis in order to examine the effective style of hedge funds. Hereby, the return based style analysis using traditional asset classes is augmented by index options to more appropriately characterize the risk of the hedge funds. Fung and Hsieh (2004) propose an asset based style factor model that can explain up to 80 percent of the monthly variation in hedge fund portfolios. More recently, Teo (2009) suggests augmenting the factor model of Fung and Hsieh (2004) with broad Asian equity indexes to study Asian focused hedge funds.

This paper aims to contribute to the literate in several dimensions. First, we make use of the return based style analysis framework suggested by Agarwal and Naik (2000) and Dor et al. (2003) to identify the effective style factors for Asia-focused hedge funds. To our knowledge, next to Teo (2009) this is one of the first empirical studies to apply this technique to the Asian hedge fund industry. Our model also differs from Teo (2009) who follows an approach similar to an APT (arbitrage pricing theory) model. In contrast, our approach is based on Sharpe's (1992) return based style analysis, in which there is no intercept term and the sum of coefficients is equal to one. Further, instead of averaging individual hedge fund returns as in Teo (2009), we adopt the HFRI Emerging Markets: Asia ex-Japan Index to represent the universe of Asia-focused hedge funds. Another contribution to the literature of this paper is the focus on backtesting the proposed models in an extensive out-of-sample forecasting and risk analysis. We apply both parametric and nonparametric models and apply a variety of performance measures using VaR and density forecasts in combination with loss functions to examine the ability of the models to give an appropriate quantification of the risk for the considered hedge fund index.

Conventional style analysis usually includes the broad range of traditional asset classes across the world. Since our focus is on Asian hedge funds, we augment the considered style factors by including the MSCI emerging markets Asia index, the MSCI Pacific excluding Japan index (developed markets in pacific region exclude Japan) and the MSCI Japan index in the return based style analysis for better explaining returns of Asia-focused hedge funds. Several studies on hedge funds show that the returns exhibit option-like features (Glosten and Jagannathan, 1994; Mitchell and Pulvino, 2001; Fung and Hsieh, 2001). Reasons for this are that hedge fund managers trade dynamically and are not limited to investing in a specific class of assets only. Hence the nonlinear payoff of a hedge fund may result from explicitly investing in derivatives or implicitly trading dynamically. In order to include the nonlinearity in hedge fund returns in the return based style analysis, the literature suggests using actively traded index options as nonlinear factors for the mapping of hedge fund returns; see e.g. Fung and Hsieh(2001), Agarwal and Naik (2004) and Teo (2009). Other studies suggest augmenting the return of traditional asset classes with the returns of synthetic options on these traditional asset classes, see e.g. Loudon et al. (2006). In this paper, we augment the trend-following factors created by Fung and Hsieh (2001) in style analysis to capture the option-like payoff of hedge fund's dynamic trading. The trend-following factors are created by using combinations of exchanged-traded put and call options in stock, bond, short term interest rate, currency and commodity. In summary, the selected style factors in this paper include global asset indices to cover the broad range of asset classes that Asian hedge fund managers can invest and trend-following style factors to capture the option-like payoff resulting from hedge fund managers' dynamic trading activities.

Similar to Agarwal and Naik (2000) and Dor et al (2003), we relax the style analysis conditions by allowing negative weights of style factors since hedge fund managers often take short positions to exploit arbitrage opportunities or hedge the portfolio against market movements. Ideally, the factors used in the style analysis are independent; however, in practise, the chosen factors will fall short of ideal and sometimes will have high correlations with other factors. Therefore we need to eliminate the redundant factors which can be replicated by others and keep the remaining factors independent as much as possible. To address this issue, we further employ the two-step procedure proposed by Lobosco and Dibartolomeo (1997) to determine the statistical significance of factor weights. Finally, as shown by many researchers (e.g. Fung and Hsieh 2004; Bollen and Robert 2008), hedge fund managers change trading strategies over time; therefore we perform style analysis using rolling estimation to provide insights of time variation of hedge fund exposures.

Our findings suggest that Asian hedge funds show significant positive exposures to emerging equity markets, especially emerging markets in Asia, and also hold significant portion of portfolio in cash and high credit rating bonds while short sell world government bond and high yield emerging market bonds. Further, small but statistically significant exposures to trend-following factors show the option-like payoff pattern of Asian hedge funds. In general, the style analysis can explain up to 82% of the variance of hedge fund returns, indicating a high explanatory power of our model. Finally, the style analysis on rolling window sheds light on how hedge fund managers change risk exposures over time in response to changing market conditions and arbitrage opportunities.

Second, the ultimate goal for identifying the underlying risk exposures of the hedge fund is to evaluate the risk of the hedge funds. In our analysis we use the Value-at-Risk (VaR) measure, defined as the maximum loss with a given confidence level over a given period of time. VaR can provide information about the risk in the extreme tails of a distribution. This is of particular importance, since many hedge funds exhibit a nonlinear payoff structure, that is, hedge funds may face great losses under certain extreme events although they have an average low standard deviation. The nonlinear exposures also lead to a situation where the normality assumption of expected returns that suggests the use of the standard deviation as the only risk measure is no longer justified. Therefore, for hedge funds, VaR, as a complementary tool for measurement of the risks, can better capture the behaviour of hedge funds in some extreme events.

Many methods have been proposed to calculate VaR (see Duffie and Pan 1997, Hull and White 1998, Jorion 2000). In general, they can be categorised as nonparametric and parametric approaches. Nonparametric VaR makes no assumption about the shape of the distribution of returns. For example, historical simulation as a nonparametric method assumes that historical returns can provide an appropriate evaluation of the risk and therefore estimates the VaR based on the empirical distribution of past observations. On the other hand parametric VaR assumes that the distribution of returns belongs to a parametric family, such as normal distribution. It usually applies a two-step approach: first, it is assumed that the portfolio variance is governed by certain specifications, such as a covariance matrix specification of the underlying assets, or a time-varying variance of the aggregate portfolio. It further makes distributional assumption about the portfolio returns and then calculates the VaR based on the estimated parametric dependence structure and return distributions. For example, portfolio variance can be modelled as a GARCH process or exponential weighted moving average (EWMA) specification. The distribution of portfolio returns is often assumed to be normal or from the Student's t distribution.

In this paper, we examine VaR for the considered hedge fund index using both parametric and nonparametric techniques. For the parametric approach, two methods are used to estimate the time-dependent portfolio variance: covariance matrix forecasts estimated from a wide variety of multivariate volatility models and aggregate portfolio variance forecasts. Loss functions are employed to evaluate the quality of the competing volatility models. Portfolio returns are assumed to be either from the normal or Student t distribution. For the nonparametric approach, VaR is calculated using historical simulation with a rolling window including 100 months of observations. To evaluate the different approaches with respect to their ability to appropriately quantify the risk, we employ different methods to determine the accuracy of the VaR forecasts. Next to examining the coverage and numbers of exceptions for the considered VaR models, we also investigate density forecasts and the magnitude of the occurred exceptions. Practitioners and researchers are interested not only in the frequency of the VaR exceptions, but also in the magnitude of the loss when the VaR is violated. Therefore, we employ a hypothesis test initially proposed by Berkowitz (1999) to examine whether the magnitudes of the observed violations are consistent with those implied by the proposed VaR models.

Our empirical results show that in general, the direct hedge fund index variance forecast (H-EWMA and H-GARCH models) outperform the forecast based on covariance matrix specification in term of hedge fund variance forecasting. However, under the VaR loss functions, the results show that VaR model based on the Student t distribution outperform those based on a normal distribution regardless of the chosen model for the volatility. These results suggest that the distributional assumption for the returns might be of greater importance than the model that is used for volatility dynamics. Because our out-of-sample data covers the Global Financial Crisis period when hedge funds also suffered from significant negative returns, it indicates that in this paper the performance of a VaR model is dominated by its ability to capture the tail distribution of hedge fund returns correctly. Moreover, we find that most of the considered VaR Models perform well with respect to the magnitude of VaR exceptions.

The remainder of the paper is structured as follows. Section two describes the hedge fund data used in this paper. Section three presents the style analysis technique used in this paper as well as the empirical results for the considered style factors. Section four presents the examined approaches for our risk analysis and evaluates the out-of-sample Value-at-Risk forecasts for the considered models. Finally, section five concludes.

3.2 Data

When performing analysis of hedge funds, data can be collected by either averaging individual hedge fund returns or using hedge fund indices directly. It is important to keep in mind that hedge fund indices can inherit biases existing in the hedge fund databases. Hedge fund data are susceptible to selection bias, survivorship bias and instant history bias as discussed by Fung and Hsieh (2004). Hedge fund managers voluntarily report the returns to the data vendors, so selection bias can arise if the hedge fund data collected in the database cannot represent the whole universe of the hedge funds. Survivorship bias occurs if the database only contains information on operating funds. Defunct funds may stop reporting to the database because of bad performance, termination or other reasons like e.g. mergers. When a fund is included in a database, its past track record is appended to the database, which creates instant history bias. Funds often undergo an incubation period before reporting to the database. Funds with good performance then go on to list in various databases for seeking new investors, while unsuccessful funds will not enter the databases. Thus backfilling the past performance into the database may generate an upward bias. Recognizing these biases, some database vendors construct the indices with the care to mitigate the effects of these errors inherited from the databases. When working with hedge fund indices, it is essential to choose those indices that are less prone to these biases.

In this paper, we choose to work with an Asian hedge fund index rather than individual hedge fund returns. There are two major providers for Asian hedge fund indices: Eurekahedge and Hedge Fund Research (HFR). The Eurekahedge database mainly includes funds with investments in the Asia-Pacific region, while HFR is a large global hedge fund database. Eurekahedge provides the explicit information on the fund main investment region; in contrast, HFR classifies a fund as Asian hedge fund if the fund has more than 50% of its investments in the Asia ex-Japan region. Eurekahedge and HFR started to collect hedge fund return data from January 2000 and January 1990, respectively. There are some differences in constructing the indices between the two databases. For example, unlike Eurekahedge, HFR has a requirement that included funds have at least $50 Million under management or have been actively traded for at least twelve months. In the HFR index, the historical performance of a new constituent fund will not affect the finalized historical performance of the index. In contrast, Eurekahedge backfills the new constituent funds with past performance and rebalances the index value, which is prone to instant history bias. Considering that the HFR has a longer performance history and is less prone to instant history bias, we decide to use a HFR Asia index (HFRI Emerging Markets: Asia ex-Japan Index) as a proxy to investigate the style factors and risk of Asia-focused hedge funds. We consider monthly returns of the index for the period January 1990 to April 2010.

3.3 Style Analysis of Asia-Focused Hedge Funds

This section provides empirical results on the conducted style factor analysis on Asia-focused hedge funds. In a first step we identify appropriate style factors to use in the style analysis. In a second step we apply the style analysis framework proposed by Sharpe (1992) to identify the risk exposures of Asia-focused hedge fund managers. Furthermore, we employ the two-step procedure proposed by Lobosco and Dibartolomeo (1997) to determine the statistical significance of factor weights. Finally, we perform the style analysis using a rolling estimation framework in order to examine the time variation of the factor weights and hedge fund exposures.

3.3.1 Style factors of Asian Hedge Funds

It has been noted by many researchers that hedge fund returns are related to returns from traditional asset classes (e.g. Fung and Hsieh 2001, 2002, 2003; Agarwal and

Naik 2000, 2004; Mitchell and Pulvino 2001; Dor et al. 2003). Fung and Hsieh (2002) find that fixed income hedge funds are typically exposed to interest rate spreads. This is a result of many fixed income hedge funds rather holding long positions in high yield bonds and hedging the interest rate risk by shorting treasury bills or bonds. Further, Fung and Hsieh (2003) show that equity long/short hedge funds tend to take long positions in low capitalization stocks and short positions in large capitalization stocks, such that hedge fund returns are typically exposed to the spread between large cap and small cap stocks. In an extension of their prior work, Fung and Hsieh (2004) propose a model of hedge fund returns using seven identified asset based style (ABS) factors. For diversified hedge fund portfolios, the seven ABS factors can explain up to 80 percent of monthly return variations. The seven ABS factors include two equity ABS factors (equity market return and spread between small-cap stock returns and large-cap stock returns), two fixed income ABS factors (change in 10 year Treasury yields and change in the yield spread between 10 year T-bonds and Moody's Baa bonds) and three trend following ABS factors (lookback straddles on bonds, currencies and commodities).

Agarwal and Naik (2000) conduct a generalised style analysis of various hedge fund indices. To cover the broad range of the asset classes hedge fund managers may invest in, they use the S&P 500 composite index, the MSCI world index excluding US and MSCI emerging market index to proxy the global equity market exposures. They further choose the Salomon Brothers (SB) Government and Corporate Bond index, the SB World Government Bond index and the Lehman High Yield index to assess the bond market exposure. Finally, they include a number of commodity and currency indices to account for the hedge funds' exposure to these variables. Similarly, Dor et al. (2003) perform a return-based style analysis to examine the effective style of hedge fund managers by using traditional asset classes and index options. They select the asset classes aiming to cover the equity and fixed income investment in the US and outside the US. For instance, they use 3-month Treasury bills as cash equivalent, intermediate and long term bonds and US corporate bonds to represent fixed income investment in the US, the Russell 1000 and 2000 index to represent equity investment in US as well as four global equity and fixed income indices to represent foreign investments. Applying principal component analysis, Teo (2009) shows that the Asia exclude Japan equity market index and Japan equity market index both are highly correlated with the returns of Asia equity hedge funds.

In addition to an exposure to traditional asset classes, many researchers argue that due to dynamic trading, hedge fund returns often exhibit non-linear option-like exposures to standard asset classes (Fung and Hsieh 1997, 2001; Agarwal and Naik 2004). Further, Agarwal and Naik (2004) illustrate that the payoffs of a large number of equity-oriented hedge funds actually resemble a short position of a put option on the market index. To capture this option-like feature of hedge fund returns, Fung and Hsieh (2001) create style factors by using combinations of exchanged traded put and call options in stocks, bonds, interest rates as well as currency and commodity markets. Further, Bollen and Robert (2008) employ the five trend-following factors used in Fung and Hsieh (2001) to investigate the time-series variation in hedge fund risk exposures. Similarly, Agarwal and Naik (2004) use actively traded at-the-money (ATM) and out-of-the-money (OTM) European call and put options on the S&P 500 composite index as option based risk factors to capture the option-like features of hedge fund returns. Following Agarwal and Naik (2004), Teo (2009) uses OTM European call and put options on the Nikkei225 traded on the Singapore Stock Exchange and calculates the time series of returns for the option trading strategy in a similar way. Given the lack of actively traded options for the identified index factors, Loudon et al (2006) create pseudo option-like payoff profiles for a subset of index factors to model the nonlinear exposures that fixed income hedge funds may face.

In this paper, we select global asset indices to cover the investment regions including Asia and the rest of the world and employ the five trend-following factors used in Fung and Hsieh (2001) to capture the option-like payoff of hedge fund dynamic trading strategies. The global asset indices included in our style analysis cover cash, equities and bonds markets. In particular, we use 3-month Treasury bills as cash equivalent. To proxy the exposure to Asia and global equities, we include the MSCI emerging markets Asia index, the MSCI Pacific excluding Japan index (developed markets in the pacific region excluding Japan), the MSCI Japan index, the S&P 500 index, the MSCI Europe index (developed markets in Europe) and the MSCI emerging markets excluding Asia index. To capture the exposures to bonds, we consider the Bank of America Merrill Lynch US High Bond index, the JP Morgan emerging markets bond Asia index, the CGBI broad investment grade (BIG) index and the CITI world government bond index. The five trend following factors are the returns of a private trend following strategy (PTFS) lookback straddles in bonds, currencies, short term interest rates, commodities and stocks. In total, we use eleven asset indices and five trend-following factors. Appendix A provides a more detailed description of the selected style factors.

3.3.2. Style Analysis

After having identified the style factors, we can conduct a return based style analysis for the hedge fund returns. Sharpe (1992) proposed an econometric technique to determine the mutual fund's investment style which requires a time series of historical fund returns. This technique involves a constrained regression that uses *N* asset classes to replicate the historical return pattern of a fund. The style analysis framework for modelling the fund return is as follows:

$$r_t = \sum_{i=1}^{N} w_i F_{i,t} + e_t \tag{1}$$

where r_t is the fund return at time t, $F_{i,t}$ is the return of the i^{th} style factor at time t, $i = 1, \dots, N$, w_i is the corresponding factor weight, and e_t represents the error term.

Style analysis has been initially proposed to analyse mutual funds. Because the weights of the replicated asset classes should add up to unity and mutual fund managers are not allowed to take short positions, Sharpe (1992) imposed two constraints on the coefficients w_i:

$$\sum_{i=1}^{N} w_i = 1,\ \forall i \tag{2a}$$

$$w_i \geq 0,\ \forall i \tag{2b}$$

When applying return based style analysis to hedge funds, the constraint of nonnegative coefficients is usually released to allow hedge fund managers also to take short positions in the various asset classes (Agarwal and Naik 2000; Dor et al 2003).

Based on Eq. (1), the excess return of the hedge fund over the sum of the weighted factor returns can be expressed as $e_t = r_t - \sum_{i=1}^{N} w_i F_{i,t}$. Sharpe (1992) suggests choosing the optimal weights for w_i through minimising the term e_t or rather the variance of e_t subject to constraint (2). This can be achieved for example by quadratic programming. To evaluate the effectiveness of the style analysis, we use the coefficient of determination (R^2) or adjusted R^2 given by $R^2 = 1 - \frac{var(e_p)}{var(r_p)}$ and ${R_{adj}}^2 = 1 - \frac{T-1}{T-N} \times \frac{var(e_p)}{var(r_p)}$, where N is the number of style factors, T the number of observations, $var(e_p)$ the variance of the residuals and $var(r_p)$ is variance of the hedge fund returns. Often, for hedge fund analysis these measures are interpreted as R^2 indicating the proportion of return variance attributable to investment styles while the unexplained part $(1 - R^2)$ is attributable to the fund manager's skill. In contrast to R^2, the adjusted R^2 has the advantage of imposing a penalty an increased number of style factors.

Ideally, the factors used in style analysis need to be independent; however, in practise, the chosen factors will fall short of the ideal and sometimes will have high correlations with other factors. To address this issue, we therefore employ a two-step procedure initially proposed by Lobosco and Dibartolomeo (1997) to determine the statistical significance of the factor weights. In a first step we conduct the analysis using all style factors and then calculate the standard deviation of the residuals (σ_e). Then we perform a style analysis for each style factor using the remaining style factors as explanatory variables calculating the standard deviation of the residuals (σ_i) for style factor i. The latter style analysis is estimated with the constraint that the sum of weights is one. The standard error for the weight of style factor i is given by $\frac{\sigma_e}{\sigma_i\sqrt{N-k-1}}$, where N is the number of observations and k is the number of style factors with non-zero weight. A low standard error indicates that the factor is difficult to be replicated by other style factors. The t-statistic for each factor is given by $\frac{w_i\sigma_i\sqrt{N-k-1}}{\sigma_e}$, where w_i is the weight for factor i . Based on the calculated t-statistics, using a 5% significance level non-significant factors are excluded from the model. This two-step procedure is repeated until the remaining factors are all statistically significant.

3.3.3 Empirical Results for the Style Analysis

As mentioned above we investigate monthly returns of the HFRI Emerging Markets: Asia ex-Japan Index obtained from the HFR database. The style factors including eleven asset indices and five trend-following factors are obtained from Datastream and David Hsieh's Hedge Fund Data Library respectively. Both hedge fund and style factor returns are considered for the period January 1994 to December 2009 including a total of 192 observations. Table 1 reports the descriptive statistics for the hedge fund and style factors returns. The average monthly hedge fund index return is 0.62% while the standard deviation of monthly returns is 3.79%. The bond factors, in general, appear to have positive mean, negative skewness and high kurtosis. Among these, the Asia bond index has the highest return but also exhibits the highest negative skewness and kurtosis indicating that the lower tail of the distribution is longer than the upper tail, as well as a heavy-tailed distribution. Similar results are obtained for the equity factors apart from the Japanese equity index, which yields a negative mean, positive skewness and low kurtosis during the considered time period. The five trend-following factors appear to have the largest standard deviation among all the style factors. We further note that the Asian hedge fund, world government bond index and Japan equity index are normally distributed during the sample period.

Table 2 provides the results for the conducted style analysis: the first column shows the weights with standard errors for all style factors. The second column shows the results of the style analysis after dropping the insignificant factors using the recursive procedure described above. We find that Asian hedge funds show significant style weights on the three-month T-bill, the world government bond index, the US broad investment grade index, the Asian bond index, Japan equity, emerging market (Asia and the rest of emerging market) equity, and three trend-following indices on short term interest rates, currencies and stocks. In particular, Asian hedge funds show a significant positive exposure to emerging equity markets with 34.5% to the Asia market index and 10.3% to the MSCI emerging markets excluding Asia index and a small positive exposure to the Japanese equity index. Since the hedge fund index studied in this paper includes hedge funds investing in emerging markets with primary focus on Asia and typically less than 10% exposure to Japan, our finding is consistent with the classification of the fund index. The long position in emerging equity is also consistent with the typical short selling restrictions in emerging equity markets. The Asian hedge funds also show positive style weight on three-month T-bill and US corporate bond index, but negative style weights on world government bond and emerging market bond Asia indices. The net exposure to the bond market is approximately 45%. This suggests that Asian hedge funds hold significant portions of portfolio in cash and high credit rating bonds while they short sell world government bonds and high yield emerging market bonds. Further, small but statistically significant exposures to trend-following factors show the option-like payoff pattern of Asian hedge funds. The style analysis can explain up to 82% of the variance of hedge fund returns, the remaining unexplained variance being attributed to managers' trading skill. Moreover, the explanatory power of the regression model remains almost unchanged after eliminating the insignificant style factors.

Unlike mutual funds that follow a defined investment strategy and therefore are not allowed to change their investment styles, hedge funds are generally free to change trading strategies and asset allocation to different asset classes. Assuming the style weights are constant, the above style analysis shows an average risk exposure of Asian hedge funds over the sample period from January 1994 to December 2009. To investigate the hedge funds' dynamic risk exposure over time, we perform the style analysis using a rolling window of 72 months. Figure 1 shows the style changes for the HFRI Emerging Market-Asia exclude Japan index over time. We find that Asian hedge funds experience significant shifts in risk exposure over time. Furthermore, the major style factors are the three-month T-bill, emerging market Asia bond index, emerging market equity and Japan equity. Figure 2 shows the adjusted R^2 when factors weights change over time. The adjusted R^2 is 80% on average, indicating a high explanatory power of the rolling-period style analysis. In the next section, the results of the rolling period style analysis will be used as inputs for an extensive risk analysis of the considered hedge fund index.

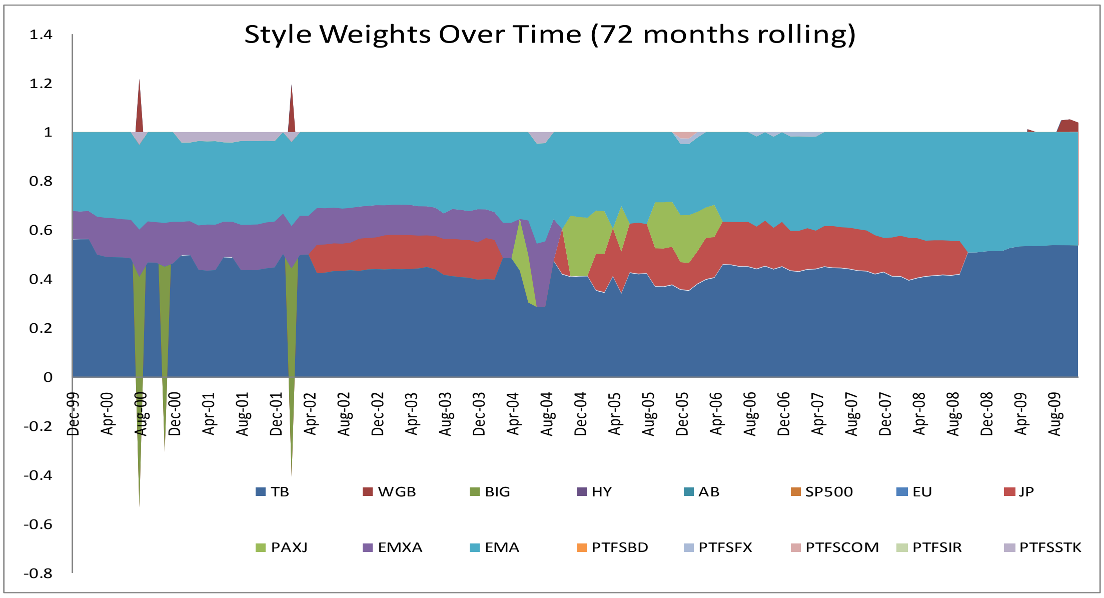

Figure 1: The dynamic exposure to the considered style factors of the HFRI emerging markets Asia excluding Japan index based on a rolling window approach with length of 72 months

	Mean	Median	Max	Min	Standard Deviation	Skewness	Kurtosis	Normality Test
HF	0.62	0.72	12.37	-11.02	3.79	-0.07	3.43	1.68
TB	0.29	0.36	0.51	0.00	0.15	-0.50	1.78	19.89*
WGB	0.42	0.49	3.38	-1.91	0.89	-0.05	3.36	1.13
BIG	0.51	0.61	4.44	-3.44	1.14	-0.17	4.17	11.95*
HY	0.54	0.85	7.15	-15.42	2.54	-1.63	11.94	724.69*
AB	0.72	0.85	9.10	-17.64	2.80	-2.42	17.42	1,851.24*
SP500	0.45	1.14	9.23	-18.56	4.54	-0.95	4.70	51.84*
EU	0.44	1.08	12.36	-23.98	5.09	-1.09	6.05	112.59*
JP	-0.12	-0.31	15.43	-16.00	5.63	0.07	2.89	0.24
PAXJ	0.22	0.85	17.43	-28.90	6.45	-0.88	5.82	88.78*
EMXA	0.60	1.96	17.55	-40.82	8.06	-1.46	7.87	257.50*
EMA	-0.02	0.05	19.44	-27.65	7.83	-0.47	3.69	10.71*
PTFSBD	-1.38	-4.82	68.86	-25.36	14.73	1.46	6.00	140.13*
PTFSFX	0.19	-4.31	90.27	-30.13	19.82	1.37	5.63	114.86*
PTFSCOM	-0.30	-2.90	64.75	-23.04	13.92	1.26	5.54	102.62*
PTFSIR	3.12	-3.14	221.92	-30.60	28.89	4.09	25.57	4,609.18*
PTFSSTK	-4.73	-6.32	46.15	-30.19	12.84	0.98	4.88	59.08*

Table 1: Descriptive statistics of hedge fund and style factor returns

The table shows the mean, median, standard deviations, minimum and maximum returns, skewness, kurtosis and results for a normality test for the hedge fund index and selected style factors during the period January 1994 to December 2009. The hedge fund index is the HFRI emerging markets Asia excluding Japan (HF). The style factors are three-month T-bill (TB), CITI world government bond index (WGB), CGBI broad investment grade index (BIG), Bank of America Merrill Lynch US High Bond index (HY), JP Morgan emerging markets bond Asia index (AB), S&P 500 index (SP500), MSCI Europe index (EU), MSCI Japan index (JP), MSCI Pacific excluding Japan index (PAXJ), MSCI emerging markets excluding Asia index (EMXA), MSCI emerging markets Asia index (EMA), bond PTFS (PTFSBD), currency PTFS (PTFSFX), commodities PTFS (PTFSCOM), short term interest rate PTFS (PTFSIR) and stock PTFS (PTFSSTK). The normality test is the Jarque-Bera Test which has a χ^2 *distribution with 2 degree of freedom under the null hypothesis of normal distribution. The 5% critical value is 5.99. The asterisk indicates statistical significance at 5%.*

	HF	HF
TB	**78.3**	**74.2**
	(15.4)	(14.7)
WGB	**-64.2**	**-64.1**
	(28.6)	(28.3)
BIG	**50.1**	**51.3**
	(22.9)	(22.3)
HY	1.0	
	(6.4)	
AB	**-17.6**	**-16.6**
	(6.4)	(6.1)
SP500	-2.5	
	(5.1)	
EU	-6.0	
	(4.8)	
JP	**7.8**	**7.4**
	(2.7)	(2.6)
PAXJ	6.6	
	(4.4)	
EMXA	**10.4**	**10.3**
	(2.8)	(2.3)
EMA	**33.0**	**34.5**
	(3.0)	(2.3)
PTFSBD	-1.0	
	(0.9)	
PTFSFX	**1.3**	**1.5**
	(0.7)	(0.6)
PTFSCOM	1.3	
	(0.9)	
PTFSIR	**-1.1**	**-1.0**
	(0.5)	(0.5)
PTFSSTK	**2.5**	**2.5**
	(1.0)	(1.0)
Adjusted R^2	82.64	82.49

Table 2: Style analysis of Asian hedge fund index

This table shows the results for style analysis of the Asian hedge fund index from January 1994 to December 2009. The first column shows the weights with standard errors for all style factors. Standard errors for style weight are in parentheses. The weights significant at 5% level are expressed in bold font. The second column shows the results of the style analysis after dropping the insignificant factors through repeated procedure. All data are in percentage. The adjusted coefficient of determination R^2 is reported as well.

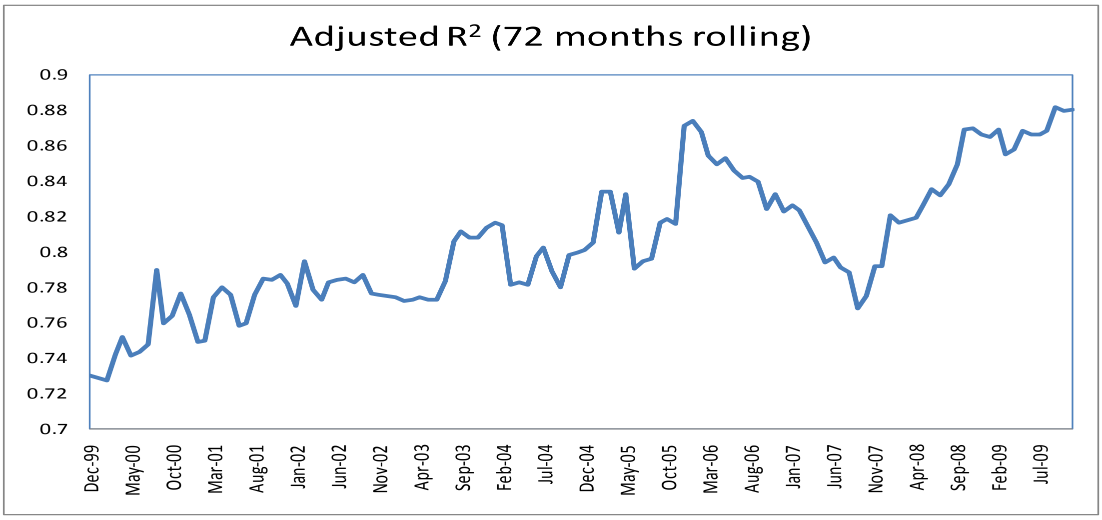

Figure 2: Adjusted R^2 of the rolling window style analysis.

3.4 Value-at-Risk Analysis

In the previous section we have identified risk factors for the considered Asian hedge fund index and examined the dynamic nature of the risk exposure to the identified factors. In this section we conduct a thorough Value-at-Risk analysis for the considered hedge fund index using various benchmark models and backtesting techniques. In particular the performance of different approaches to modelling the volatility of the index and the style factor returns are considered.

3.4.1 Modelling the conditional volatility

In order to evaluate the performance of the style factor analysis with respect to risk quantification, an adequate approach for modelling the conditional variance of the index and factor returns is required. Therefore, we start our analysis with a description of the considered models for volatility in the empirical analysis. Let y_t and $r_{i,t}$ denote the return of the hedge fund index and style factor i at time t. Investigating the autoregressive nature of returns, we find that, generally, the considered time series do not indicate significant ARMA dynamics. An exception is the return series of the three-month T-bill, which yields significant first-order autocorrelation. Further, we conduct augmented Dickey-Fuller unit root tests for the three-month T-bill to test the stationarity of the data series, and find that it has a unit root such that the series is non-stationary. Therefore, we express the returns of the three-month T-bill as

$$r_{b,t} = r_{b,t-1} + \varepsilon_{b,t} = r_{b,t-1} + \eta_t\sqrt{h_{b,t}} \tag{3}$$

Since our focus is on volatility forecasting, we model the hedge fund index and the style factors using the following model:

$$y_t = u_t + \epsilon_t = u_t + \eta_t\sqrt{h_t} \tag{4}$$

$$r_{i,t} = \mu_{i,t} + \varepsilon_{i,t} = \mu_{i,t} + \eta_t\sqrt{h_{i,t}} \tag{5}$$

where u_t and $\mu_{i,t}$ are the conditional mean for the hedge fund index and factor *i* at time *t*. Further, h_t and $h_{i,t}$ are conditional variance for hedge fund index and factor *i* at time *t*, and η_t is an *iid* process with zero mean and unit variance. Using a rolling window analysis, therefore, the conditional mean for the three-month T-bill is its past return at time $t-1$, while the conditional mean for the hedge fund index and the other style factors is equal to the mean return over the past 72 months. Let further $w_t = [w_{1,t}, w_{2,t}, \dots, w_{n,t}]$ denote the style weights vector at time *t* estimated from the rolling window style analysis, such that the forecasted hedge fund conditional variance at time $t+1$ in covariance matrix specification is given by

$$h_{t+1} = w_t \times H_{t+1} \times w_t' + h_{e,t} \tag{6}$$

where H_t is a *n*x*n* matrix with *n* being the number of significant non-zero style factors. Taking $n = 2$, for example, $H_{t+1} = \begin{bmatrix} h_{11,t+1} & h_{12,t+1} \\ h_{21,t+1} & h_{22,t+1} \end{bmatrix}$, while $h_{e,t}$ is the conditional variance of the corresponding residuals of the rolling window style analysis, which is assumed to be normally distributed with variance $\sigma_{e,t}^2$.

To generate appropriate covariance matrix forecasts, we then apply equally weighted moving average, exponentially weight moving average (EWMA) and GARCH-BEKK models.

The equally weighted moving average (MA) model puts equal weights on the past priod observations, taking the form:

$$H_{t+1} = \sum_{k=0}^{72} \varepsilon_{t-k}\, \varepsilon_{t-k}{}' \tag{7}$$

where $\varepsilon_t = [\varepsilon_{1,t}, \varepsilon_{2,t}, \dots, \varepsilon_{n,t}]$ is the vector containing the style factor's innovations at time *t*. In contrast, the exponentially weight moving average (EWMA) model is based on exponentially decreasing weights, i.e., more weight is given to more recent observations:

$$H_{t+1} = \sum_{k=0}^{72} (1-\lambda)\lambda^k \varepsilon_{t-k}\, \varepsilon_{t-k}{}' \tag{8}$$

where λ is the decay factor that is set equal to 0.97 following the weight suggested by RiskMetrics (JP Morgan 1996).

The third method used in this paper for forecasting the covariance matrix is a multivariate GARCH model. Multivariate GARCH models provide estimates for the conditional covariance as well as the conditional variances in contrast to univariate models and have gained high popularity in modelling and forecasting multivariate time series. For example, Gibson and Boyer (1998) compare the correlation forecasting ability of three sophisticated models (two GARCH models and a two-state Markov switching model) and two simple moving average models and find that the sophisticated models (a diagonal GARCH and a Markov switching approach) produce better correlation forecasts than the simple moving averages. Multivariate GARCH models specify equations for the behaviour of the variance covariance matrix through time. Several different multivariate GARCH formulations have been proposed in the literature, including the VECH, the diagonal VECH and the BEKK model, see e.g. Bauwens et al. (2006) for a survey on the most important developments in multivariate GARCH modelling. In our analysis we suggest to use a GARCH BEKK (Baba-Engle-Kraft-Kroner) model (Engle and Kroner, 1995). This model overcomes some of the difficulties of the VECH model by ensuring that the conditional variance-covariance matrix is always positive definite. The model has the form

$$H_{t+1} = CC' + \sum_{j=1}^{p} \sum_{k=1}^{K} A_{kj}' H_{t+1-j} A_{kj} + \sum_{j=1}^{q} \sum_{k=1}^{K} B_{kj}' \varepsilon_{t+1-j} \varepsilon_{t+1-j}' B_{kj} \qquad (9)$$

where A_{kj}, B_{kj} are parameter matrices and C is a lower triangular matrix. The decomposition of the constant term into a product of two triangular matrices (CC') is conducted to ensure the positive definiteness of the conditional variance-covariance matrix (H_{t+1}). For example, for $q = p = K = 1$ the BEKK model becomes

$$H_{t+1} = CC' + A'H_t A + B'\varepsilon_t \varepsilon_t{}' B \qquad (10)$$

The diagonal BEKK model is a further simplified version of Eq. (10) where A and B are diagonal matrices. It is a restricted version of the diagonal VECH model such that the parameters of the covariance equations for h_{ijt} $(i \neq j)$ are products of the parameters of the variance equations (equations for h_{iit}). To illustrate the diagonal BEKK model, consider the simple GARCH(1,1) model in a bivariate case, where the diagonal BEKK model becomes:

$$H_{t+1} = \begin{bmatrix} c_{11} & c_{12} \\ 0 & c_{22} \end{bmatrix} + \begin{bmatrix} a_{11} & 0 \\ 0 & a_{22} \end{bmatrix}' H_t \begin{bmatrix} a_{11} & 0 \\ 0 & a_{22} \end{bmatrix} + \begin{bmatrix} b_{11} & 0 \\ 0 & b_{22} \end{bmatrix}' \begin{bmatrix} \varepsilon_{1,t}^2 & \varepsilon_{1,t}\varepsilon_{2,t} \\ \varepsilon_{2,t}\varepsilon_{1,t} & \varepsilon_{2,t}^2 \end{bmatrix} \begin{bmatrix} b_{11} & 0 \\ 0 & b_{22} \end{bmatrix} \qquad (11)$$

In our analysis, we employ a GARCH(1,1)-BEKK model in order to generate forecasts of the covariance matrix through time. Then the conditional variance of the hedge fund index at time *t* can be forecasted using the covariance matrix forecast of the underlying style factors in combination with the estimated weights.

An alternative approach would be to estimate the conditional variance of the hedge fund index directly from its historical return observations. This approach reduces the computational effort provided that the historical observations are sufficient for estimation. In this paper, we employ EWMA and GARCH(1,1) models to forecast the hedge fund conditional variance, taking the following forms:

$$h_{t+1} = \sum_{k=0}^{72}(1-\lambda)\lambda^k \epsilon_{t-k}\, \epsilon_{t-k}{}' \tag{12}$$

$$h_{t+1} = \alpha_0 + \alpha_1\epsilon_t^2 + \beta h_t \tag{13}$$

where ϵ_t denotes the innovation of the hedge fund index at time *t* and λ is equal to 0.97.

In summary, we use two approaches to estimate the hedge fund conditional variance: a covariance matrix forecast based on the style factors and a forecast being only based on historical returns of the hedge fund. As mentioned above for the covariance matrix specification based on the style factors we employ three different methods: equally weighted moving average (henceforth F-MA); exponentially weighted moving average (henceforth F-EWMA) and GARCH-BEKK (henceforth F-BEKK) models. To derive forecasts based on historical returns only, we apply two approaches: the EWMA approach (henceforth H-EWMA) and a GARCH(1,1) model (henceforth H-GARCH).

3.4.2 Statistical loss functions for volatility models

To evaluate the out-of-sample performance of the considered forecast models, we adopt a variety of statistical loss functions that have different interpretations and therefore provide a more complete evaluation of the competing models; see e.g. Bollerslev and Ghysels (1996) for more details on the choice of appropriate loss functions. The loss functions considered in our empirical analysis are:

$$MSE_1 = \frac{1}{T}\sum_{i=1}^{T}\left(\sigma_{t+1} - h_{t+1}^{1/2}\right)^2 \tag{14}$$

$$MSE_2 = \frac{1}{T}\sum_{i=1}^{T}(\sigma_{t+1}^2 - h_{t+1})^2 \tag{15}$$

$$MAE_1 = \frac{1}{T}\sum_{i=1}^{T}\left|\sigma_{t+1} - h_{t+1}^{1/2}\right| \tag{16}$$

$$MAE_2 = \frac{1}{T}\sum_{i=1}^{T}|\sigma_{t+1}^2 - h_{t+1}| \tag{17}$$

$$GMLE = \frac{1}{T}\sum_{i=1}^{T}\left(\ln(h_{t+1}) + \frac{\sigma_{t+1}^2}{h_{t+1}}\right) \tag{18}$$

$$LL = \frac{1}{T}\sum_{i=1}^{T}(\ln(\sigma_{t+1}^2) - \ln(h_{t+1}))^2 \tag{19}$$

$$HMSE = \frac{1}{T}\sum_{i=1}^{T}\left(\frac{\sigma_{t+1}^2}{h_{t+1}} - 1\right)^2 \tag{20}$$

Where σ_{t+1}^2 is the realised hedge fund variance at time $t+1$ given by $(y_{t+1} - \bar{y})^2$, y_{t+1} the hedge fund return at time $t+1$ and $\bar{y}$ is the mean hedge fund return during the out-of-sample period. Note that the mean-squared error (MSE) in (14) and (15) and mean absolute error (MAE) in (16) and (17) penalise the errors symmetrically, while logarithmic loss function (LL) in (19) and the heteroscedasticity-adjusted MSE in (20) have the particular features of penalising forecast errors asymmetrically. Further, GMLE in (18) corresponds to the loss implied by a Gaussian quasi-maximum likelihood function. Theoretically, the volatility model that yields the minimum value for a particular loss function is considered to be the best model. However, as pointed out by Bollerslev et al. (1994), the criteria being used to select the best model are not always straightforward when several loss functions are being considered. Once the volatility model that generates the lowest value under a given loss function is said to be the best model, the Diebold-Mariano test (1995) can be applied to test for significant differences between the models. This is a pairwise test of equal predictive ability (henceforth EPA) of two competing models, to find out whether the competing model has the same predictive power as the best model. Under the null hypothesis of equal forecasting accuracy of two competing models, the Diebold-Mariano statistic given by $\frac{\bar{d}}{\sqrt{\hat{V}(\bar{d})}}$ is asymptotically normally distributed. Hereby, $\bar{d}$ denotes the sample mean of the loss difference between the two competing models and $\hat{V}(\bar{d})$ is an estimate of the asymptotic variance of $\bar{d}$.

3.4.3 Empirical results for the considered volatility models and loss functions

Table 3 reports the out-of-sample evaluation of the competing volatility models, according to the statistical loss functions introduced in Section 4.2. The evaluation of the one month ahead volatility forecasts is based on 120 out-of-sample observations and a rolling window of 72 months. As indicated in Table 3, the H-EWMA model performs best with respect to the MSE, MAE and LL loss functions, while the H-GARCH model performs best with respect to the GLME and HMSE loss functions. The F-EWMA is the second best model for the MSE, MAE and LL loss functions, while the H-EWMA and F-BEKK model are the second best for the GMLE and HMSE loss functions, respectively. Although the H-EWMA is not consistently the best model for all of the considered loss functions, we decided to choose it as the benchmark model for the conducted Diebold-Mariano test since it provided the best results for five out of the seven considered loss functions. Table 4 reports the results for the Diebold-Mariano tests where the H-EWMA is tested against the other competing models under the null hypothesis of equally predictive ability. We find that the H-EWMA and H-GARCH models provide the same level of forecast accuracy except for the MAE criterion where the H-EWMA is significantly better. In general, the hedge fund index variance forecasts being based on past returns (H-EWMA and H-GARCH models) seem to outperform the forecasts based on the covariance matrix specification. This is probably due to the difficulties in modelling the conditional variance of the regression residuals, which is attributable to the hedge fund managers' skills. However, also the F-EWMA model that used the style factors to forecast the conditional variance of the hedge fund index returns provides appropriate results and comes second for most of the considered loss functions. To apply this technique might of particular interest for risk managers when newly created funds with a short history of return observations are being considered. In such cases the estimation of conditional variance EWMA or GARCH models may not be feasible due to lack of data.

Model	MSE_1	MSE_2	LL	GMLE	MAE_1	MAE_2	HMSE
F-MA	10.0106	753.5978	9.0383	6.2431	3.4162	23.8276	4.9643
F-EWMA	**9.4507**	**728.2086**	**8.6407**	6.2236	**3.3130**	**23.0510**	5.0119
F-BEKK	11.8436	931.5314	9.3659	6.2715	3.6962	27.1971	**4.1585**
H-EWMA	<u>**9.1445**</u>	<u>**713.4914**</u>	<u>**8.4755**</u>	**6.1584**	<u>**3.2408**</u>	<u>**22.5607**</u>	4.4970
H-GARCH	10.9369	794.8419	9.1468	<u>**6.1323**</u>	3.5280	25.7920	<u>**2.9251**</u>

Table 3: Out-of-sample evaluation of volatility models

Evaluation of one month ahead volatility forecasts based on 120 out-of-sample observations and a rolling window of 72 months. The minimum value for each loss function is in bold font and underlined, and the second smallest value is just in bold font.

Model	MSE_1	MSE_2	LL	GMLE	MAE_1	MAE_2	HMSE
F-MA	-2.06*	-2.19*	-1.99*	-1.49	-1.85	-1.82	-0.63
	(0.040)	(0.028)	(0.047)	(0.136)	(0.065)	(0.069)	(0.529)
F-EWMA	-2.10*	-2.00*	-1.51	-2.67**	-2.32*	-2.24*	-1.80
	(0.035)	(0.045)	(0.130)	(0.008)	(0.020)	(0.025)	(0.071)
F-BEKK	-1.97*	-1.39	-2.95**	-1.92	-2.43*	-2.08*	1.27
	(0.049)	(0.166)	(0.003)	(0.054)	(0.015)	(0.038)	(0.205)
H-GARCH	-1.86	-0.96	-1.89	0.30	-2.00*	-2.10*	1.37
	(0.063)	(0.338)	(0.059)	(0.768)	(0.046)	(0.036)	(0.171)

Table 4: Diebold-Mariano test (benchmark model: H-EWMA)

*This table shows the statistics and corresponding p-values (in parentheses) for the conducted Diebold-Mariano test. Other competing models are tested against H-EMWA model under the null hypothesis of equal predictive ability. * and ** represent the rejection of the null hypothesis at 5% and 1% respectively.*

3.4.4 Value-at-Risk framework and loss functions

The derived forecasts for the volatility of the considered hedge fund index can also be used as an input for a Value-at-Risk (VaR) analysis. The one-month ahead hedge fund index VaR at the α% confidence level of model i can then be denoted by:

$$VaR^i(\alpha) = u_{t+1}^i + \Phi(\alpha)\sqrt{h_{t+1}^i} \tag{21}$$

where u_{t+1} and h_{t+1} are the forecasted conditional mean and variance estimated at time t with model i; $\Phi(\,)$ is a cumulative distribution function, which is often assumed to be the Gaussian or Student t distribution. Using the style factor covariance matrix specification, u_{t+1} is given by $w_t \times \mu_t{}'$, where $w_t = [w_{1,t}, w_{2,t}, \dots, w_{n,t}]$ and $\mu_t = [\mu_{1,t}, \mu_{2,t}, \dots, \mu_{n,t}]$ are the vectors containing the weights and the conditional means of the significant style factors at time t. On the other hand, when only considering historical hedge fund index returns, u_{t+1} is simply equal to u_t, the conditional mean of the hedge fund index at time t. In this paper, using the techniques described in the previous sections, we use a normal and Student t distribution in order to estimate the hedge fund VaR at the 1% and 5% confidence level.

Choosing an appropriate distributional assumption is vital for accuracy of VaR estimation. As discussed by Lopez and Walter (2000), for a portfolio of foreign currency, the performance of VaR models depend more on the assumption of return distribution than on the conditional volatility forecasting. Generally, the normal distribution is the most commonly used distribution in VaR estimation. However, empirical distributions of hedge fund returns are usually fat-tailed, i.e., great losses have a higher likelihood than suggested by the Gaussian distribution. Therefore, we also employ a Student t distribution for the VaR estimation of the considered index. The t-distribution has the form

$$f(\epsilon_t) = \frac{\Gamma\left(\frac{v+1}{2}\right)}{\sqrt{\pi}\Gamma\left(\frac{v}{2}\right)}[h_t(v-2)]^{-0.5}[1 + \frac{\epsilon_t^2}{h_t(v-2)}]^{-\frac{v+1}{2}} \tag{22}$$

where $v > 2$ is the degree of freedom that affects the tail thickness of the distribution, and $\Gamma(\,)$ denotes the Gamma function. Using the residuals of the considered volatility models during the in-sample period, we estimate the degree of freedom parameter as $v = 6$. In the following we compare the VaR results for the considered models in combination with the assumption of a Gaussian or Student t distribution for the hedge fund index returns.

When investigating the appropriateness of different VaR models, another common approach is to estimate VaR based on historical simulation. In this case the VaR is calculated from the empirical distribution of historical returns only, not assuming any parametric model for the returns or volatility. Historical simulation is a particularly popular approach in the industry such that we decided to compare our results for VaR quantification based on the considered parametric models also to a nonparametric approach. Hereby, we use a rolling window of the past 100 return observations in order to construct the nonparametric empirical distribution of hedge fund returns and subsequently estimate the VaR.

To evaluate the different VaR models in their ability to forecast extreme losses at a specified confidence level, we thus employ the historical simulation approach, parametric models for the index returns as well as models based on the estimated style factors in order to determine the accuracy of the VaR forecasts. This is done by using a three-step procedure initially proposed by Christoffersen (1998) and Christoffersen and Diebold (2000).

The first step is to evaluate the VaR estimation based on the unconditional coverage. Here the null hypothesis is that the unconditional coverage $\hat{\alpha} = x/T$ is equal to p, with x being the number of VaR exceptions at a given confidence level p and T the total number of VaR forecasts in the out-of-sample period. Then the test statistics is given by $LR_{uc} = -2\ln[\frac{p^x(1-p)^{T-x}}{\alpha^x(1-\alpha)^{T-x}}]$, follows a $\chi^2(1)$ distribution. The second step is to test for independence of VaR exceptions in order to examine whether exceptions are spread evenly through the period used for backtesting. Then the LR statistic for the test of independence is $LR_{ind} = -2\ln\left[\frac{(1-\pi_2)^{(T_{00}+T_{00})}(1-\pi_2)^{(T_{01}+T_{11})}}{(1-\pi_{01})^{T_{00}}\pi_{01}{}^{T_{01}}(1-\pi_{11})^{T_{10}}\pi_{11}{}^{T_{11}}}\right]$, following a $\chi^2(1)$ distribution, where $\pi_{ij} = \Pr\{I_t = i | I_{t-1} = j\}(i, j = 0,1)$, I_t to be 1 if VaR is exceeded and to 0 otherwise, $\pi_{01} = \frac{T_{01}}{T_{01}+T_{00}}$, $\pi_{11} = \frac{T_{11}}{T_{10}+T_{11}}$, $\pi_2 = \frac{T_{01}+T_{11}}{T_{01}+T_{00}+T_{10}+T_{11}}$, T_{ij} is the number of observations that a i followed by a j in the I_t series. Note that when hedge fund returns exhibit heteroskedasticity, evaluation of VaR models based on the test for unconditional coverage only may not be sufficient, because a VaR model providing an appropriate unconditional coverage may still yield an incorrect conditional coverage. Thus, the third step is to test the conditional coverage by using the statistics $LR_{cc} = LR_{ind} + LR_{uc}$ that follows a $\chi^2(2)$ distribution. As pointed out by Christoffersen (1998) and Christoffersen and Diebold (2000), a model that passes both the unconditional and conditional coverage test can be considered as adequate for VaR estimation.

3.4.5 Empirical results for Value-at-Risk quantification

Table 5 presents the unconditional coverage (UC), i.e., the percentage of exceptions, the LR statistics for unconditional coverage test (LR_{uc}), the independence test (LR_{ind}) and the conditional coverage test (LR_{cc}) for both 95% and 99% VaR one month ahead forecasts. If the fraction of empirically observed exceptions is greater than the theoretical number of exceptions at the 1% and 5% significance level, it indicates that the model is inadequate. For 99% VaR, the unconditional and conditional coverage tests reject most parametric VaR models under normal distribution assumption and historical simulation approach, while for 95% VaR, all the models pass the three tests. Further, all the models pass the independence test, showing that VaR exceptions seem do not cluster abnormally. From the table, the results show that VaR models based on t-distribution assumption clearly outperform those based on a normal distribution assumption. Recall that for volatility forecasting, the H-EWMA and H-GARCH models performed best for most of the considered loss functions. However, for the conducted VaR analysis we find that with respect to unconditional coverage during the out-of-sample period in particular the H-GARCH and F-BEKK models combined with the assumption of a t-distribution for the returns yield the best results. For these models, the number of exceptions is usually equal or even lower than the corresponding probability level while all three LR tests fail to reject the null hypothesis of adequate model specification. In contrast to these parametric methods, the nonparametric historical simulation approach performs rather poorly. In Section 4.3, the empirical results indicated that the H-GARCH and F-BEKK models rank among the two best models with respect to volatility forecasting for the HMSE loss function. This function assigns higher weight to an incorrect low variance forecast when the actual realised variance is high. Hence, the volatility model that closely captures the tail features of the distribution should perform best for HMSE loss criterion. Moreover, we find that the Student t distribution is significantly more suitable than the normal distribution to capture the fat-tailed distribution of hedge fund returns. Taking into account that our out-of-sample data covers the Global Financial Crisis period during the years 2008 and 2009 when also hedge funds suffered significant losses, the empirical results indicate that the performance of VaR models is particularly dominated by its ability to capture the tail of the return distribution.

In Figure 3 and 4 we also provide a plot of the actual hedge fund returns and the 95% and 99% VaR estimates based on the considered H-GARCH-t and F-BEKK-t models. Both models react quite significantly to the change in market condition during the Global Financial Crisis. However, the H-GARCH-t model seems to respond even quicker to the changes than the F-BEKK-t model.

Model	95% VaR UC	LR_{uc}	LR_{ind}	LR_{cc}	99% VaR UC	LR_{uc}	LR_{ind}	LR_{cc}
H-EWMA-n	7.50	1.38	0.16	1.54	4.17	6.79*	1.82	8.62*
H-EWMA-t	**5.00**	0.00	1.18	1.18	1.67	0.45	0.07	0.52
H-GARCH-n	7.50	1.38	0.16	1.54	**0.83**	0.04	0.02	0.05
H-GARCH-t	**5.00**	0.00	1.18	1.18	**0.83**	0.04	0.02	0.05
F-MA-n	8.33	2.36	0.04	2.39	5.00	9.91*	1.18	11.09*
F-MA-t	5.83	0.17	0.71	0.88	**0.83**	0.04	0.02	0.05
F-EWMA-n	9.17	3.56	0.00	3.56	5.00	9.91*	1.18	11.09*
F-EWMA-t	**5.00**	0.00	1.18	1.18	1.67	0.45	0.07	0.52
F-BEKK-n	9.17	3.56	0.00	3.56	3.33	4.10*	2.71	6.81*
F-BEKK-t	**5.00**	0.00	1.18	1.18	**0.83**	0.04	0.02	0.05
Historical Simulation	6.67	0.64	0.38	1.02	3.33	4.10*	2.71	6.81*

Table 5: VaR out-of-sample evaluation: 95% and 99% VaR

*Unconditional coverage (UC), i.e., the percentage of exceptions as well as the LR statistics for unconditional coverage test (LR_{uc}), independence test (LR_{ind}) and conditional coverage test (LR_{cc}) for both 95% and 99% VaR estimates. * indicates rejection of the null hypothesis at the 5% significance level. The minimum value of the unconditional coverage is highlighted in bold letters.*

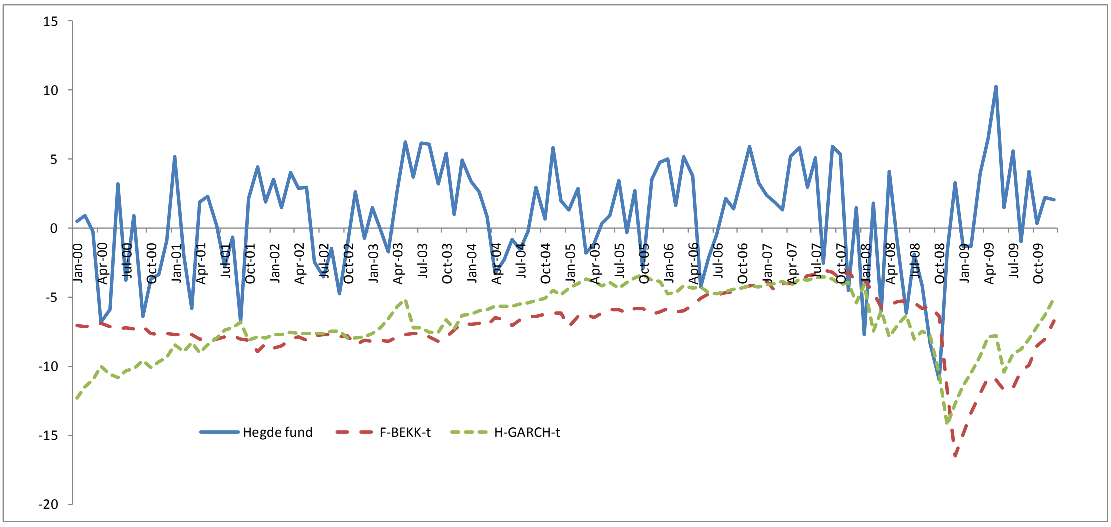

Figure 3: Returns of HFRI Emerging Market-Asia exclude Japan index and 95% VaR forecasts for the considered out of sample period January 2000 to December 2009. The graph compares the computed 95% VaR for the H-GARCH-t model (green dotted line) and F-BEKK-t model (red dotted line).

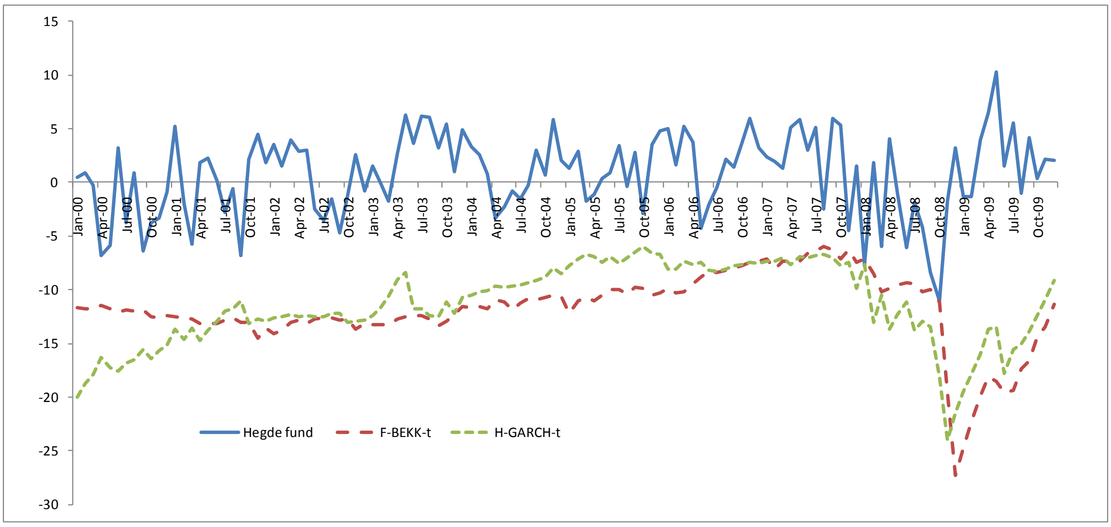

Figure 4: Returns of HFRI Emerging Market-Asia exclude Japan index and 99% VaR forecasts for the considered out of sample period January 2000 to December 2009. The graph compares the computed 99% VaR for the H-GARCH-t model (green dotted line) and F-BEKK-t model (red dotted line).

3.4.6 Magnitude of VaR exceptions

Generally, in the academic literature and practice, most evaluations of VaR estimates are based on the frequency of the VaR exceptions. However, also the magnitude of VaR exceptions is of particular interest to risk managers and financial institutions. This is even of higher importance when risk management practices focus also on expected shortfall instead of VaR only. In this section, we employ a hypothesis test proposed by Berkowitz (2001) focussing on the expected loss in comparison to the actually observed loss when the VaR is exceeded.

A difficulty in evaluating the performance of VaR models is the small number of observed violations. For example, a 99% VaR should provide only approximately one violation in every 100 observations if it is correctly specified. Therefore, as stressed by Kupiec (1995), a large sample size is required to verify the accuracy of a VaR model. An alternative to focusing on the low frequency of VaR exceptions only, is to apply the Rosenblatt (1952) transform to the predicted return distribution

$$\hat{F}(y_t) = \int_{-\infty}^{y_t} \hat{f}(x)dx \tag{23}$$

where y_t is the realised return at time *t* and $\hat{f}(x)$ is the loss density function generated by the model used for forecasting. Rosenblatt shows that if the distribution is correctly specified this will transform the observed returns into a series of *iid* random variables. Thus, the accuracy of the VaR model can be tested under the null hypothesis that the probability integral transforms $\hat{F}$ are *iid* and distributed uniformly on [0,1]. As suggested by Crnkovic and Drachman (1996), the Kuiper statistic based on the distance between the empirical and the theoretical cumulative distribution function of the uniform distribution can be used in order to test for uniformity. However, a small sample size is not suitable for this test since a large number of points is required to calculate the distance. Therefore, instead of testing the uniformity, Berkowitz (2001) transforms $\hat{F}$ into standard normal series and tests the accuracy of VaR models by constructing likelihood-ratio (LR) tests. Focusing on the magnitude of the VaR exceptions, Berkowitz proposes a LR test based on the censored likelihood, such that the shape of realised lower tail is compared with the forecasted lower tail so as to determine whether the observed VaR exceptions are in line with the underlying VaR model. Moreover, Berkowitz points out that the proposed likelihood-ratio test is well suited for sample sizes as small as 100. Let $z_t = \Phi^{-1}\left(\hat{\mathrm{F}}(\mathrm{y_t})\right)$ denote the inverse of the standard normal distribution function of $\hat{\mathrm{F}}(\mathrm{y_t})$, $VaR = \Phi^{-1}(\alpha)$ denote the cut-off point, i.e., VaR=-1.96 for 5% lower tail of standard normal distribution and z_t^* denote the further transformation of z_t given by

$$z_t^* = \begin{cases} VaR \; if \; z_t \geq VaR \\ z_t \quad if \; z_t < VaR \end{cases}.$$

Then the log-likelihood function can be expressed as

$$L = \sum_{z_t^*=VaR} Log\left(1-\Phi\left(\frac{\mathrm{VaR}-\mu}{\sigma}\right)\right) + \sum_{z_t^*<VaR}\left(-\frac{1}{2}log(2\pi\sigma^2) - \frac{(z_t^*-\mu)^2}{2\sigma}\right) \qquad (24)$$

where μ and σ denote the mean and standard deviation of the transformed standard normal series z_t. Under the null hypothesis that $\mu = 0$ and $\sigma = 1$, the likelihood-ratio test statistic is given by $LR = 2\big(L(\mu,\sigma) - L(0,1)\big)$, which is approximately $\chi^2(2)$ distributed.

The test statistics for the LR test are reported in Table 6. Interestingly, for none of the models the null hypothesis that $\mu = 0$ and $\sigma = 1$ can be rejected, indicating that the mean and the variance of the observed violations is consistent with those implied by the considered VaR models. That is, all the models appear to perform well regarding to the magnitude of VaR exceptions.

Model	95% VaR	99% VaR
H-EWMA-n	-1.18	0.68
H-EWMA-t	1.11	0.13
H-GARCH-n	-1.04	-0.20
H-GARCH-t	2.69	0.76
F-MA-n	0.52	2.15
F-MA-t	1.97	0.48
F-EWMA-n	-0.74	1.36
F-EWMA-t	0.95	0.10
F-BEKK-n	-0.11	0.50
F-BEKK-t	2.19	0.77
Historical Simulation	-0.22	-0.38

Table 6 Magnitude of VaR exceptions

*LR statistics for magnitude of VaR exception test for both 95% and 99% VaR. * indicates the rejection of the null hypothesis at 5% significance level.*

3.5 Conclusion

In this paper, we identify style factors for Asia-focused hedge funds represented by the HFRI Emerging Market-Asia exclude Japan index. Hereby, we make use of the style analysis framework initially suggested by Agarwal and Naik (2000) and Dor et al. (2003). Furthermore, we employ a two-step procedure proposed by Lobosco and Dibartolomeo (1997) to test for the significance of the considered style factors. A rolling window style analysis is performed to provide further insights into the dynamic structure of style factor weights and risk exposures. This is one of the first empirical studies applying these techniques with particular focus on the Asian hedge fund industry.

The empirical results show that the most significant equity factors relating to the HFRI Emerging Market-Asia exclude Japan index are emerging equity markets, especially emerging markets in Asia. The two factors representing global and Asian emerging markets together account for a weight of approximately 45% on average. The risk exposures are consistent with the investment objective of the hedge fund strategy. With respect to the fund's exposure to bond markets, we find that Asia-focused hedge funds indicate positive exposures to cash and high credit rating bonds but negative exposures to world government and emerging market bonds. In general, these fixed income factors account for a weight of 45%. The rolling window style analysis captures the hedge fund managers' style drift in responding to dynamic trading and changing market situations. For both static and rolling period style analysis, our model provides a high explanatory power for returns of the hedge fund index.

We further conduct an extensive analysis with respect to the ability of the models to provide appropriate forecasts for volatility and Value-at-Risk of the index. Hereby, we use the identified factors and factor weights of the rolling window style analysis in combination with a multivariate GARCH, moving average or exponentially weighted moving average (EWMA) model. The results are also compared to an approach that applies a univariate EWMA and GARCH model directly to the returns of the index. With respect to volatility forecasting the models are compared based on a set of different loss functions. We find that none of the models performs best for all of the considered loss functions or significantly outperforms all of the other models. However, overall the best results are obtained for three of the considered models: the EWMA and GARCH model using the actually observed returns of the hedge fund index as well as a model using the estimated style factor weights in combination with an EWMA scheme for the volatility.

In a second step, based on hypothesis tests for unconditional and conditional coverage, we further evaluate the performance of the considered models with respect to VaR estimation. Hereby, also different assumptions for the return distribution are applied. Finally, the magnitude of the observed VaR exceptions is compared to those implied by the estimated VaR models. Our results indicate that the accuracy of the VaR models is dominated by its ability to capture the tail distribution of the hedge fund returns. Moreover, the performance of the models in VaR prediction seems to be dependent rather on the distributional assumption for the returns than on the chosen approach for volatility modelling: all models assuming a Student t distribution for the returns of the hedge fund index are significantly better than their counterparts assuming a Gaussian distribution. Overall, the best models for VaR estimation are a GARCH BEKK model based on the underlying style factors and a GARCH model that is based on the hedge fund returns only. Our findings further suggest that, in VaR forecasting, all parametric models outperform a simple historical simulation approach being purely based on past return observations. Finally, all of the considered VaR models perform reasonably well in forecasting the magnitude of the loss conditional on a VaR exception.

Overall, our findings suggest that style analysis in combination with an appropriate parametric model for the identified factors provides an appropriate quantification of the risk for the considered Asian hedge fund index. We also find that multivariate models based on identified style factors and style weights significantly outperform a historical simulation approach with respect to volatility or VaR forecasting. On the other hand, our analysis indicates that they do not necessarily outperform simpler models like a univariate GARCH or EWMA model being directly applied to the hedge fund return series. However, they provide important insights on the exposures and investment style of a fund and indicate how fund returns can be replicated by observable market factors. In a time-varying setting style analysis also provides information on how the weights of the different style factors potentially change through time as a reaction to different market conditions. Finally, style analysis might be useful for risk management when only a short period of observations is available for the fund itself while the identified style factors provide a much longer history that can be employed for estimating VaR or other risk measures. Therefore, we believe that style analysis approach should also be of particular help when individual hedge funds with a short track record are analysed and the use of hedge fund returns only for risk analysis will fail due to the lack of historical data. This issue should be thoroughly investigated in future research.

References

Ackermann, C., McEnally, R., Ravenscraft, D., 1999. The performance of hedge funds: risk, return, and incentives. Journal of Finance 54(3), 833-874.

Agarwal, V., Naik, N.Y., 2000. Generalised style analysis of hedge funds. Journal of Asset Management 1, 93-109.

Agarwal, V., Naik, N.Y., 2004. Risks and portfolio decisions involving hedge funds. Review of Financial Studies 17(1), 63-69.

Amenc, N., Martellini, L., Vaissié, M., 2002. Benefits and risks of alternative investment strategies. Journal of Asset Management 4(2), 96-118.

Bauwens, L., Laurent, S., Rombouts, J.V.K., 2006. Multivariate GARCH models: a survey. Journal of Applied Econometrics 21, 79-109.

Berkowitz, J., 2001. Testing density forecasts, with applications to risk management.

Journal of Business and Economic Statistics 19(4), 465-474.

Bollen, N.P.B., Robert, E.W., 2008. Hedge fund risk dynamics: Implications for performance appraisal. Journal of Finance 64, 985-1035.

Bollerslev, T., Engle, R.F., Nelson, D.B., 1994. ARCH models, in Engle, R.F. and McFadden, D., eds. The Handbook of Econometrics Volume 4, 2959-3038. Amsterdam: North-Holland.

Bollerslev, T., Ghysels, E., 1996. Periodic autoregressive conditional heteroskedasticity. Journal of Business and Economic Statistics 14, 139-184.

Brown, S.J., Goetzmann, W.N., Ibbotson, R.G., 1999. Offshore hedge funds: survival and performance 1989-1995. Journal of Business 72(1), 91-118.

Christoffersen, P.F., 1998. Evaluating interval forecasts. International Economic Review 39, 841-862.

Christoffersen, P.F., Diebold, F., 2000. How relevant is volatility forecasting for financial risk management. Review of Economics and Statistics 82, 12–22.

Crnkovic, C., Drachman, J., 1996. A universal tool to discriminate among risk measurement techniques. Risk 9, 138–143.

Diebold, F.X., Mariano, R., 1995. Comparing predictive accuracy. Journal of Business and Economic Statistics 13, 253-264.

Dor, A.B., Jagannathan, R., Meier, I., 2003. Understanding mutual funds and hedge fund styles using return-based style analysis. Journal of Investment Management 1(1), 97-137.

Duffie, D., Pan, J., 1997. An overview of value at risk. Journal of Derivatives (Spring), 7-49.

Engle, R.F., Kroner, K.F., 1995. Multivariate simultaneous generalised GARCH. Econometric Theory 11, 122-150.

Fung, W., Hsieh, D.A., 1997. Empirical characteristics of dynamic trading strategies: the case of hedge funds. The Review of Financial Studies 10(2), 275-302.

Fung, W., Hsieh, D.A., 2000. Performance characteristics of hedge funds and commodity funds: natural vs. spurious biases. Journal of Quantitative and Financial Analysis 35(3), 291-307.

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

Fung, W., Hsieh, D.A., 2002. The Risk in fixed-income hedge fund styles. Journal of Fixed Income 12(2), 16-27.

Fung, W., Hsieh, D.A., 2003. The risk in equity long/short hedge funds. Working Paper, London Business School and Duke University.

Fung, W., Hsieh, D.A., 2004. Hedge fund benchmarks: a risk based approach. Financial Analyst Journal 60, 65-80.

Gibson, M.S., Boyer, B.H., 1998. Evaluating forecasts of correlation using option pricing. Journal of Derivatives, Winter, 18-38.

Glosten, L.R., Jagannathan, R., 1994. A contingent claim approach to performance evaluation. Journal of Empirical Finance 1, 133-160.

Hull, J., White, A., 1998. Value at risk when daily changes in market variables are not normally distributed. Journal of Derivatives 5(3), 9-19.

Jorion, P., 2001. Value-at-risk: the new benchmark for controlling market risk. Chicago: McGraw-Hill.

JP Morgan, 1996. RiskMetrics technical document.

Kupiec, P.H., 1995. Techniques for verifying the accuracy of risk measurement models. Journal of Derivatives, winter, 73-84.

Liang, B., 1999. On the performance of hedge funds. Financial Analysts Journal 55(4), 72-85.

Lobosco, A., Dibartolomeo, D., 1997. Approximating the confidence intervals for Sharpe style weights. Financial Analysts Journal 53(4), 80-85.

Lopez, J.A., Walter, C.A., 2000. Evaluating covariance matrix forecasts in a Value-at-Risk framework. Federal Reserve Bank of San Francisco, working paper.

Loudon, G., Okunev, J., White, D., 2006. Hegde fund risk factors and the value at risk of fixed income trading strategies. Journal of Fixed Income 16(2), 46-61.

Mitchell, M., Pulvino, T., 2001. Characteristics of risk and return in risk arbitrage. Journal of Finance 56 (6), 2135 – 2175.

Rosenblatt, M., 1952. Remarks on a multivariate transformation. Annals of

Mathematical Statistics 23, 470–472.

Sharpe, W.F., 1992. Asset allocation: management style and performance measurement. Journal of Portfolio Management 18(2), 7-19.

Teo, M., 2009. The Geography of hedge funds. The Review of Financial Studies 22(9), 3531-3561

Appendix A: Description of considered style factors for the style analysis of Asian-focused hedge funds

Factors	Description
3-month Treasury Bill	Monthly yield on U.S. Treasury securities at 3-month constant maturity.
Bank of America Merrill Lynch US High Bond	Index for US high yield bonds (below investment grade).
CGBI Broad Investment Grade	Index for US investment grade bonds. The index includes treasuries, agency debt, corporate, non-corporate credit, mortgage-backed securities, and asset-backed securities.
CITI World Government Bond	A market capitalization weighted bond index consisting of the government bond markets of the multiple countries.
S&P 500	The S&P 500 is a free-float capitalization-weighted index of the prices of 500 large-cap common stocks actively traded in the United States
MSCI Europe	The market capitalization weighted index measures the equity market performance of the 16 developed markets in Europe.
MSCI Japan	MSCI Japan measures the performance of the Japanese equity market.
MSCI emerging markets excluding Asia	The market capitalization weighted index measures the equity market performance of the emerging markets excluding Asia. There are 13 emerging market countries included in this index.
MSCI emerging markets Asia	The market capitalization weighted index measures the equity market performance of the emerging markets in Asia. The index consists of the following emerging market countries: China, India, Indonesia, Korea, Malaysia, Philippines, Taiwan, and Thailand.
MSCI Pacific excluding Japan	The market capitalization weighted index measures the equity market performance of the developed markets in the Asia Pacific region excluding Japan. The index consists of the following developed market countries: Australia, Hong Kong, New Zealand and Singapore.
Trend-Following Risk Factors	The returns of private trend following strategy (PTFS) lookback straddles in bond, currency, short term interest rate, commodity and stock.

4. Agency Theory and Financial Planning Practice

Geoffrey Kingston (contribution: 50%) and Haijie Weng (contribution: 50%)

This paper has been presented at:

- Financial Risk Day, Macquarie University, Sydney, March 2012

This paper has been accepted for publication in the forthcoming issue of *Australian Economic Review.*

Abstract

We extend an influential contribution to the literature on agency theory and then use this extension, along with other theoretical contributions, to shed light on agency problems affecting funds management and financial planning in Australia. The case for pure fee for service in actively managed funds and plans turns out to be weak. The amount of money exposed to risk by an active manager should be less than the entire investible wealth of the client, especially in the case of investors on the cusp of retirement. Asset-based fees on actively managed funds should include a fulcrum component.

4.1 Introduction

This paper compares and contrasts mainstream agency theory with financial planning practice in Australia. It appears to be the first attempt to do so. It extends an influential mainstream contribution to the literature on agency theory and then uses this extension, in conjunction with other theoretical contributions, to shed light on actual contracts between investors, financial planners, licensees and product providers. The case for pure fee for service in actively managed funds and plans turns out to be weak, at least when the manager and the provider are treated as a consolidated entity, that being a reasonable first-order approximation in the Australian case. The amount of money exposed to risk by an active manager should be less than the entire investible wealth of the client, especially in the case of recently-retired investors. Asset-based fees on actively managed funds should be symmetric in gains and losses relative to a benchmark, contrary to current mainstream practice.

The background to this paper is the continued growth in financial planning and funds management as the baby boomers move towards retirement. At present only 13 per cent of Australians are at least 65 years of age, and 7 out of 10 retired households rely principally on the Age Pension. Only 15 to 20 per cent of Australians have received financial advice from planners at some point during their lives. However, the 65-plus population is projected to hit 23 per cent of the population by 2050 and self-funded retirements are becoming more widespread, so the number of Australians receiving advice from financial planners should rise. Funds under advice in Australia stand at $519 billion (Rainmaker Group 2013). There are 760 advisory groups, 8,300 financial planning practices, and 18,200 financial planners. There has been ongoing vertical integration within the industry, as small practices enter into 'sponsorship' relationships, primarily with the big-four banks and the major insurance companies. At least 80 per cent of financial planners are sponsored. On the other hand, self-managed superannuation funds account for 31 per cent of total superannuation assets of $1.4 trillion (Australian Prudential Regulation Authority 2012). This alone suggests that the market for advice is contestable, as a consequence of a substantial competitive fringe of comparatively self-reliant investors.

The financial planning industry has come under scrutiny in the wake of the global financial crisis (2008-10) and the collapse of Storm Financial in 2009. The 2007 budget had abolished taxes on the earnings of superannuation funds in drawdown mode and allowed higher personal contributions. As a consequence there were strong inflows into superannuation during the 2007 financial year – decisions which worked out badly for many investors in the wake of the global financial crisis. The year 2009 saw two official inquiries into industry practices. The Ripoll inquiry reported in 2009 and the Cooper inquiry reported in 2010. These served as inputs to the government's Future of Financial Advice (FoFA) and MySuper reforms. Questions raised in the Australian debate on financial planning include these:

1. Should fee-for-service supplant asset-based fees?
2. Should commissions from product providers to planners be banned?
3. Do recommended asset allocations tend to be too risky for clients on the cusp of retirement?
4. Do financial plans tend to be 'cookie cutter' ones rather than customised to the particular circumstances of clients?
5. Do typical fee structures encourage 'closet indexing' by fund managers?
6. Has there been inadequate disclosure of dollar (rather than percentage) amounts charged in fees?

We arrive at affirmative answers to all these questions except the first one.

Section 2 sheds light on the first, third, fourth and fifth questions by extending the model of Dybvig et al. (2010). That model is a direct descendent of the classic agency model due to Ross (1973). It derives optimal contracts in financial plans when both the investor-principal and the planner-agent have log utility. It does not distinguish between planners and managers, and this is useful to the extent that the managed funds industry shows strong vertical integration, as is the case in Australia. Efficient fee structures always involve asset fees, and generally tie a component of remuneration to portfolio performance relative to a suitable passive benchmark, to discourage closet indexing.

We introduce generalised log utility[1] into the setup of Dybvig et al. Generalised log utility has the realistic implication that relative risk aversion is a declining function of wealth, unlike its log, quadratic, power and exponential competitors.[2] It is the simplest way to capture habit-dependent utility whereby a retiree is concerned to prevent her living standard falling below some pre-determined level, and is consistent with a desire to ‘keep up with the Joneses’. It can rationalise conservative asset allocations on the cusp of retirement whereas simple log utility generates aggressive allocations. Generalised log utility captures the concern of some investors with preventing shortfalls in wealth below some subjective reference level, and the present value of protected future consumption is the natural interpretation of that level. In this way it sheds light on the question of excessively risky allocations for people on the cusp of retirement.

The proposed model is of a contract between an investor and a unified fund manager/financial planner entity and not a standalone financial planner. It is important to emphasize that this assumption is realistic in the Australian context because of strong vertical integration in the financial planning/fund managing industry as described in Section 3. Approximately 70-80% of financial planning practices are owned by large institutions which also manage funds.

Section 3 examines Australian industry practice. Unsurprisingly, typical contracts set out in actual Statements of Advice and Product Disclosure Statements turn to be much richer than could be captured by a single theory. Accordingly, Section 3 draws informally on the results of Stoughton et al. (2011), Bateman et al. (2007) and Grossman and Stiglitz (1976), in addition to the formal theory of Section 2. These contributions shed light on the agency problems raised by intermediated investment management, multiple time periods, and general equilibrium. Put another way, Section 2 does not shed light on all six questions, and therefore needs to be beefed up by other theories, at least informally.

Take our second question, on commissions. It presupposes a three-way split between investors, advisers and investment managers (notwithstanding the considerable vertical integration in Australia.) Stoughton et al. (2011) do introduce such a split (in contrast to Dybvig et al.) and it sheds light on commissions. Investors can engage an adviser, or pay a fixed cost to access actively managed funds without intermediation by advisers, analogous to Australia's self-managed funds. Investors divide into sophisticated or unsophisticated ones, depending on whether they anticipate equilibrium outcomes in the financial planning industry and are impervious to promotional material. Commissions from managers to advisers can take the form of cash or soft-dollar compensation such as conferences in resort locations. All this helps explain Australian practice.

4.2 Agency Theory

This section extends the theory of fee structures for actively managed funds that mitigate agency problems when both the principal and the agent have generalised-log utility functions. One new result is that efficient fees include a fixed component reflecting the agent's protected consumption. This generates a new rationale for a flat component of fees, analogous to fee-for-service. But the optimal contract retains roles for asset-based fees. Another new result is that, from both an investor and manager standpoint, the participation decision is not all-or-nothing; the amount placed with the active manager is equal to the investor's wealth less the present value of the total protected consumption of the investor and the manager. Remaining wealth is allocated to safe assets. In practice, this suggests that an investor should place part of her retirement money in term deposits rather than allow all of it to be actively managed.

Our setup retains some features of Dybvig et al. (2010). Notably, the asset-based fees derived there are retained here as a component of the overall fee structure, including a symmetrical asset-based fee for performance relative to a passive benchmark.

Dybvig et al. (2010) derive an optimal contract for portfolio managers using an agency theory model. In this model, it is assumed that the signals about future prices can be influenced by a portfolio manager's efforts. An investor is then required to find an optimal contract to provide the manager incentive to exert effort and to use the signal in the investor's interest. Further, Dybvig et al. (2010) assume both investor and manager have a log utility and consider three scenarios: 1) the first-best world where the manager's effort is observable; 2) the second-best world where the manager's signal is observable but the effort is not observable and 3) the third-best world where neither the effort nor the signal is observable. The authors have shown that in the first best world the optimal contract is a fixed proportion of the end-of-period assets under management independently of the signal. In the second-best world, the optimal contract is a fixed fraction of the end-of-period assets under management plus a bonus in proportion to the excess portfolio return over a benchmark. The bonus component is to give the manager incentives to exert efforts. In the third-best world, the manager's fee is no longer a liner combination of the portfolio and a benchmark performance, as it is in the second-best world, but contains an additional nonzero payoff conditional on the signal. The excess return component in the fee structure will still provide the manager the incentive but will tend to make the manager overly conservative when making investment decisions.

We extend the utility function for both the investor and manager in Dybvig et al (2010) to a generalised log utility. A constant is added in the log utility which we consider as the protected consumptions of the investor and manager.

Dybvig et al. (2010) consider three optimization problems corresponding to increasingly severe agency problems. In the first-best case, agency problems are absent. In the second-best case the manager reveals truthfully the observed signal to the investor but has private information about her effort level. In the third-best case the adverse-selection problem and the moral-hazard problem are both present. It is the second-best case which yields the most interesting results, so we disregard the third-best case, and comment only briefly on the first-best case. Our main result is this:

4.2.1 Proposition

The optimal contract between an investor and an active manager whose effort level cannot be verified by the investor first carves out the total protected wealth of the investor and the agent. It then subjects the remaining wealth of the investor to a fee structure with a flat component and two asset-based components. One asset fee is a standard proportional fee on fund earnings. The other is a symmetrical fulcrum-style performance fee:

$$\phi(s,\omega) = \underline{C}_m + \frac{(w_0 - \underline{w})\lambda_R}{1+\lambda_R}[R^P + \left(\frac{\lambda_\varepsilon}{\varepsilon\lambda_R}\right)(R^P - R^B)]\,. \tag{1}$$

On the left-hand side of equation (1), $\phi(s,\omega)$ is the fee paid by an investor when the manager's unobserved effort ε $(0 \le \varepsilon \le 1)$ generates a private signal $s \in S$ about future returns, and the state of the world is $\omega \in \Omega$. On the right-hand side, $\underline{C}_m$ is the protected consumption of the manager, w_0 is investible wealth, $\underline{w}$ is the present value of the total protected consumption of the investor and the manager, λ_R is a Lagrange multiplier on a participation constraint, R^P is the return to the actively managed portfolio, λ_ε is a Lagrange multiplier on an incentive-compatibility constraint, and R^B is the return to a passively-managed (zero-effort) benchmark portfolio.

4.2.2 Proof

See Appendix 1

4.3 Financial Planning Practice

This section compares and contrasts our claimed optimal structure (1) with actual fee structures and associated advice documented by the Financial Planning Association and Morningstar. Consistent with (1), actual fees contain both flat and proportional components. On the other hand (and as one might expect) there appears to be little or no advice to the effect that investors set aside part of their wealth in safe assets. Rather, the plan discussed here recommends that investors elevate their pre-existing exposures to growth assets. Moreover, there appears to be little or no use of fulcrum fees by either planners or fund managers. Performance fees exist and are mostly set in practice by managers rather than planners. They are not of the fulcrum variety, as they are neither symmetrical nor based on the natural benchmark, i.e. the best passively managed allocation for investors with age and wealth comparable to that of the actual client. Consistent with these gaps between prescriptive theory and actual practice, there is evidence of the closet indexing that fulcrum fees would discourage.

The Financial Planning Association is the dominant industry association for Australian financial planners. Roughly two thirds of licensed planners belong to it. The FPA has promulgated an 'Example' Statement of Advice on behalf of a hypothetical couple aged 57 and with a dependent teenage daughter (FPA 2008). The couple's accumulated superannuation is $550,000. The associated model plan places the breadwinner into salary sacrifice and a transition to retirement pension, thereby reducing the couple's short-term annual tax bill from $38,975 to $22,941.[3] The model plan makes persuasive recommendations for retaining life insurance associated with the client's pre-existing superannuation fund at work. It says: 'The FPA liaised with the Australian Securities and Investment Commission regularly during the development process to arrive at this final version' (FPA 2008, p2).

4.3.1 Fee structure

Initial advice is charged out at $8,277, after tax and on a fee-for-service basis. This initial fee appears to be primarily in exchange for receiving the tax benefits of salary sacrifice and a Transition to Retirement pension. There may be scope to unbundle the initial fee from ongoing fees (in exchange for asset-management services) in the event the couple decides to stick with their pre-existing superannuation fund while adopting the planner's recommendations concerning salary sacrifice and a TTR pension structure. This fee is paid in the first instance to the licensee, who retains 2 per cent of it, and the remainder goes to the planner. In this way, the bubble containing .98 x $8,272 refers to the flat-rate fee actually received by the planner. If the couple does switch its superannuation balances into the fund recommended by the model plan, several ongoing or asset-based annual fees become payable. The investor pays 1.89 per cent pa of assets under management to the product provider. The provider pays 0.6 per cent pa of assets under management to the licensee, from their management fees, and to pay the cost of ongoing advice. Thus the bubble containing .98 x 0.6 per cent refers to the asset fee actually received by the planner.[4] The provider may also pay an additional 0.2 per cent to the licensee, for recommending their products, along with soft-dollar benefits, 'typically between $10,000 and $20,000' pa. Figure 1 summarises these payments.

[Figure 1 here]

Judging by the FPA's model plan, fee-for-service in practice appears to be confined to initial tax advice and does not extend to portfolio formation. Mainstream agency theory − our Section 2 model included − typically does not prescribe pure fee-for-service for an actively-managed portfolio, and this accords with industry practice. The purpose of the two asset-based fees identified in Section 2 is to elicit effort from the manager/adviser that is commensurate with the earnings potential of the asset.

The acknowledged commissions of 0.6 and 0.2 per cent pa, along with the soft-dollar benefits, accord with the theory of Stoughton et al. (2011). Commissions have been contentious on the argument that an agent should not try to serve two principals simultaneously – they create the possibility of a conflict of interest between adviser and client. Moreover, the model plan does not mention any requirement that investors periodically 'opt in,' leaving it open to the criticism of inertia selling.

Elderly couples could be particularly susceptible. For example, if the person responsible for managing household finances passes away before her partner, it could take a considerable time before the surviving partner becomes aware of trail fees in the family's financial plan.

Regulatory Guide 246, promulgated by ASIC (2013), generally bans 'conflicted remuneration' in plans written from 1 July 2013 onwards. This would seem to ban future use of the commissions of 0.6 and 0.2 per cent, as well as the soft-dollar payments. The associated FoFA reforms also require planners to offer advice that is in the best interests of their clients, again from 1 July 2013. FoFA initially proposed requiring clients to 'opt in' every two years. In March 2012, however, this was watered down; membership of an industry association with an ASIC-approved code of ethics now exempts a planner from opt-in.

4.3.2 Asset allocation and asset fees

Table 1 summarises the plan's proposed asset allocation and part of its proposed fee structure. It itemises and breaks down the figure of 1.89 per cent shown in Figure 1.

[Table 1 here]

Table 1 shows the model plan recommends a fund-of-funds portfolio in which each individual fund carries an asset-based fee. The riskier funds on the menu generally carry higher fees. This is on the face of it an incentive to recommend risky products. The generally sizeable fees suggest that most sub-funds envisage adding value via active management. The FPA's model says that the exposure of the elderly couple's superannuation to 'growth' assets (shares plus commercial property) is too low in their pre-existing superannuation fund. It recommends that at least 70 per cent of couple's portfolio, and possibly as much as 95 per cent, be invested in growth assets.[5] The reason is to 'take advantage of market opportunities and your investment time horizon' (p4). By contrast the average Self-Managed Superannuation Fund allocates 30 per cent to safe interest-bearing assets, typically term deposits.

The long investment horizon faced by the elderly couple is actually a reason for caution, and the FPA's advice is at odds with indications that financial plans are fragile around the point of retirement. Bengen (2001) appears to have initiated this line of research, basing his fragility finding on historical simulations with actual returns data for the United States. He noticed that, once you are retired, the sequence of investment returns becomes critical: a market down followed by a market up tends to do more damage than it would have done early in working life. One relevant argument is that if you suffer a big hit early on, you still need to draw down your account balance for living expenses, further depleting it. Even if markets do eventually bounce back you cannot expect to recoup your losses.

Theory based on the notion of a 'protected' consumption level (e.g. our Section 2 model) supports this fragility argument: if your annual expenditure over an expected lifetime cannot fall below some minimum standard, then your asset allocation initially needs to be conservative, reflecting a high present value of protected lifetime expenditure. In practice, protected consumption corresponds to 'ultra' necessities such as electricity, gas and timely medical procedures. The longer your expected time in retirement, the more conservative your initial allocation needs to be. The present value of your protected consumption falls as your remaining years run out, so your proportionate allocation to risky investments can progressively be lifted, provided your risky investments have not underperformed. Another reason for planning an upward-sloping equity-age profile in retirement is that bequests tend to be luxury expenditures and can therefore perform a shock-absorber role late in life.[6]

Another reason for caution on the part of the FPA's hypothetical couple is that the FPA's model plan does not address the couple's apparent lack of retirement flexibility. The plan assumes that the couple's sole breadwinner will continue to work until age 65, or for 3 or 4 years past the recent average retirement age for males. As a consequence, the hypothetical household cannot count on being able to work for longer if investment markets fall before the planned retirement date. Moreover, a setback in health or employment could see the breadwinner forced into retirement before age 65.

Section 2 is consistent with a 'shortfall' notion of risk whereby potential investor losses are capped by placing some wealth in passively-managed safe assets.[7] The FPA's plan appears to recommend that only 5 per cent of the investor's superannuation be unambiguously invested in this way.

The FPA submitted to the Cooper review that portfolio restrictions are unwarranted, particularly in the case of investors on the cusp of retirement: 'Lifestyle [also known as 'glidepath' or 'target-date'] options per se are not necessarily an appropriate strategy for super fund members to adopt. For example, the 10/30/60 rule indicates that the majority of the growth of an investment portfolio occurs during the retirement stage' (FPA 2009, p12). This rule says that 10 per cent of your nominal investment earnings in retirement come from contributions, 30 per cent comes from investment earnings before retirement, and 60 per cent comes from investment earnings after retirement.[8] The FPA evidently sees this rule as reinforcing the case for a comparatively aggressive allocation on the cusp of retirement.

There are problems with both the 10/30/60 rule and the FPA's application of it. First, it appears to be based on comparisons of nominal (rather than real) contributions and investment earnings at widely separated dates, whereas neither nominal contributions nor nominal investment earnings are commensurate across widely separated dates. Second, an allocation's potential for generating high expected nominal earnings at retirement is far from being a sufficient statistic for evaluating it. The extra information needed includes an appropriate forward-looking adjustment for inflation, the client's risk aversion (which could be time-varying), and the price of risk (the ratio of the equity premium to the variance of returns to risky assets). Finally, a comfortable retirement usually necessitates drawing down superannuation balances rather than attempting to maintain them intact.

FoFA should help discourage highly risky allocations. It bans asset-based fees on the borrowed component of sums invested by geared investors, which should help avoid a repetition of the Storm Financial affair. It also seeks to reduce operational and counterparty risks in the managed funds industry, by setting up compensation schemes for investors.

4.3.3 Performance fees

Theory and practice diverge on the question of performance fees. Our Section 2 theory retained a key feature of Dybvig et al. (2010): the performance component of an asset-based fee should be symmetrical in the outperformance of the actively managed portfolio over the zero-effort benchmark portfolio. This symmetry is the defining feature of so-called fulcrum fees. Figure 2 portrays a fulcrum fee.[9] Current regulatory practice in the United States towards managed funds is that fulcrum fees have become the only legal performance fees outside hedge funds. However, this regulatory change saw a big switch by fund managers, away from conventional legal structures and towards hedge funds.

Fee schedules for hedge funds offered in the US are instead tend to be the '2-20' variety: managers do not pay clients if they underperform the agreed benchmark, always receive a fixed asset-based fee of 2 per cent and, in addition, receive an additional fee of 20 per cent whenever the portfolio outperforms the benchmark (Cochrane 2012). Cochrane observes that the corresponding payoff profile resembles the payoff profile for a long position in a call option. The value of a call is an increasing function of the value of the underlying asset, so an option-style payoff motivates managers to exert more effort. But the value of an option is generally an increasing function of the volatility of the underlying asset, so that an option-style payoff introduces an agency problem whereby the manager is tempted to form an excessively risky portfolio, unless sufficiently deterred by a concern for reputation.[10]

Eighteen performance fees used in Australian managed funds are analysed by Whitelaw et al. (2011a).[11] Seventeen of the funds define the benchmark return as returns to the Standard and Poors/Australian Stock Exchange's accumulation indices for its largest 200 or 300 stocks, plus a hurdle rate ranging from zero to five percentage points. No fund offers a fulcrum–style performance fee. On the other hand, all but two offer some variant or other of a 'high watermark' feature. This means performance fees cannot be collected until some or all of any underperformance relative to the benchmark has been recovered. Partially offsetting high watermarks are resets whereby about half the funds in the sample allow their pre-existing high watermarks to be cancelled periodically, putting the manager back in a position to receive fees for outperformance relative to an agreed benchmark without having to make good previous losses.

The heavy solid line in Figure 2 is a stylised portrayal of the performance fee structure of 'fund number 9' in the sample of Whitelaw et al. The figure is stylized because it does not portray the short-run dynamics of high-watermark and reset features. Fund number 9 is typical of the sample. Its fee for outperforming the benchmark is 20 per cent, its benchmark is one plus annualised growth in the S&P/ASX 200, and its base fee is 0.75 per cent, calculated before extracting the performance fee. The heavy solid line shows fund number 9's performance fee before allowing for high watermarks and resets. The heavy dashed line is a stylised portrayal of the average performance fee after allowing for high watermarks and resets. The combined effect of these two features is to push the manager's embedded call option towards being out-of-the-money. This will tend to weaken incentives both for exerting effort and recommending a risky portfolio.

[Figure 2 here]

The sample of Whitelaw et al. (2011a) suggests that the agency problems affecting performance fee schedules in US hedge funds are also present in performance fee schedules for Australian actively-managed funds. Active managers in both countries apparently prefer to rule out downside rather than submit to fulcrum contracts, notwithstanding the relevant prediction (or prescription) of agency theory. In fairness to product providers, one reason for the absence of fulcrum contracts could be investor resistance to the higher fees that might be required. Moreover, an institution offering fulcrum contracts would need to hold capital against the contingency of having to compensate investors for underperforming.

FoFA does not address the question of performance fees, including fulcrum ones. Worth investigating are compromises which would see underperforming managers reimburse investors up to a cap.

4.3.4 Active versus passive management

Theory and practice diverge on the question of the relative merits of active and passive management (paralleling the question of performance fee design.) Grossman and Stiglitz (1976) model an economy in which the expected returns to active and passive management are equalised. Grossman and Stiglitz assume (like Section 2 above) that information acquisition is costly and research investments are rewarded. Active investors drive asset prices towards fair value, and just cover costs in doing so. Yet the international evidence is that active managers perform about 100 basis points pa less well on average than this parity-like condition suggests.

The FPA's model quotes a management expense ratio of 1.89 per cent pa of funds under advice-cum-management (Table 1). Internationally, Vanguard is the best-known provider of index funds. Its Australian Growth Index Fund quotes a management expense ratio of 0.36 per cent. The FPA's recommended fund-of-funds therefore suggests a considerable degree of active management. This inference is consistent with this finding of Whitelaw et al. (2011b), based on their study of 75 large-cap Australian share strategies: 'There is a discernible relationship between active share score and fees' (p2). By the same token, 'a number of vehicles' have 'relatively high fees and low active share scores' (p2). Moreover, the problem of closet indexing appears to be particularly acute in the case of large-cap Australian share funds, which are 'are among the least active globally' (p2).

Whitelaw et al. point out that the ASX/S&P Accumulation index is 'tightly constrained and top-heavy' (p4). Notably, the top 10 holdings in the Australian index accounting for 52 per cent of index capitalisation. This creates an unfortunate interaction with the strong home bias of the FPA's model plan, which allocates only

10 per cent to international shares (Table 1). If you want both an active management style and strong diversification then you would probably be better off with a more internationally diversified portfolio than that proposed by the FPA's model plan.

This problem − of active managers who are actually among the least active internationally − may well derive from our unusual approach of compulsory pre-funded superannuation. MySuper seeks to mitigate the problem. It seeks to ensure that low-cost and diversified funds are available to people who are not actively engaged with their superannuation. It also proposes designing workplace forms and procedures that nudge inactive investors into ticking a box that steers them into low-cost and diversified funds. However, MySuper does not mandate the 'lifestyle' allocations that would reduce risk on the cusp of retirement. In the United States, by contrast, 36 per cent of accumulation fund members are in target-date funds. A big Vanguard plan has seen a 'large steepening of the age-equity allocation gradient' (Mitchell and Utkus, 2012). The share of growth assets can now drop by 20 percentage points or more over working life.

4.3.5 Disclosure

The FPA's 'Example SOA' discloses dollar amounts payable during the first year of the contract. Not disclosed, however, are dollar amounts projected after the first year. For example, the projected total dollar amount payable in the first year of the breadwinner's projected retirement at age 65 can be estimated at $36,181. That amounts to more than half of the couple's 'target' retirement income of $70,000 pa (Kingston 2009).[12] The trail fees include annual asset-based commissions from the product provider. Thus, the projected $36,181 fee payable in the client's retirement year includes a commission of $2,906 from the fund-of-funds to the financial planner (and the licensee.)

The model plan does put a figure on a 'target' level of retirement income, namely $70,000 pa. But it gives no numeric indication of the possible dispersion of actual incomes around this target. An argument for plans based on the notion of 'protected' consumption is that they do address this question of 'lifestyle' risk. Kingston (2009) examines strategies for putting a floor of $50,000 pa in real terms under the couple's spending in retirement, based on the couple's average life expectancy.[13]

Noted earlier was the terse and (in our view) flawed justification for reweighting the portfolio to a comparatively aggressive one with at least 70 per cent in growth assets. The possible incentive effects of asymmetric performance fees are not mentioned at all. The FPA's 'Example SOA' instead refers to a Product Disclosure Statement (access restricted) for details of performance fees.

FoFA does not address questions of disclosure of risks and does not seek to restrict portfolio allocations. On the other hand, it does mandate annual disclosure of dollar amounts payable in fees.

4.4 Conclusion

Our extension of the model of Dybvig et al. (2010) shed light on the tension between agency theory and financial-planning practice. To the extent that investors on the cusp of retirement have concerns about their wealth falling short of some predetermined value, they should simply allocate their wealth partly to safe interest-bearing assets before contracting with the active manager/financial planner. In practice, however, this appears not to be the norm. Rather, planners are entrusted with the bulk of the superannuation balances of their clients and derive most of their income from asset-based fees. As a consequence, fee income from a given client tends to hit a maximum at the outset of the client's retirement. This tempts planners to overweight high-fee growth assets at that particular point of the client's life cycle. By contrast, theory suggests that the case for a high weight on growth assets is stronger at ages well before retirement and, possibly, late in retirement too.

Likewise, theory suggests that fulcrum contracts are the right type of performance fee. In practice, however, we typically see option-type payoff profiles, which tend to promote excessive exposure to growth assets, at least at the outset of a client's retirement.

The proposed model does successfully capture the fact that fees for clients of financial planners or managed funds typically have both flat and proportional (asset based) components. However, option-like payoffs are also commonly seen in a manager's contract. Note that further analysis of the non-linear features of these contracts is beyond the scope of this paper and will be left to future research. .

There is a public interest in financial plans with less aggressive asset allocations for elderly clients, in particular, the taxpayer interest. The Age Pension is indexed to the maximum of wage and price inflation, and is payable for the remaining life of the pensioner, subject to means tests. It can be viewed as public retirement income insurance. It tempts advisers to recommend aggressive asset allocations, since the taxpayer becomes in effect a part guarantor of the client's core retirement income stream. In this way, the Age Pension fall-back promotes moral hazard in advice on asset allocations, with taxpayers picking part of the tab for unsound or unlucky advice. Thus, Harmer (2009, p15) noted that 'Age Pension applications in December 2008 were around 50 per cent higher than the number recorded in October of the same year.'

Between June 2007 and June 2012 the share of self-managed funds in total superannuation balances rose from 27 per cent to 31 per cent. Our Section 2 theory suggests that part of the explanation may be contract designs that were always suboptimal from the standpoint of investors but with weaknesses that remained latent until the financial crisis hit.

The next review of financial advice should examine ways of requiring financial advisers to disclose and respond to the fragility of financial plans for investors on the cusp of retirement. A good start would be this: require Statements of Advice for clients aged over 55 to disclose the percentage allocation to Australian-dollar-denominated interest-bearing securities rated at least 'high quality' by one of the major credit rating agencies.

Appendix 1: Proof of Section 2 Proposition

For convenience we reproduce here, as equation (A1), the relevant Section 2 proposition:

$$\phi(s,\omega) = \underline{C}_m + \frac{(w_0 - \underline{w})\lambda_R}{1+\lambda_R}[R^P + \left(\frac{\lambda_\varepsilon}{\varepsilon\lambda_R}\right)(R^P - R^B)] . \qquad \text{(A1)}$$

On the left-hand side of equation (A1), $\phi(s,\omega)$ is the fee paid by an investor when the manager's unobserved effort ε $(0 \leq \varepsilon \leq 1)$ generates a private signal $s \in S$ about future returns, and the state of the world is $\omega \in \Omega$. On the right-hand side, $\underline{C}_m$ is the protected consumption of the manager, w_0 is investible wealth, $\underline{w}$ is the present value of the total protected consumption of the investor and the manager, λ_R is a Lagrange multiplier on a participation constraint, R^P is the return to the actively managed portfolio, λ_ε is a Lagrange multiplier on an incentive-compatibility constraint, and R^B is the return to a passively-managed (zero-effort) benchmark portfolio.

Proof of equation (A1) begins with a mixture model of the joint density f of s and ω, conditional on effort ε:

$$f(s,\omega;\varepsilon) = \varepsilon f^I(s,\omega) + (1-\varepsilon) f^U(s,\omega) \qquad \text{(A2)}$$

where f^I is the informed (effort-conditioned) distribution of portfolio returns, and f^U is the uninformed density. The uninformed density has the property

$f^U(s,\omega) = f^s(s) f^\omega(\omega)$, where f^ω and f^s are the marginal distributions of f^U with respect to ω and s respectively.

The agency problem here is simultaneously to choose three things: (i) the utility u_i of the investor, given by $u_i(s,\omega) = \ell n(C_i(s,\omega) - \underline{C}_i)$, where $C_i(s,\omega)$ is the consumption of the investor and $\underline{C}_i$ is the protected consumption of the investor; (ii) the utility u_m of the manager, given by $u_m(s,\omega) = \ell n(C_m(s,\omega) - \underline{C}_m)$, where $C_m(s,\omega)$ is the consumption of the manager; and (iii) the manager's effort level ε, to maximise the investor's expected utility. The relevant maximum problem, then, is

$$\max_{u_i(s,\omega),\, u_m(s,\omega),\, \varepsilon} \iint u_i(s,\omega)(\varepsilon f^I(\omega|s) + (1-\varepsilon) f^\omega(\omega)) f^s(s) d\omega ds, \quad \text{(A3)}$$

subject to constraints. One is a budget constraint. For mathematical convenience we use the Dybvig et al. transformation of consumption levels into exponential functions of utility levels:

$$(\forall s \in S) \quad \int (\exp(u_i(s,\omega)) + \exp(u_m(s,\omega))) p(\omega) d\omega = w_0 - \underline{w}, \quad \text{(A4)}$$

where $p(\omega)$ is the pricing density for a claim that pays a dollar in state ω, and $\underline{w} \equiv \int (\underline{C}_i + \underline{C}_m) p(\omega) d\omega$ is the present value of total protected consumption. A second constraint ensures participation by the manager:

$$\iint u_m(s,\omega)(f^I(\omega|s) - f^\omega(\omega)) f^s(s) d\omega ds - c'(\varepsilon) = 0, \quad \text{(A5)}$$

where $c(\varepsilon)$ is the cost of manager effort, and the prime superscript of the function c in (A5) denotes a derivative. A third and final constraint ensures the incentive-compatibility of effort:

$$\varepsilon = \arg\max_{\varepsilon'} \iint u_m(s,\omega)(\varepsilon' f^I(\omega|s) + (1-\varepsilon') f^\omega(\omega)) f^s(s) d\omega ds - c(\varepsilon'). \text{(A6)}$$

Dybvig et al. come up with a lemma that enables replacement of $u_i(s,\omega)$ in the above problem by the investor's indirect utility. We need a minor extension of it to the case of protected consumptions: the expected utility conditional on s for the investor will be shown to equal

$$\log\left(B_i\left(s\right)\frac{f^{\omega}\left(\omega\right)+\varepsilon(f^{I}(\omega|s)-f^{\omega}(\omega))}{p\left(\omega\right)}\right), \tag{A7}$$

where the term $B_i\left(s\right)$ is given by

$$B_i\left(s\right) \equiv w_0 - \underline{w} - \int \exp(u_m(s,\omega)p(\omega)d\omega \tag{A8}$$

and has the interpretation of the investor's share of the budget net of the present value of total protected consumption.

Proof of equation (A7) follows Dybvig et al. The optimal solution must satisfy the sub-problem of maximizing (A3) subject to (A4). Differentiating the Lagrangean for this problem with respect to $u_i\left(s,\omega\right)$ gives

$$[\varepsilon f^{I}(s,\omega) + (1-\varepsilon)f^{\omega}(\omega)]f^{s}\left(s\right) = \lambda_B\left(s\right)p\left(\omega\right)\exp\left(u_i\left(s,\omega\right)\right) \tag{A9}$$

where $\lambda_B\left(s\right)$ is the multiplier to the budget constraint. Integrate equation (A9) with respect to ω and rearrange to get

$$\lambda_B\left(s\right) = \frac{f^{s}\left(s\right)}{B_i\left(s\right)}.$$

Substitute this into equation (A9) to get equation (A7), as required for this paper's counterpart of the lemma in Dybvig et al.

Following Dybvig et al. we set out three definitions of equilibrium returns. The gross portfolio return conditional on observing s is

$$R^{P} \equiv \frac{\varepsilon f^{I}(\omega|s) + (1-\varepsilon)f^{\omega}(\omega)}{p\left(\omega\right)}. \tag{A10}$$

The gross portfolio return without observing s is termed the benchmark return and is given by

$$R^{B} \equiv \frac{f^{\omega}(\omega)}{p\left(\omega\right)}. \tag{A11}$$

Finally, the return under maximum effort $(\varepsilon = 1)$ is

$$R^I \equiv \frac{f^I(\omega|s)}{p(\omega)}. \tag{A12}$$

These definitions give the intuitive decomposition $R^P = \varepsilon R^I + (1-\varepsilon)R^B$.

Equation (A7) enables computation of the investor's expected utility, namely

$$\int \log\left(w_0 - \underline{C} - \int \exp(u_m(s,\omega)p(\omega)d\omega\right)f^s(s)\,ds$$
$$+ \iint \log\left(\frac{\varepsilon(f^I(\omega|s) + (1-\varepsilon)f^\omega(\omega))}{p(\omega)}\right)\left(\varepsilon f^I(s,\omega) + (1-\varepsilon)f^\omega(\omega)f^s(s)\right)ds\,d\omega.$$

(A13)

Differentiate the Langrangean associated with the problem of maximising equation (A13) with respect to $u_m(s,\omega)$ and subject to equations (A5) and (A6). This gives the first-order condition

$$\begin{aligned}\frac{\exp(u_m(s,\omega))\,p(\omega)}{B_i(s)} &= \lambda_R(f^\omega(\omega) + \varepsilon(f^I(\omega|s) - f^\omega(\omega))) \\ &\quad + \lambda_\varepsilon\left(f^I(\omega|s) - f^\omega(\omega)\right)\end{aligned} \tag{A14}$$

where λ_R and λ_ε are the Lagrange multipliers to (A5) and (A6). Multiply both sides by $B_i(s)$ and integrate both sides with respect to ω to get an expression for the manager's share of the budget net of the present value of total protected consumption, namely, $B_m(s) \equiv w_0 - \underline{w} - B_i(s)$:

$$B_m(s) = \lambda_R B_i(s). \tag{A15}$$

Apply equations (A4) and (A15) to get

$$B_i(s) = \frac{w_0 - \underline{w}}{1 + \lambda_R}. \tag{A16}$$

Equations (A14) and (A16) together imply

$$u_m(s,\omega) = \log\left(\frac{(w_0-\underline{w})\lambda_R}{1+\lambda_R}\frac{f^{\omega}(\omega)+\left(\varepsilon+\frac{\lambda_\varepsilon}{\lambda_R}\right)(f^{I}(\omega|s)-f^{\omega}(\omega))}{p(\omega)}\right). \tag{A17}$$

Taking exponentials of both sides gives equation (A1), as required to complete the proof of it.

A.1 Numerical Analysis

This section gives a numerical analysis of optimal contracts. Paralleling Dybvig et al. we made the following assumptions for $f^{s}(s)$, $f^{\omega}(\omega)$, $f^{I}(\omega|s)$ and $p(\omega)$, where the market state ω and signal s have zero means and standard deviations σ and have correlation $\rho > 0$. In what follows, r is the risk free rate and μ is the mean return on the market:

$$f^{s}(s) = \frac{1}{\sigma\sqrt{2\pi}}\exp\left(-\frac{s^2}{2\sigma^2}\right) \tag{A18}$$

$$f^{\omega}(\omega) = \frac{1}{\sigma\sqrt{2\pi}}\exp\left(-\frac{\omega^2}{2\sigma^2}\right) \tag{A19}$$

$$f^{I}(\omega|s) = \frac{1}{\sigma\sqrt{2\pi(1-\rho^2)}}\exp\left(-\frac{(\omega-\rho s)^2}{2\sigma^2(1-\rho^2)}\right) \tag{A20}$$

$$p(\omega) = e^{-r}\frac{1}{\sigma\sqrt{2\pi}}\exp\left(-\frac{(\omega+\mu-r)^2}{2\sigma^2}\right). \tag{A21}$$

Set $\mu = 0.1$, $\sigma = 0.2$, $\rho = 0.5$, $r = 0.05$, $\varepsilon = 0.5$, $w_0 = 100$ and, initially, $\underline{C}_i = \underline{C}_m = 5$. By choosing nonnegative Lagrange multipliers λ_R and λ_ε, we can plot the investor's wealth and the manager's fee for the first-best and second-best problems. The first-best case arises as we let λ_ε tend towards zero in Eq. (A1). Because effort can be observed in this case, there is a zero shadow price of tightening the incentive-compatibility constraint.

The figures show the second-best fee minus the first-best fee. Figure A.1 shows the manager is rewarded when signal and market outcomes are both high and is therefore induced to exert effort. Figure A.2 considers the effects of an increase in the manager's protected consumption, from 5 to 10 units. The manager is now rewarded less when signal and market outcomes are high, indicating that higher protected consumption leads to less effort. More generally, comparison of the vertical axes of the two figures shows that with high protected managerial consumption there is now comparatively little difference between effort levels regardless of signal and market outcomes.

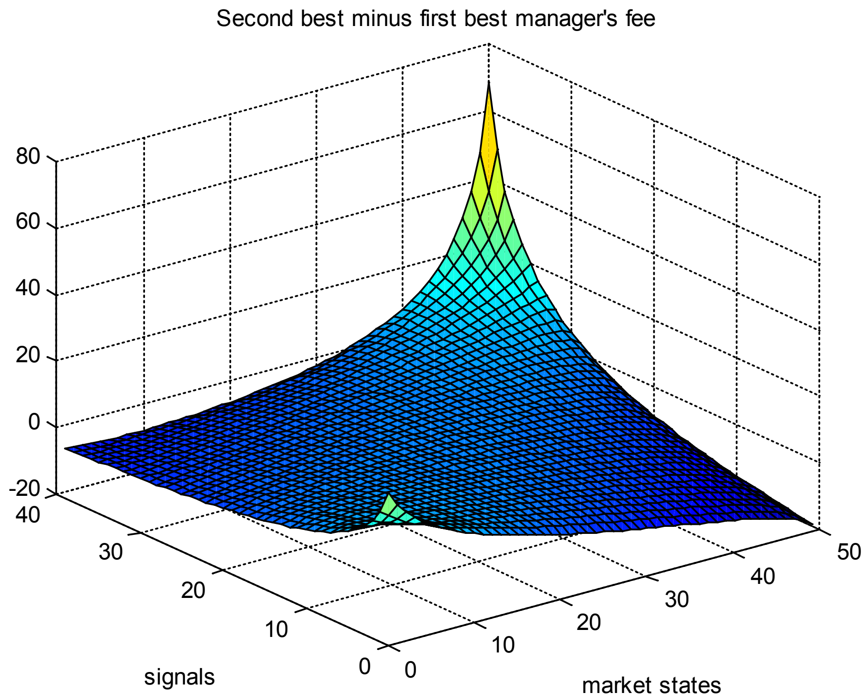

Figure A.1 Manager's payoff: Second-best minus first-best levels

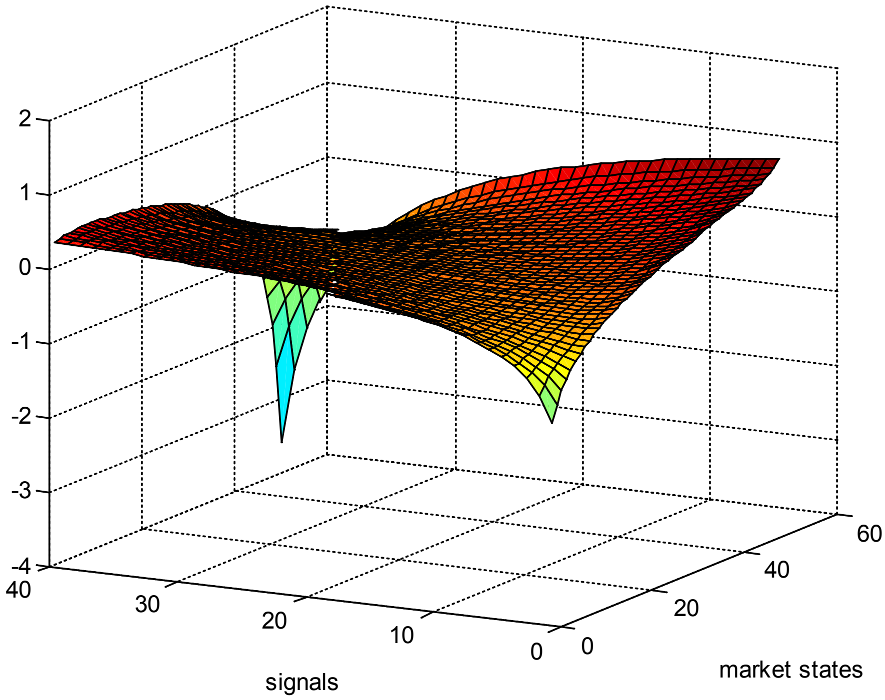

Figure A.2 Difference in second-best minus first-best fee when $\underline{C}_m$ is increased

A.2 Discussion

Moving from simple to generalised log utility changes the optimal contract in two ways. First, the efficient fee structure incorporates a flat component $\underline{c}_m$. Second, the proportional component $\lambda_R / (1+\lambda_R) R_P$ is not based on the investor's entire wealth w_0 but on an amount net of the present value of total protected consumption $\underline{w}$.

The first-best case arises as we let λ_ε tend towards zero in equation (A1). Because effort can be observed in this case, there is a zero shadow price of tightening the incentive-compatibility constraint.

The concept of protected manager consumption $\underline{c}_m$ is vague, and this lack of specificity is a source of strength as well as weakness. Notably, it facilitates alternative interpretations. For example, $\underline{c}_m$ might take the indirect form of a minimum investible amount stipulated by the manager, in conjunction with a proportional asset fee. We typically see such minima in the prospectuses of actively managed funds. In this way, the flat component could be interpreted as the fixed cost of operating an account. Alternatively, $\underline{c}_m$ could represent the wage costs of providing ancillary services such as tax minimisation, which could be more 'commoditised' than skilful active management.

The benchmark portfolio return R^B is also usefully flexible. It can be interpreted as the return to an index portfolio of equities, or as the return to an equities-plus-cash portfolio that has had no value added via efforts to time the market. In the same way, the risks entailed by active management can be interpreted as originating solely from investment risk, or also from operational and other non-investment risks specific to active management. We can interpret the uninformed distribution $f^U(s,\omega)$ as including such risks and the effort level ε as including efforts to reduce them.

Endnotes

1. Also known as the Stone-Geary utility function.
2. See e.g. Wachter and Yogo (2010) for a review of the evidence on relative risk aversion falling with wealth (which they rationalise by a distinction between necessities and luxuries rather than 'protected' consumption − a device which can be interpreted as playing the role of necessities without the complication of a relative price between necessities and luxuries.)
3. There is currently a cap of $25,000 p.a. on concessional contributions by people with over $500,000 in super, so tax benefits on this scale are not currently available.
4. Neither this fee nor the one shown in the bubble containing 98 x $8,272 ought be interpreted as separate payments from the licensee to the planner. Rather, routing payments via the licensee in this way is presumably for the purpose of mitigating operational risks.
5. 'It's my understanding you are willing to implement a less conservative strategy to meet your objectives...I have allocated approximately 30% to cash and income funds to cover pension payments'– FPA (2008, p4). The generic assets in 'income' funds

are not disclosed.

6. See Bateman et al. (2007), Ding et al. (forthcoming) and Kingston (2009). Constant-mix allocations, through time and across the major asset classes, are associated with constant relative risk aversion (also known as 'power' utility). Actual proportionate allocations to risky assets tend to rise with an investor's wealth (recall Section 1).

7. More precisely, funds placed under active management were given by the investor's initial investment 'cushion', namely $w_0 - \underline{w}$.

8. For an exposition, see e.g. Russell (2008). Ironically, this influential exposition also endorses a glide-path approach – the FPA treats the 10/30/60 rule as an argument against glide paths.

9. In terms of our Section 2 theory, the slope of the ray in the figure is given by $\lambda_\varepsilon / \varepsilon\lambda_R$ and outperformance is given by $R^P - R^B$.

10. In terms of options analysis, if a call has a knockout feature and the value of the underlying asset is sufficiently close to the knockout price, then increases in volatility will reduce the value of a call. The real-options analogue here is that an investor might switch to a different fund if the value of the original fund has fallen sufficiently. There is a substantial literature on these considerations, which we ignore, to save space.

11. Whitelaw et al. are concerned with the comparative expense of performance fee structures rather than the implications for incentives (i.e. the concerns here.)

12. Financial planning practices have typically sold on multiples of three or four times annual revenues. By contrast, accounting practices have typically sold on multiples of two.

13. Kingston builds on estimates in Bateman et al. (2007). In the terminology of dynamic asset allocation, this strategy is constant-proportion portfolio insurance with a finite horizon. The case of generalised log utility corresponds in practice to a multiple of about one. This is conservative compared to typical infinite-horizon CPPI strategies, which typically have multiples in the range of 3 to 5. Such strategies are

not particularly conservative.

References

Australian Prudential Regulation Authority 2012, Statistics: Annual Superannuation Bulletin, viewed June 2013, http://www.apra.gov.au/Super/Publications/Pages/annual-superannuation-publication.aspx.

Australian Securities and Investments Commission 2013, Regulatory Guide 246: Conflicted Remuneration, viewed June 2013, <http://www.asic.gov.au/asic/pdflib.nsf/LookupByFileName/rg246-published-4-March-2013-B.pdf/$file/rg246-published-4-March-2013-B.pdf.>

Bateman, H., Kingston, G., and Thorp, S. 2007, 'Financial Engineering for Australian Annuitants', in H. Bateman (ed.), Retirement and Scary Markets, Edward Elgar, Cheltenham, pp. 124-148.

Bengen, W. 2001, 'Conserving Client Portfolios during Retirement, Part IV', Journal of Financial Planning (May issue), pp. 110-119, viewed June 2013, <http://www.fpanet.org/>.

Cochrane, J. 2012, 'Hedge Funds' (slides), Booth School of Business, University of Chicago, viewed June 2013, <http://faculty.chicagobooth.edu/john.cochrane/.>

Ding, J., Kingston, G. and Purcal, S. forthcoming, 'Dynamic Asset Allocation when Bequests are Luxury Goods', Journal of Economic Dynamics and Control.

Dybvig, P., Farnsworth, H., and Carpenter, J. 2010, 'Portfolio Performance and Agency', Review of Financial Studies, vol. 23, pp. 1-23.

Financial Planning Association 2008, 'Simplifying Statements of Advice: FPA Example SOA', viewed June 2013, <http://www.scribd.com/doc/60713243/SOA-Transition-to-Retirement>.

Financial Planning Association 2009, 'Superannuation Review: Phase Two – Operation and Efficiency', viewed June 2013, <survey.fpa.asn.au/files/Policy09CooperPhase2submission.pdf>.

Grossman, S., and Stiglitz, J. 1976, 'Information and Competitive Price Systems', American Economic Review vol. 66, pp. 246–253.

Harmer, J. 2009, Pension Review Report, Commonwealth of Australia, Canberra.

Mitchell, O., and Utkus, S. 2012, ‘Target-Date Funds in 401(k) Retirement Plans’, working paper, Wharton School, University of Pennsylvania.

Kingston, G. 2009, ‘Financial Plans for Baby Boomers: How much Risk?’, Economic Papers, vol. 28, pp. 65-74.

Rainmaker Group 2013, ‘Financial Planning Report’, viewed June 2013, <http://www.rainmaker.com.au>.

Ross, S. 1973, ‘The Economic Theory of Agency: The Principal's Problem’, American Economic Review vol. 62, pp. 134-139.

Russell Investments 2008, ‘The Russell 10/30/60 Retirement Rule’, viewed June 2013, <http://www.russell.com/CA_FP/PDF/Invest4Retir_Investr_Final.pdf>.

Stoughton, N., Wu, Y. and Zechner, J. 2011, ‘Intermediated Investment Management’, Journal of Finance, vol. 66, pp. 947-980.

Wachter, J. and Yogo, M. 2010, ‘Why do Household Portfolio Shares Rise in Wealth?’, Review of Financial Studies, vol. 23, pp. 3929-3965.

Whitelaw, T., O’Donnell, K. and Gray, P. 2011a, ‘Best Practices in Managed Fund Performance Fees’, unpublished paper, Morningstar.

Whitelaw, T., O’Donnell, K. and Gray, P. 2011b, ‘Active Share: The Activeness of Large-Cap Australian Share Fund Managers’, unpublished paper, Morningstar.

Portfolio share (per cent)	Investment sector	Investment options	Management fee (per cent)
5	Cash	Cash	1.13
5	Income	Income extra	1.77
20	Income	Income fund	1.92
5	Listed property securities	Property securities fund	1.66
10	Australian shares	Australian active equity	1.86
17	Australian shares	Boutique Australian shares	1.96
8	Australian shares	Australian equity long/short	2.24
20	Australian shares	Australian small companies	1.91
10	International shares	Global value equity	2.01
100 per cent			*1.89 per cent*

Table 1: Asset Allocation and Product Fees in the 'Example SOA'

Source: FPA (2008).

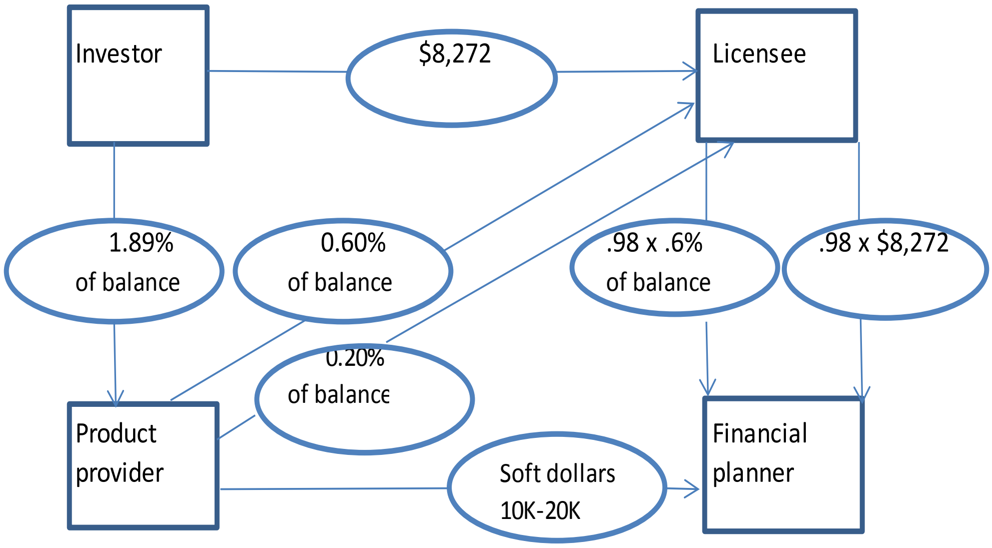

Figure 1: Fee Structure in the 'Example SOA'

Source: constructed from data in FPA (2008).

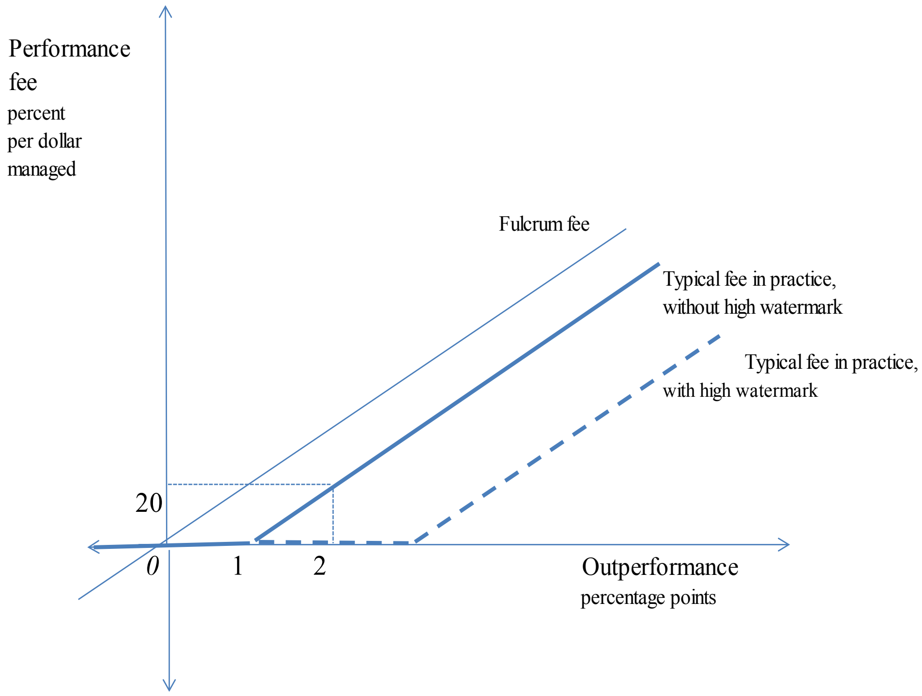

Figure 2: Alternative Performance Fee Structures

5. Backfilling Financial Data with an Iterative PCA-based Imputation

Haijie Weng (contribution: 90%) and Stefan Trück (contribution 10%)

This paper has been presented at:

- Internal seminar at the Royal Bank of Scotland, Singapore, January 2013

Abstract

In this paper an iterative PCA-based data imputation algorithm is proposed for handling missing values in financial time series. The designed backfilling algorithm generates satisfactory results for both synthetic and real data, covering equity, rates and FX asset classes. Further, the proposed model outperforms other commonly used approaches for data imputation. The model serves as a robust tool for risk managers to backfill missing values in financial data, since complete data is a prerequisite for generating correct VaR numbers. It is worth noting that the performance of the model depends on the fraction of the missing data and the noise of the data set.

5.1 Introduction

The value-at-risk (VaR) concept has emerged as one of the most prominent measures of downside market risk, where VaR is defined as the lower end of a 99% confidence interval for a given time horizon (typically a day or two weeks). In 1997, a Market Risk Amendment to the Basle Accord permitted banks to use VaR estimates for setting bank capital requirements related to trading activity. To calculate VaR, historical simulation has been adopted by most of banks as the standard industry approach, see, e.g., Jorion (2000) and Alexander (2001).

In this approach, changes in major risk factors observed during historical periods are combined with sensitivities to generate vectors of profits and losses for a portfolio from which VaR is calculated. The key advantage of this method is that there is no explicit requirement to model joint distributions, i.e. the multivariate distribution of all of the risk factors. Typically the distribution of the individual risk factors is non-Gaussian, while correlations between the factors will be time-varying and subject to regime changes. However, for historical simulation, the joint behaviour of the risk factors is implicitly captured in historical data. For example, if two risk factors are strongly correlated, then we should expect to mostly see those factors move together in the historical price and return data.

One of the major concerns for historical simulation relates to the quality of the available data. Historical data are usually sourced from various data vendors and it is not uncommon for downloaded data to be of poor quality. The problematic data can be missing, meaningless or unlikely. In practical applications, missing data may be due to a market close, or insufficient contributor depth in the case of composite quotation, or system failure. In the case of meaningless data, although the data is present, it may violate some sort of condition. For example, negative FX spot rates, negative FX option volatility, or a normally liquid time series that suddenly shows complete staleness: all these examples of data are meaningless. Unlikely data are those that behave abnormally, for example, a sudden spike of FX volatility in a normal market environment.

Once problematic data has been detected, correction protocols have to be applied. Given the large set of time series to be cleansed and maintained, for practical application in financial institutions such a correction tool should run on a fully automated basis. The default corrective action is simply to re-query the original source system from which data is snapped. Should that fail to resolve the problematic value, the second action is to impute a value based on relevant statistical information about the time series and its nearby neighbours. For example, an error in the 7 year interest swap rate can be corrected with reference to its 5 year and 10 years swap rates.

To our best knowledge, limited research has been conducted to apply advanced and automated techniques to backfill financial time series. Karelmo (2010) uses a basic PCA based algorithm to fill the missing observations in corporate bond time series. Mailleta and Merlin (2009) propose a way that does not require any hypothesis and is totally data driven to complete the missing values in hedge fund monthly return time series. Minsky et al. (2010) simply backfill missing return data for a hedge fund by randomly selecting its peers' returns. All the studies are either based on a relatively simple approach or developed for a particular asset class. In this paper, by adopting the regularised PCA algorithm proposed by Josse et al. (2011), we attempt to propose an advanced PCA-based backfill procedure for financial time series and to test the imputation performance using time series of various asset classes. The main focus of this paper is on backfilling missing observations, which result from removal of problematic data (i.e., missing, meaningless and unlikely data). There is a clear gap between the regularised PCA algorithms in Josse et al. (2011) and backfilling missing financial data, a problem faced by many practitioners and researchers. In this paper, we further extend and enrich their model with additional functions that cater for dealing with financial time series. Our extensions lead to an algorithm that allows for backfilling missing financial data and test the performance against other competing models.

This paper focuses on backfilling missing values through Principal Component Analysis (PCA) and a detailed literature review on using this approach for missing data is provided in section 5.2. However it is worth reviewing other strands of literature dealing with missing data. For example, Beckers and Rixen (2003) adopt empirical orthogonal functions (EOF) to infer missing data from oceanographic data series. They calculate the missing data from an optimal number of EOFs determined by a cross-validation technique. One advantage of EOF is that they do not require a priori information about the error covariance structure and are parameter free. An application of EOF to reconstruct incomplete oceanographic data sets is, for example, also found in Alvera-Azcarate et al. (2005).

Another strand of the literature uses machine learning to impute missing data. For example, Breiman (2001) proposes a machine learning methodology called "Random Forest". In the applied algorithms, the decision trees grow iteratively and data are classified. The rules learned from these classifications are used for the imputation of missing values (see also Nourani et al. (2008), Rustum and Adeloye (2007), Kim and Pachepsky (2010)). Other methods that have been proposed in the literature to deal with missing values also include the so-called inverse distance method by Xia et al. (1999) and the simple arithmetic averaging (SAA) method by Xia et al. (1999, 2001).

In this paper we decide to focus on correlation based imputation methods, including PCA-based imputation and Expectation Maximization methods, rather than the above reviewed methods. Admitting that those methods may also generate good quality of imputation results and have less requirements on the correlation structure among the time series, we argue that correlation based methods for missing data are still widely used in the financial industry. This is also due to the fact that very often financial time series for related or similar products exhibit strong correlation. We also argue that correlation based imputation methods would be easier to be endorsed by practitioners, given that the main purpose of this paper is to propose a tool to efficiently and accurately backfill the missing financial time series data for the practitioners.

The paper is organized as follows. Section 2 provides a brief review of principal component analysis (PCA) and its applications in data imputation. Section 3 proposes an iterative PCA-based date imputation procedure for financial time series. Section 4 conducts empirical analysis of the proposed method on various data sets. Section 5 compares the proposed model with other data imputation methods and discusses the results. Finally, Section 6 concludes.

5.2 Principal Component Analysis (PCA)

Principal component analysis (PCA) is a well-established technique for reducing the dimensionality of a large set of data while retaining as much variability as possible. A substantial number of PCA-based backfilling routines have been proposed by practitioners and researchers, see e.g. Alexander (2009), for backfilling missing data in various areas, including climate records, bioscience, software engineering, just to mention a few. The popularity of PCA comes from three important properties. First, it is the optimal linear scheme for compressing a set of high dimensional vectors into a set of lower dimensional vectors. Second, the model parameters can be computed directly from the data – for example by diagonalizing the sample covariance. Third, given the model parameters, the original data can be easily reinstated from the compressed data without much loss of information: they require only matrix multiplications.

The Principal Component representation is defined as follows: Let X be a $T \times n$ matrix of random variables where T and n are the number of rows and columns of matrix X respectively, and V denotes its corresponding $n \times n$ covariance matrix.

Furthermore, let W be the $n \times n$ orthogonal matrix of V. The $T \times n$ matrix P, where columns that correspond to principal components (as an exact linear combination) of X, is then given by the relation: $P = XW$ or equivalently $X = PW^T$.

By selecting the first k columns of P and W and thus creating P^* and W^*, an approximation of X can be obtained through the relation $X^* = P^* W^{*T}$. This is the dimension reduction benefit of PCA.

The principal components are retrieved through an eigenvalue decomposition of a covariance matrix of a set of observable variables. The i^{th} eigenvalue λ_i of V is obtained by taking the sum of squares of each element in the corresponding i^{th} principal component. Given that the total variation is explained by the sum of the eigenvalues of V, one can easily derive an expression for the fraction of the variability which is explained by the first k components: $\frac{\lambda_1 + \lambda_2 + ... + \lambda_k}{\lambda_1 + \lambda_2 + ... + \lambda_n}$. For more details on PCA, refer to Alexander (2009), for example.

In the presence of missing data, a number of alternatives have been suggested, see, e.g. Jolliffe (2002) for a more detailed review. A quick and simple method is to replace missing values by the mean value calculated from the available observations. A more sophisticated approach of imputation is regression-based PCA backfilling, proposed by Grung and Manne (1998). The authors suggest a regression based method in which factors and factor loadings are obtained by a two-step regression procedure. Another alternative offered by imputation is to assume a distribution based on the observed data and to simulate the missing values from the distribution. This procedure is repeated multiple times and the variability of the missing value rather than a single value is estimated (Schafer 1997).

A different class of procedure is based on maximum likelihood estimation (Little and

Rubin 1987). Under the assumption of a multivariate normal distribution, an expectation maximization (EM) algorithm together with an iterative procedure is applied to estimate the missing values. The iterative PCA algorithm (also named EM-PCA) consists of initialising the missing values, performing PCA on the available observations, filling-in missing values with the reconstruction formula and iterating until convergence. To overcome the shortcoming that PCA is highly sensitive to outliers, Stanimirova et al. (2007) propose an Expectation-Maximization Spherical Principal Component Analysis (EM-SPCA) to backfill missing values in a real data environment. Their approach uses a robust PCA based method combined with the expectation maximization algorithm to deal with missing values and outlying observations simultaneously. The proposed method works well for highly contaminated data containing different amounts of missing elements. The authors claim that EM-SPCA outperforms standard EM-PCA in the case where outliers exist. Schneider (2001) proposes a regularised EM algorithm which is suggested to be more suitable in cases where the number of variables exceed the number of observations. The regularized EM algorithm is based on an iterated analysis of linear regressions of variables with missing values on variables with available values, with regression coefficients estimated by ridge regression.

The quality of the prediction of the missing values will deteriorate in accordance with an increasing number of principal components, creating an overfitting problem. Problems of overfitting are exacerbated with increasing numbers of missing values. To overcome this problem, Josse et al. (2011) propose an iterative regularized EM-PCA algorithm which imposes a “shrunk” imputation step to remove the noise-causing instabilities in the predictions.

Another suggested PCA based backfilling routine is presented by Kondrashov and Ghil (2006). They use Singular Spectrum Analysis (SSA) to fill the gaps in several types of spatial-temporal data sets. Making use of the spatial and temporal correlations, they iteratively produce estimates of missing data points. The algorithm is demonstrated on both simulated and real data and yields promising results for single missing value and even longer continuous gaps. The main challenge to apply SSA is to define the optimal parameters, which depend on the distribution of missing data, as well as on the variance distribution between oscillatory modes and noise.

5.3 Proposed Iterative PCA Imputation

This section describes the proposed PCA-based data imputation algorithm. It consists of three major steps: clustering, selecting the number of components and imputation as described below.

5.3.1 Clustering

Before applying PCA, one may want to categorise all the time series into different groups. The objective of grouping is to have high co-linearity among all the time series in the same group, so that PCA is suitable. Intuitionally, time series can be first grouped by asset classes (IR, FX, Equity, Credit, etc.). The data set by asset class can be further sub-grouped based on specific features of each asset class. For example, for interest rates we can further group the data by currencies (USD, EUR, JPY, etc.), while equity time series can be grouped by region, market capitalisation, industry, etc.

Alternatively, one can group the time series based on observed correlations, an approach that is also adopted in this paper. Grouping by correlations serves the purpose of reducing the noise brought by other less correlated time series. Moreover, in each subgroup, the number of principal components is reduced when a lower number of time series is considered, such that also the computational time is reduced. For each time series with missing values, we can find the most correlated time series from two available options: choosing the top *m* most correlated series or choosing the series with absolute correlation greater than a certain threshold value. The determination of the threshold value is either based on long term historical correlations between two asset classes or a discretionary figure that risk managers consider to be reasonable.

5.3.2 Selecting the number of components

The next step is to determine the number of principal components used in the PCA analysis. On the one hand, if too few components are kept, relevant information may not be taken into account in the analysis. On the other hand, an excessive number of components could also be problematic, since the components with small eigenvalues that only explain a small fraction of the variability may be considered as noise. Then, considering too many components might cause an overfitting problem. In the following, we adopt the general cross-validation approximation proposed by Josse and Husson (2012) to select the number of components. Once the number of components is specified, it is kept fixed in the imputation step.

5.3.3 Imputation step

We adopt the regularised PCA algorithm proposed by Josse et al. (2011) for the step relating to imputation. The regularised PCA algorithm can help to overcome the overfitting problem in the missing data framework and improves the estimation of principal components and the prediction of missing values. Further we can repeat the imputation several times by using the previously imputed data. Our results show that repeating imputations can further improve the results by a small margin.

We define a $m \times n$ matrix *E*, containing the 1-day time series of changes in the risk factors that is calculated from raw data. Hereby, *n* is the number of time series falling in the same group in the clustering step and *m* is the length of the time series. In the matrix *E*, missing positions are identified and defined as a subspace $\Omega = (t, j)$, where $x(t, j)$ are the missing observations with *t* and *j* being the row and column number in the matrix *E*. Each column of the matrix *E*, is standardized by subtracting its mean and then dividing by its standard deviation[1]. This standardization ensures that time series are mean zero distributed and have a standard deviation of one. The missing values are then initialized with zero. The original set of observations can be restored from this standardized version of the data set by simply multiplying the reconstructed vectors with the standard deviation and adding back the originally estimated mean.

The imputation procedure can now be initiated with the standardized matrix $\hat{E}_0$, but is described for the general matrix, $\hat{E}_i$, obtained in the ith iteration, by the algorithm below:

1. A principal component analysis is performed on matrix $\hat{E}_i$ such that principal components (PCs) P_i and loadings W_i are obtained.
2. The first k ranked PCs are selected together with corresponding loadings, resulting in a subset of PCs and loadings denoted by P_i^* and W_i^*. An estimate $\hat{E}_i^*$ of $\hat{E}_i$ is obtained by $\hat{E}_i^* = P_i^* W_i^{*T}$.
3. From the estimate $\hat{E}_i^*$, values are taken only from positions initially identified as missing, i.e., $\forall(t, j) \in \Omega$, and imputed into corresponding positions of the matrix $\hat{E}_i$, resulting in the new matrix $\hat{E}_{i+1}$.
4. Step 1 to 4 are repeated until convergence or until the maximum number of iterations is reached.

It is also worth noting that the matrix E consists of daily data for the changes in the considered variables, which are calculated from the raw price data, e.g., equity prices, zero rates and FX rates.

[1] To obtain series that are stationary and IID, normally univariate ARMA and GARCH filters need to be applied on each of the time series to remove autocorrelation and conditional volatility. For simplicity, we merely standardize the series by subtracting the mean and dividing by the standard deviation of each time series and assume the resulting series are stationary and IID.

Since missing data are essentially from the raw data, additional steps are required to ensure the imputed 1-day change data are in line with the raw data. For example, we observe the raw data of series *j* at time *t*, denoted as $y(t,\ j)$, is missing. In the matrix *E*, $x(t,\ j)$ and $x(t+1,\ j)$ are therefore missing since $x(t,\ j)$ and $x(t+1,\ j)$ are calculated as $y(t,\ j)-y(t-1,\ j)$ and $y(t+1,\ j)-y(t,\ j)$ respectively. We therefore need to adjust $x(t,\ j)$ and $x(t+1,\ j)$ to ensure that $x(t,\ j)+x(t+1,\ j)=y(t+1,\ j)-y(t-1,\ j)$.

We apply the following method to adjust the imputed data: Assuming there are *n* consecutive missing raw data starting from *t-n+1* and ending on *t*, such that all observations for series *j* between $y(t+1, j)$ and $y(t-n+1, j)$ are missing. We then define D as the deviation of the sum of the imputed changes $x\_imputed(t+i, j)$ for the missing observations and the actual difference between $y(t+1, j)$ and $y(t-n+1, j)$:

$$D = y(t+1, j) - y(t-n+1, j) - \sum_{i=0}^{n+1} x\_imputed(t+i, j).$$

To match the sum of the imputed changes with the observed difference between $y(t+1, j)$ and $y(t-n+1, j)$, we add $D/n+1$ to each imputed data.

A flow chart for the applied algorithm is provided in Figure 1.

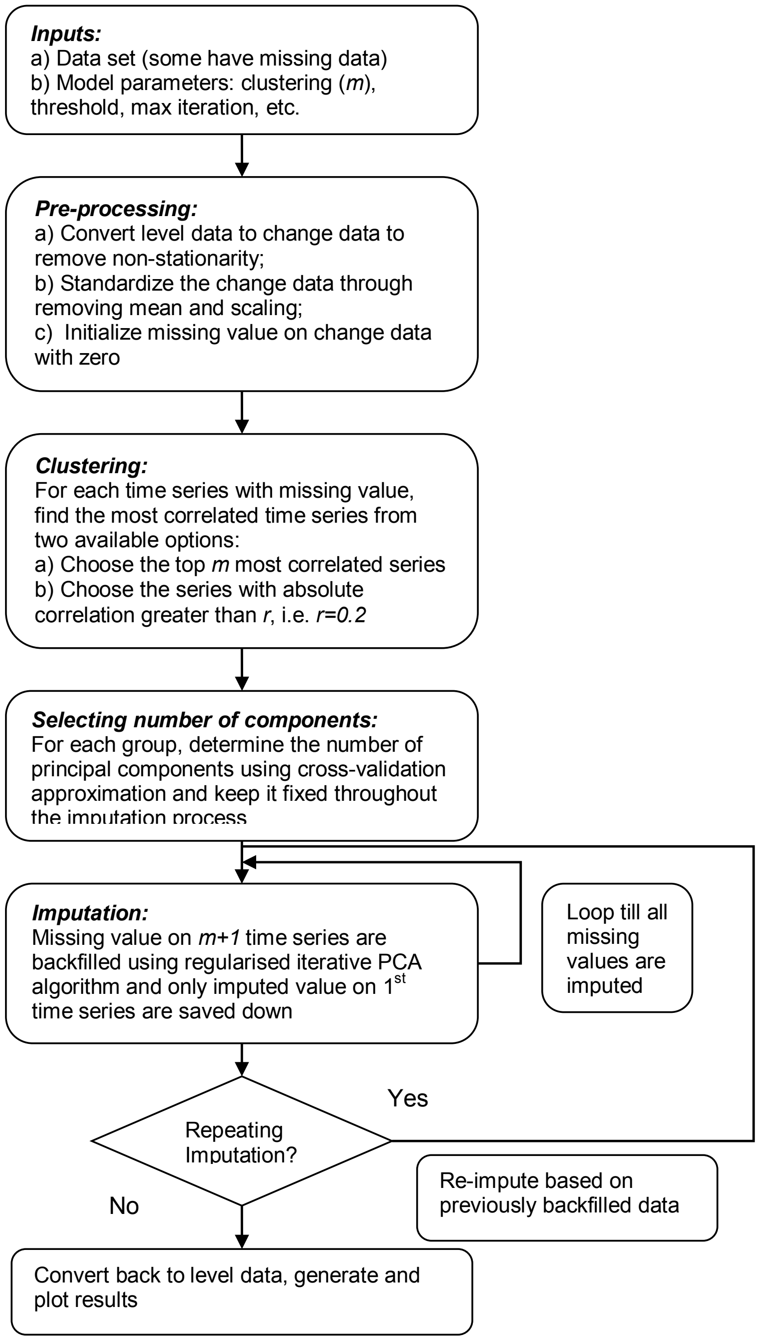

Figure 1: Flow chart of the proposed imputation algorithm that outlines the pre-processing step, the clustering step, the selection of the number of components and the imputation step.

5.4 Model Validation

To ensure the validity of the proposed backfilling procedure and to further reach understanding of its limitations, imputation validation tests are performed on two sample data sets (synthetic and real data) with missing data points imputed using the proposed backfilling algorithm.

To measure the performance of the backfilling procedure, the mean and maximum of the absolute difference between true and estimated value among all imputed positions, and 90% and 95% percentiles of the absolute difference distribution are reviewed. The sample data structure and validation process, together with obtained results for the applied imputation algorithm are described in this section.

5.4.1 Synthetic Data Set

The synthetic data sets are simulated using the following underlying model:

$$Y_{mn} = F_{ms}U_{sn} + \varepsilon_{mn}$$

The *m* by *s* matrix F and *s* by *n* matrix U are simulated from a standard normal distribution $N(0,1)$. Then each column of the *m* by *n* matrix $F_{ms}U_{sn}$ is divided by its standard deviation, estimated by column in order to control the signal to noise ratio. Finally a noise term ε_{mn} is added by simulating from a normal distribution with mean equal to zero and standard deviation equal to σ. Therefore, the signal to noise ratio is 1/ σ. In our simulation experiment, we set *m=500*, *s=5*, and *n=30,* where *m* and *s* are the number of rows and columns for matrix F, respectively, and *n* is the number of columns for matrix U. We also let σ vary between 0.25 and 1. With the simulated data set, we randomly remove between 5% up to 50% of the data and test the performance of the imputation under different scenarios.

For each time series with any missing data, the selection of correlated time series is conducted either by picking a predetermined number of series ranked by correlation or by setting a threshold value for the correlation coefficient. The former method fixes the number of series to be clustered but implies the risk of selecting less correlated series in comparison to the latter method, which explicitly sets a minimum value for inclusion. We test the following two scenarios:

(i) Selecting the 10 series exhibiting the highest coefficient of correlation with the series that contains missing data; and,

(ii) Selecting all series with a coefficient of correlation $\rho > 0.2$.

Further, we are interested in investigating whether a recursive imputation algorithm outperforms single imputation and hence test whether repeating the imputation step five times provides better results than conducting the imputation procedure only once.

Figures 2-5 illustrate the results for the mean of the absolute differences, the maximum of the absolute difference as well as the 90% and 95% percentiles of the absolute difference between imputed and actual values for different values of the variance σ^2 for the noise term. Note that the y axis shows the value for the considered performance criteria, i.e. the mean, the maximum, the 90% and 95% percentiles of the absolute difference between imputed and actual values, while the x axis contains the percentage of missing data and the noise level. Clearly, the higher the chosen sigma, the more noise is introduced in the data and the worse is the performance of the applied methods with respect to the considered measures. As indicated, combining the 10 time series with the highest correlation per group and repeating the imputation step 5 times yields the overall best results. We find that this implementation of the backfilling algorithm typically yields the lowest values for the considered performance criteria.

In particular, under the same number of imputations, using the top 10 most correlated series appears to outperform a setup that uses a minimum correlation threshold for grouping the data. Further, we find that under the same setup for grouping the correlated series, the recursive imputation outperforms a method that only conducts a single imputation step. Thus, overall we find that a fixed group size and recursive imputation seem to provide the best results for the proposed backfilling algorithm. As expected, we also observe that the results for the imputation algorithm deteriorate with an increasing fraction of missing data and for a higher standard deviation parameter of the noise. This is true for all four considered criteria.

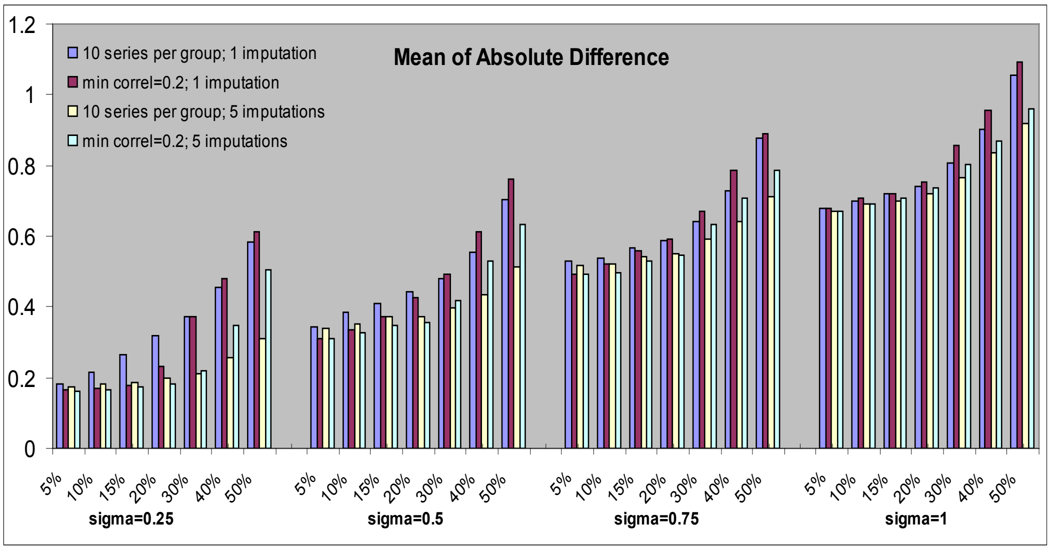

Figure 2: Synthetic data imputation results: mean of absolute difference

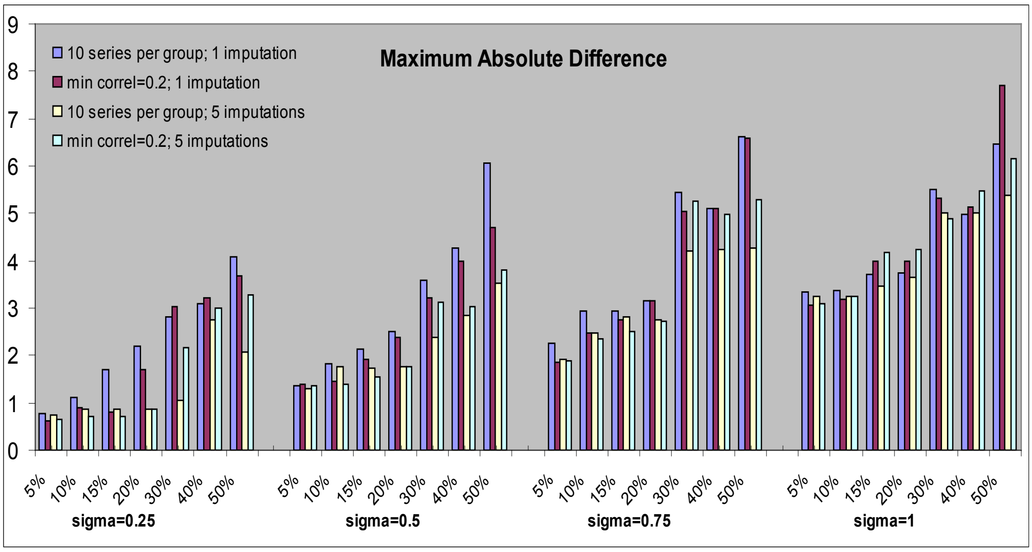

Figure 3: Synthetic data imputation results: maximum absolute difference

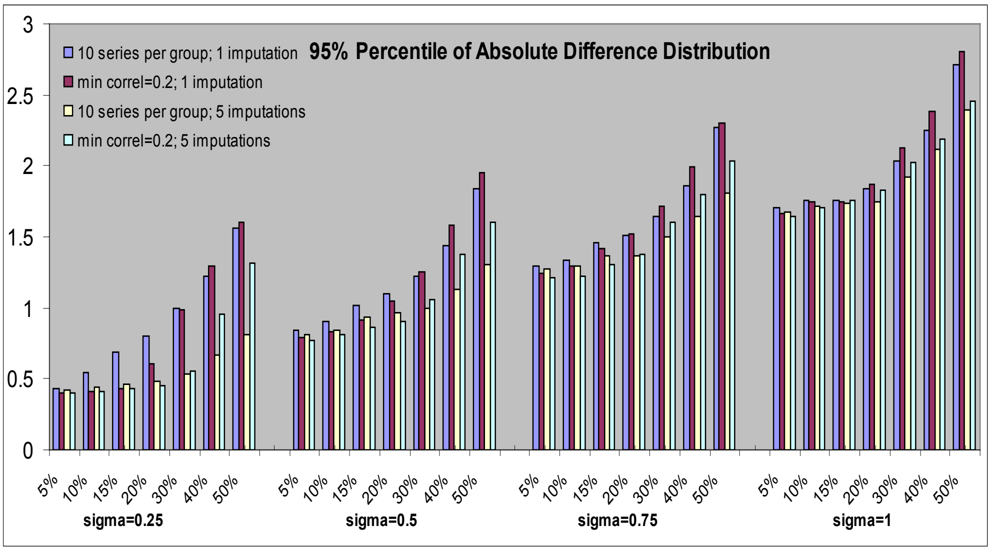

Figure 4: Synthetic data imputation results: 95% percentile of absolute difference distribution

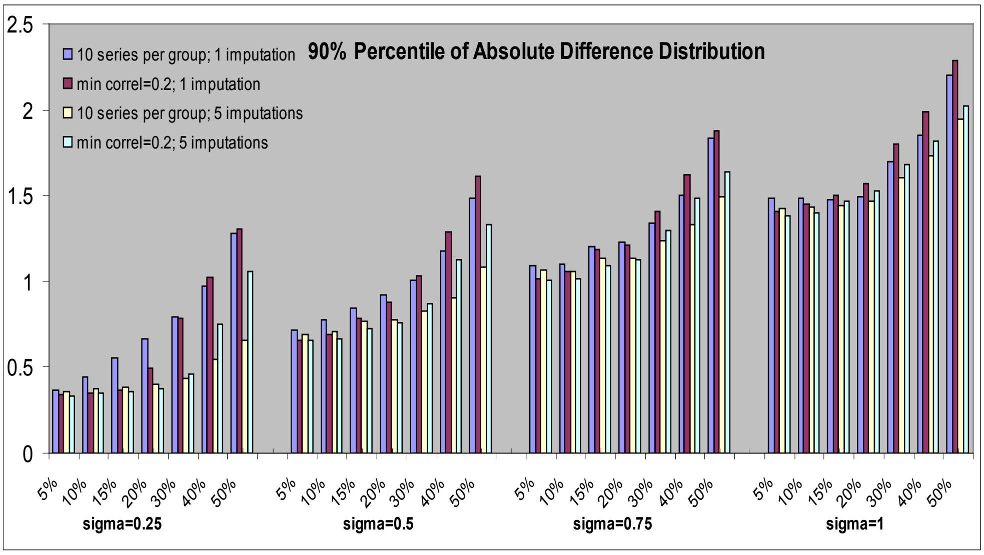

Figure 5: Synthetic data imputation results: 90% percentile of absolute difference distribution

5.4.2 CNY Data

To validate the proposed algorithm on empirical data, we design a test scenario to backfill price data on products related to the Chinese yen (CNY) exchange rate. We use information inherent in USD exchange rate products with the same currencies to extract the missing information for the CNY. We expect USD related products to contain relevant information also for the CNY, given that the CNY is pegged to the

USD. We construct a data set with 28 time series covering CNY and USD rates instruments, i.e., deposit, deposit futures and swaps. Each time series has 500 observations ranging from 09 October 2006 to 05 September 2008. To test the developed algorithm, we randomly remove 361 observations what roughly corresponds to 2% of the total observations. The results are provided in Table 1.

	Average sum of absolute difference (bps)	Maximum absolute difference (bps)	99%	95%	90%
10 series per group; single imputation	2.91	**47.0**	24.3	8.3	**5.5**
Min correlation =0.1; single imputation	3.08	51.4	25.2	8.5	6.3
10 series per group; five imputations	**2.90**	47.1	**24.2**	**8.2**	5.6
Min correlation =0.1; five imputations	3.08	51.6	25.2	8.4	6.3

Table 1: CNY imputation results using CNY and USD rate instruments (28 time series in total)

Note: The minimum value in each column is bold and underlined.

From Table 1, we observe that the combination of selecting the 10 most correlated times series and performing multiple rounds of imputation yields the best results for three of the criteria considered. This set up of the algorithm outperforms other test combinations in 3 of 5 indicative test statistics, while performing almost equally well as the best method for the other two criteria.

Further, we notice that the backfilling result for the CNY_Deposit_1W series is the worst among all the imputed series. Results regarding the quality of the imputation when excluding this series are provided in Table 2. As shown, the results are greatly improved after leaving out this particular time series. Interestingly now all four implementations of the algorithm produce rather similar results, although the fixed group size and recursive imputation setup marginally outperforms the other methods.

	Average sum of absolute difference (bps)	Maximum absolute difference (bps)	99%	95%	90%
10 series per group; single imputation	2.52	13.63	11.24	7.23	5.42
Min correlation =0.1; single imputation	2.53	13.54	11.24	7.12	5.43
10 series per group; five imputations	**2.52**	**13.51**	**11.22**	**7.10**	**5.41**
Min correlation =0.1; five imputations	2.53	13.53	11.23	7.12	5.43

Table 2: CNY imputation results without CNY_Deposit_1W series (27 time series in total)

Note: The minimum value in each column is bold and underlined.

The poor imputation results for the CNY_Deposit_1W series are due to a lack of similar time series. Therefore we further include three more series: CNY NDF (non-deliverable forward) Deposit with ON (overnight), 1 week and 6 months tenors and tabulate the results in Table 3. Since the CNY_NDF_Deposit_1W series is highly correlated with the CNY_Deposit_1W series, the results are significantly improved as compared with those in Table 1. This is true in particular for the maximum absolute difference. The results also suggest that finding highly correlated time series is vital in obtaining satisfying results for data imputation. As indicated in Table 3, it seems that choosing the top 10 most highly correlated series per group with 5 imputation steps yields the best result again.

	Average sum of absolute difference (bps)	Maximum absolute difference (bps)	99%	95%	90%
10 series per group; single imputation	2.49	13.44	10.93	6.93	5.34
Min correlation =0.1; single imputation	2.98	26.13	12.62	8.82	6.52
10 series per group; five imputations	**2.48**	**13.21**	**10.83**	**6.91**	**5.32**
Min correlation =0.1; five imputations	2.98	26.12	12.52	8.82	6.52

Table 3: CNY imputation results using additional CNY NDF (non-deliverable future) series (31 time series in total) Note: The minimum value in each column is bold and underlined

5.5 Comparison against other imputation methods

To further assess the performance of the proposed data imputation method against other competing methods, two more approaches have been selected for comparison: principal component regression and univariate linear regression. Principal component regression has enjoyed large popularity in a wide range of fields because of its capability to confront the situation that there are too many highly correlated predictor variables or too small sample size- a situation that is quite common in natural sciences. In the literature, the principal component regression has been applied to impute the missing data, for example, in the medical survey data sets (Marivate et al. 2007), in compositional data (Hron et al. 2010), and in multivariate statistical process control (Arteaga and Ferrer 2002). The univariate linear regression (ULR) is a least square based imputation, which first models the relationship, expressed as regression coefficient, between a dependent time series containing missing data and one explanatory time series with complete data and then computes the missing positions in the dependent time series using the estimated regression coefficients and the corresponding observed values in the explanatory time series. The applications of univariate linear regression in backfilling missing data are widely seen in the literature. For example, ULR has been adopted for estimation of the missing values in microarray data sets in biological research (Hellem et al. 2004) and for handling missing data in clinical trials (O'Kelly and Ratitch 2014). Rodwell et al. (2014) compare the performance of various methods, including ULR, for imputing limited-range variables in biomedical research.

Following a brief description of the two algorithms, we conduct an additional empirical study to analyse the backfilling performance for each algorithm. We cover three different types of data series each from interest rates, equity and foreign exchange asset classes, namely CNY interest rates, volatility surface parameters for EURGBP FX options and Hong Kong equity index options. We consider time periods of 7 years from 2005 to 2011 for the CNY interest rates and a two year series from 2008 to 2009 for EURGBP FX and Heng Seng equity index options. The rationale behind selecting periods spanning from 2005 to 2011 for interest rates and from 2008 to 2009 for FX and equity options can be explained in two aspects: first, both periods cover the recent financial crisis happened in 2008 and consequently contain volatile and extreme data points, which has created more challenges to the backfill algorithms and test the model performance more convincingly than using tranquil market data. Second, interest rate expresses mean reversion property in history and hence using longer period (e.g., 6 years) helps to capture the correlation between each interest rate instrument and thus improves model performance. In contrast, the correlation structures for FX and equity options are embedded in recent period and we expect less information can be extracted from data dated back long into history, and therefore we consider two years' data is sufficient enough.

In summary, PCA imputation outperforms the two competing methods, as indicated by the test results. Further, it is worth noting that the performance of the PCA imputation method and principal component regression (PCR) are comparable in some cases, such that PCR could also be used as a verification tool for the PCA imputation.

5.5.1 Principal Component Regression (PCR)

In multiple linear regression, one of the major difficulties is the problem of multicollinearity, which occurs when there are near-constant linear functions of two or more of the predictor variables. Multicollinearity is often indicated by large correlations between subsets of the variables. If multicollinearities exist, the variances of some of the estimated regression coefficients can become very large, leading to unstable and potentially misleading estimates of the regression equation. Principal component regression (PCR) has gained strong popularity as a means to overcome this problem. PCR uses the PCs of the predictor variables in place of the predictor variables. As the PCs are uncorrelated, there are no multicollinearities between them, and often the principal components with the highest variance are selected.

In the context of multiple linear regression, the least square solution for $Y = XB + \varepsilon$ is given by $B = \left(X^T X\right)^{-1} X^T Y$. The problem is often that $X^T X$ is singular because of existing multicollinearities. PCR circumvents this by decomposing X into orthogonal scores P and loadings W^T, where $X = PW^T$, and by regressing Y not on X itself but on the first α columns of the scores P. Note that typically Y and X denote returns or changes in the price data in order to remove the non-stationarity that often exists for level data. Missing values for Y are then calculated from estimated regression coefficients and the observed X. We refer to Jolliffe (2002) for more details on the application and estimation of PCR.

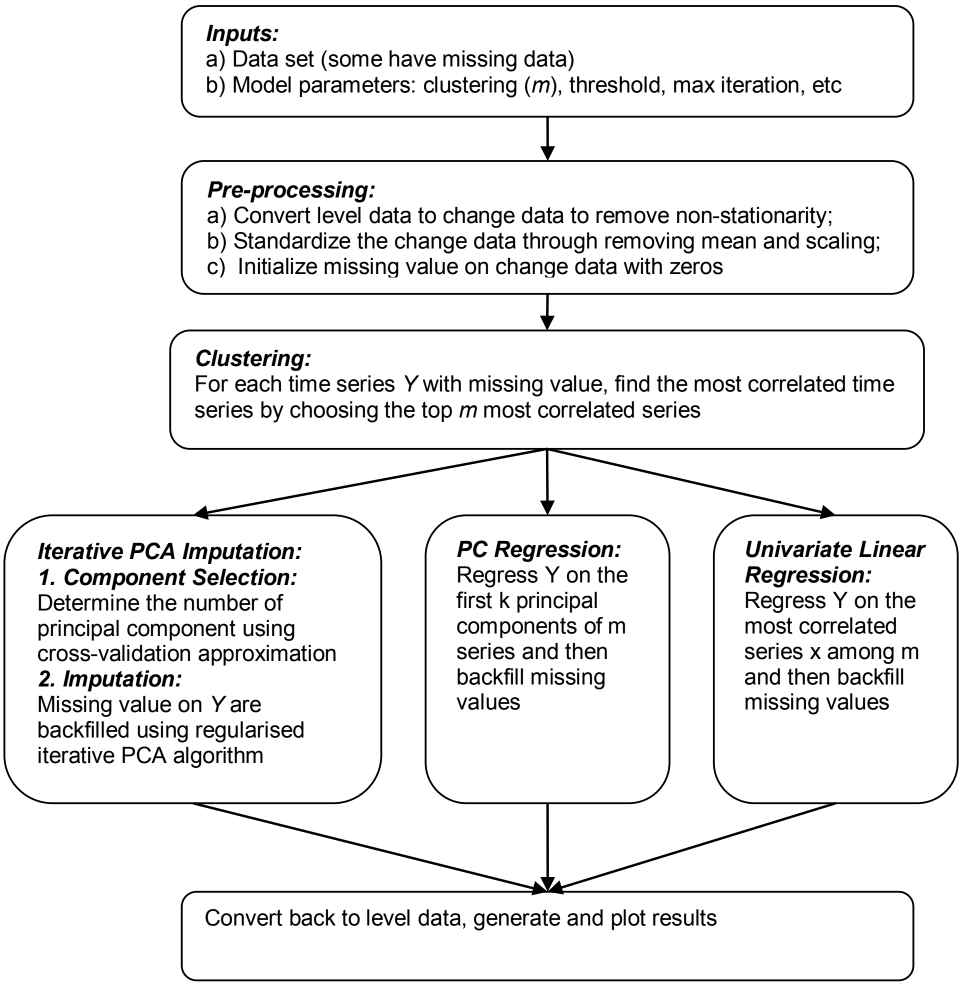

Figure 6: A flow chart illustrating the three proposed imputation algorithms: iterative PCA

imputation and the benchmark models PC regression and univariate linear regression.

5.5.2 Univariate Linear Regression

In this approach, we simply regress Y on a time series x which has the highest correlation with Y and backfill the missing positions in Y using the estimated regression coefficients and the corresponding observed values in x.
A flow chart, outlining the algorithm for each of the three imputation methods is provided in Figure 6.

5.5.3 Test Cases

In the following we compare the performance of the three algorithms over various asset classes, covering interest rates, as well as options on equity and foreign exchange rates. In particular, we backfill time series of CNY interest rate instruments, volatility surface parameters[2] for Heng Seng Index options and FX volatility surface building instruments[3] for EURGBP FX options.

5.5.3.1 CNY rates

The data sets selected for imputation include CNY interest rates (deposit rates and swap rates), USD interest rates (deposit rates and swap rates) and CNY government bonds. In this test, 50 continuous points for the considered CNY rates instruments across all its tenors (i.e., maturing in one month to ten years) are removed from the data set. Hereby, the entire time series contain 1568 data points, covering the period from January 1, 2005 to December 31, 2011, and 50 consecutive data points of all CNY interest rates time series, covering the period from January 28, 2007 to April 6, 2007 are removed for imputation. Time series for USD rates instruments and CNY government bonds are complete and contain information for the imputation of the

---

[2] The volatility surface is a plot of the implied volatility of an option as a function of its strike price. A call/put option gives the holder the right, but not the obligation, to buy/sell the underlying asset by a certain date (defined as the maturity of the option) for a predetermined price (defined as the strike of the option).
The volatility surface of an equity option (e.g. Heng Seng Equity Index Option) can be modelled using three Stochastic Alpha Beta Rho (SABR) model parameters: at-the-money volatility (the level of the surface), the rho (the skewness of the surface) and the vol of vol (the convexity of the surface).

[3] In the context of FX option (e.g., EURGBP FX Option), its volatility is usually modelled using at-the-money (ATM) volatility (the level of the surface), the risk reversal (the slope of the surface) and the straddle (the curvature of the surface). The ATM option, risk reversal and straddle are liquidly traded in the market, and hence being used for building FX option volatility surface.

missing CNY rates. Note that for the conducted analysis all CNY rates data are missing within the same period; hence no information is available among CNY rates instruments for backfilling and backfilling has to rely on information from USD rates and CNY government bonds instead.

We measure the performance of the proposed iterative PCA imputation algorithm and the two benchmark models in terms of the ratios of maximum and average absolute differences between the imputed and true value, divided by the averaged true values. Results for the considered performance criteria are reported in Table 4. A graphical representation of the results for the considered CNY rates across different tenors is provided in Appendix A. Hereby, the smallest average absolute difference or maximum absolute difference is an indicator for the best performance of a particular method. Over the considered data sets, we simply count the number of winners to rank the performance.

	Maximum absolute difference			**Average absolute difference**		
Time series with 50 missing points	IPCA	PCR	ULR	IPCA	PCR	ULR
Deposit_CNY_1M	**20.6**	26.9	33.6	**12.3**	16.5	21.8
Deposit_CNY_2M	**10.7**	13.5	17	**6.4**	8.4	11
Deposit_CNY_3M	**1.8**	1.8	2.0	**1.0**	1.1	1.1
Deposit_CNY_6M	6.6	7.2	**4.4**	2.7	2.9	**1.7**
Deposit_CNY_9M	5.1	**4.7**	5.8	**2.1**	2.2	2.4
Deposit_CNY_1Y	**5.2**	5.3	6.2	2	**1.9**	2.5
Swap_CNY_2Y	**6.3**	6.4	6.9	3.1	**2.9**	3.6
Swap_CNY_3Y	**6.8**	6.9	7.1	**1.7**	**1.7**	2.3
Swap_CNY_4Y	**9.5**	9.7	9.9	1.9	**1.8**	2.2
Swap_CNY_5Y	**10.1**	10.2	10.3	**2.0**	**2.0**	2.4
Swap_CNY_7Y	12.9	13.2	**12.6**	**2.0**	2.2	2.2
Swap_CNY_10Y	**12.7**	12.9	12.8	**2.0**	2.1	2.1

Table 4: Maximum and average absolute difference for CNY imputation

Note: IPCA stands for Iterative PCA, PCR stands for Principal Component Regression and ULR stands for Univariate Linear Regression. The maximum absolute difference and average absolute difference are measured between the imputed and the real value and expressed as ratio to the mean of *the real values in percentage. The two smallest values in each row for maximum and average absolute difference respectively are highlighted in bold and underlined*

In general, as indicated in Table 4, the Iterative PCA, or simply IPCA, is the best performer followed by Principal Component Regression (PCR). In particular, IPCA provides the best results for nine of the considered series with respect to providing the smallest maximum absolute difference, and for eight series with respect to providing the minimum average absolute difference. PCR comes second, yielding the lowest maximum absolute difference for one series and the lowest average absolute difference for five of the considered series. The simple ULR approach performs worse, and yields the best lowest maximum absolute difference for two of the series and the lowest average absolute difference for only one of the series. The performance of ULR is also unstable as it is highly sensitive to the correlation between two time series. If the correlation is low, particular in this test case where the correlations between incomplete data (i.e., CNY rates) and complete data (i.e., USD rates and CNY government bonds) are rather low in nature, the results from ULR appear to be the least accurate.

5.5.3.2 Equity Option Volatility Surface

The equity option volatility surface for a given maturity is typically modelled using three Stochastic Alpha Beta Rho (SABR) model parameters: at-the-money volatility (the level of the surface), the rho (the skewness of the surface) and the vol of vol (the convexity of the surface). Here we choose the volatility surface for Hong Kong Heng Seng Index options with expiry ranging from one week to two years for imputation. Again we select six time series and remove 50 consecutive data points to examine the performance of the suggested backfilling algorithm and the benchmark models. The entire time series contains 523 data points, covering the period from January 1, 2008 to December 31, 2009, while the 50 consecutive observations of the selected six time series that are selected for removal cover the period from March 3, 2008 to May 9, 2008. Please refer to Appendix B for plots of the time series, as well as for plots of the results for the applied backfilling algorithms. Similar to the conducted analysis for the CNY rates, we find that the suggested iterative PCA method generally outperforms the other two competing algorithms. In particular, IPCA yields the best results for 50% of the considered time series and performance criteria, while both PCR and ULR perform best for only 3, respectively, 4 of the considered cases. Interestingly, ULR performs better for the considered equity option volatility surface, since the correlation for each SABR parameter time series across different maturities, is relatively high. Therefore, we obtain clearly better results for the ULR method in comparison to our analysis for the CNY rates instruments, where correlations between CNY and USD rates were typically much lower.

	Maximum absolute difference			**Average absolute difference**		
Time series with 50 missing points	IPCA	PCR	ULR	IPCA	PCR	ULR
HSI_1Y_SKEW	3.1	4.4	**<u>2.4</u>**	1.3	1.5	**<u>1.0</u>**
HSI_2Y_SKEW	**<u>1.3</u>**	1.9	2.9	**<u>0.5</u>**	0.7	1.3
HSI_1W_CONVEX	43.1	**<u>33.8</u>**	36.6	8.7	**<u>5.8</u>**	7.2
HSI_2M_CONVEX	7.7	9.6	**<u>6.4</u>**	2.7	3.6	**<u>2.3</u>**
HSI_2M_ATM	**<u>1.6</u>**	2.3	2	**<u>0.5</u>**	0.9	0.7
HSI_2Y_ATM	**<u>0.4</u>**	0.4	1.1	**<u>0.1</u>**	**<u>0.1</u>**	0.4

Table 5: Maximum and average absolute difference for Equity option volatility surface imputation

Note: IPCA stands for Iterative PCA, PCR stands for Principal Component Regression and ULR stands for Univariate Linear Regression. The maximum absolute difference and average absolute difference are measured between the imputed and the real value and expressed as ratio to the mean of the real values in percentage. The two smallest values in each row for maximum and average absolute difference respectively are highlighted in bold and underlined. Time series are presented in an "Index Name_Maturity_SABR Parameter" format. For example, HSI_1Y_SKEW represents the skewness time series for Heng Seng Index option expiring in 1year.

5.5.3.3 FX Option Volatility Surface

Further, we consider backfilling missing data for FX option volatility surfaces. Note that the FX option volatility surface for a given maturity (i.e., 6 months or 1 year) is described by the at-the-money volatility (the level of the surface), the risk reversal (the slope of the surface) and the straddle (the curvature of the surface). The risk reversal is defined as the difference between the volatility required to price the 15-delta put (a put option having a delta of 0.15) and the volatility to price the 15-delta call (a call option having a delta of 0.15) for a given maturity. The strangle is defined as: 0.5* (the sum of the 15-delta put and the 15-delta call volatility- 2* the at-the-money volatility). Here we choose the volatility surface for the EURGBP option expiring in 6M and 1 year for imputation, and 6 times series with 50 missing points for each series are selected to backfill and report. The time series contain 523 data points, covering the period from January 1, 2008 to December 31, 2009, while the 50 consecutive data points that were removed from six of the time series to test the imputation algorithm, cover the time period from January 28, 2008 to April 4, 2008. Please refer to Appendix C for the plots of comparisons. Once again, the suggested Iterative PCA algorithm overall performs best among the considered techniques.

	Maximum absolute difference			**Average absolute difference**		
Time series with 50 missing points	IPCA	PCR	ULR	IPCA	PCR	ULR
RR_6M	**<u>9.5</u>**	10.1	10.9	**<u>1.7</u>**	2.7	3
RR_1Y	13.7	20.5	**<u>12</u>**	5.7	8	**<u>3.7</u>**
ST_6M	**<u>3.9</u>**	5.7	8.6	**<u>1.7</u>**	1.8	3.9
ST_1Y	8	7	**<u>6.1</u>**	4.5	2.7	**<u>2.1</u>**
ATM_6M	2.4	**<u>2.2</u>**	2.8	1.1	**<u>0.9</u>**	1
AMT_9M	**<u>1.6</u>**	1.7	3.4	**<u>0.7</u>**	0.8	0.9
ATM_1Y	**<u>1</u>**	1.3	1.1	0.5	0.7	**<u>0.4</u>**

Table 6: Maximum and average absolute difference for FX option volatility surface imputation

Note: IPCA stands for Iterative PCA, PCR stands for Principal Component Regression and ULR stands for Univariate Linear Regression. The maximum absolute difference and average absolute difference are measured between imputed and real data and expressed as a ratio to the mean of the real values in percentage. The two smallest values in each row for maximum and average absolute difference respectively are highlighted in bold and underlined. RR denotes risk reversal, ST denotes *strangle, ATM denotes at-the-money volatility. The last two letters denotes the maturity of the FX option.*

In particular, IPCA provides the best results for 7 out of 12 series and criteria, followed by ULR and PCR that perform best only for 5, respectively, 2 criteria. Interestingly, also for this study ULR outperforms PCR due to the high correlation among the considered time series which is important for achieving stable results for ULR. Note, however, that often for backfilling financial data, missing data are more frequently observed in less liquid or lower trading volume asset classes, like CNY rates, for which it is difficult to find highly correlated counterparts. Hence, overall one might still prefer PCR over ULR in practical applications. Given that the performances of IPCA imputation and PCR are comparable in some cases, PCR can therefore be used as a verification tool for the PCA imputation suggested here. For example, if the two approaches generate close backfilling results, we are comfortable to accept the IPCA results.

5.5.3.4 Distribution Comparison

In this section we compare the distribution of imputation errors of the three competing imputation methods. In particular, we make use of the synthetic data set described in Section 5.4.1 and conduct the following simulation experiment n=100 times: in each simulation run, we generate 30 correlated time series, each containing 500 data points. Then we randomly remove 4% of the data points for imputation, resulting in 600 missing points for imputation. For each simulated data set, we then apply the iterative PCA imputation method as well as the two benchmark models, i.e. PC regression and the univariate linear regression. For each simulation run, we then calculate the mean squared error (MSE) and the mean absolute error (MAE) based on the deviation of the imputed values from the actual simulated data points.

The empirical distribution of the imputation errors (expressed as imputed value minus the true value) for one of the simulation runs is plotted in Figure 7. Apparently, the iterative PCA imputation method outperforms the two benchmark methods and typically provides smaller deviations of the imputed from the actual values.

Table 7 provides descriptive statistics for MSEs and MAEs for the missing data based on the conducted simulation experiment with n=100 runs. Both for average MSEs and MAEs the PCA imputation method clearly provides superior results, i.e. smaller MSEs and MAEs than the two benchmark competitors. In particular, the linear univariate regression performs much worse than the PCA imputation method, but also than the PC regression approach.

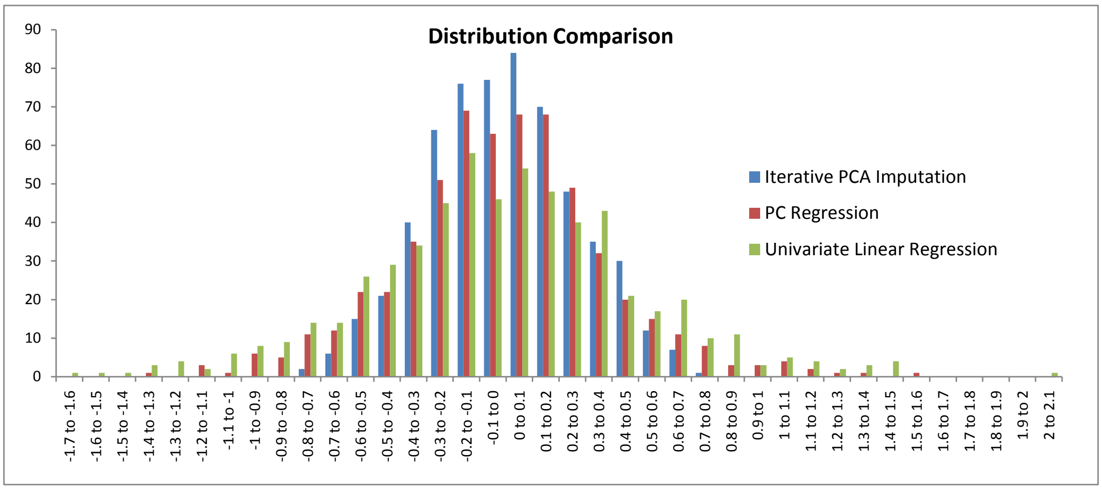

Figure 7: Empirical distribution of the imputation errors for one of the simulation runs (600 missing data points) for the three imputation algorithms: iterative PCA imputation, PC Regression and Univariate Linear Regression.

Statistics	Iterative PCA Imputation	PC Regression	Univariate Linear Regression
Mean Squared Errors (MSEs)			
Mean	0.0442	0.0876	0.2744
Min	0.0373	0.0548	0.2127
Max	0.0517	0.1730	0.3580
Standard Deviation	0.0028	0.0219	0.0307
Mean Absolute Errors (MAEs)			
Mean	0.1670	0.2283	0.4094
Min	0.1532	0.1844	0.3602
Max	0.1811	0.3107	0.4770

Standard Deviation	0.0054	0.0248	0.0235

Table 7: Descriptive statistics of Mean Squared Errors (MSE) and Mean Absolute Errors (MAE) for the missing data based on the conducted simulation experiment with n=100 runs. In each simulation run, 4% of the simulated data are removed, resulting in 600 missing data points for imputation.

Figure 8 also provides a plot of the MAE for the three based on 100 simulations runs with 600 missing data points.

In a next step we apply statistical tests to formally examine differences between the performances of the three methods. Using a Kruskal-Wallis test, we test whether the error samples generated by the three competing imputation methods are statistically different. The Kruskal-Wallis test is a nonparametric version of the classical one-way analysis of variance (ANOVA), and tests the null hypothesis that all samples are drawn from the same population, or equivalently, from different populations with the same distribution (Hollander and Wolfe, 1999). The test does not require the samples to follow a normal distribution. Both for MSE and MAE, the conducted test rejects the hypothesis that the samples are drawn from the same population at all reasonable levels of significance ($p<0.001$). We conclude that the distributions of MSEs and MAEs are statistically significantly different.

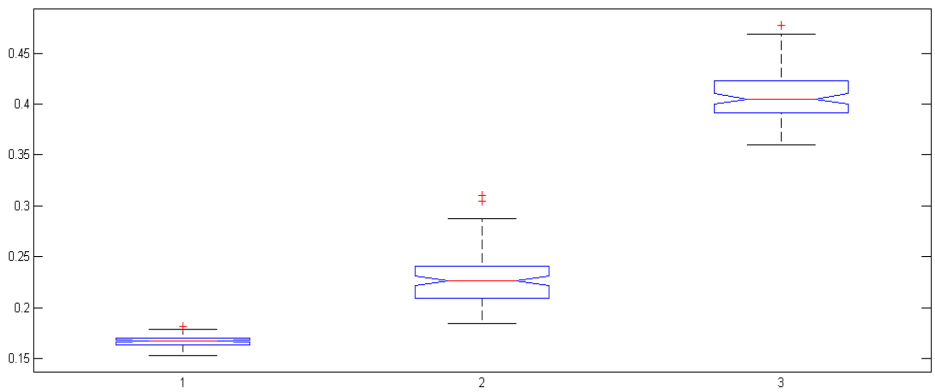

Figure 8: Empirical distribution of MAEs for imputation methods iterative PCA imputation

(left), PC Regression (middle) and Univariate Linear Regression (right) based on 100 simulations runs with 600 missing data points.

Finally, we conduct a multiple comparison procedure in order to further investigate which of the samples are significantly different, see Hochberg and Tamhane (1987). The test uses Tukey's honestly significant difference (Tukey's HSD) criterion that is optimal for the comparison of groups with equal sample sizes, to test for significant differences with respect to the performance of the methods. The test is conducted with a significance level of $\alpha = 0.05$ and uses the rank statistics of the nonparametric Kruskal-Wallis test. For both MAEs and MSEs the multiple comparison procedure suggests that the PCA imputation method provides mean ranks that are significantly different from the two benchmark methods, i.e. that the approach provides MAEs and MSEs that are significantly smaller. The procedure can also significantly distinguish between the rank distribution of the MAEs and MSEs for the other two methods, suggesting that the applied PC regression significantly outperforms the univariate linear regression. The results for the conducted multiple comparison procedure are also illustrated in Figure 9.

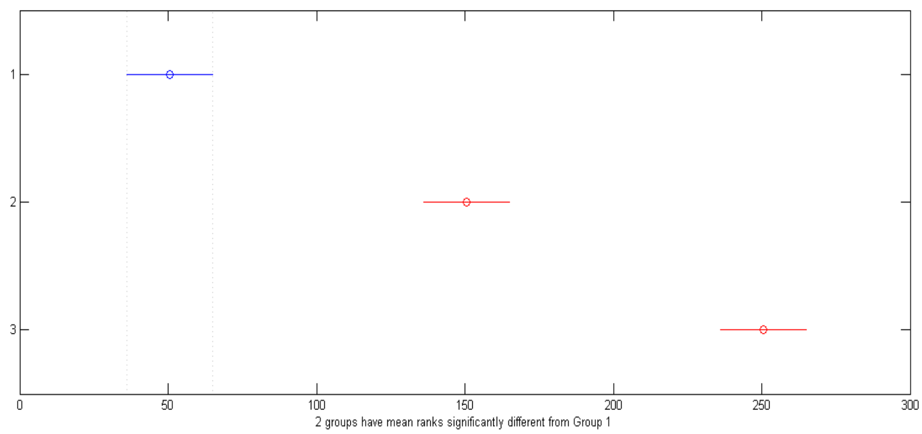

Figure 9: Results for multiple comparison procedure (Hochberg and Tamhane, 1987) using the calculated rank statistics of the Kruskal Wallis tests. The figure illustrates that the PCA imputation method provides mean ranks that are significantly different (smaller) from the

two benchmark methods, i.e. that the approach provides MAEs and MSEs that are significantly smaller.

Overall, the conducted tests provide strong evidence for a superior performance of the PCA imputation method in comparison to the two benchmark methods, namely PC regression and a univariate regression model.

5.6 Conclusion

Clean and complete historical data is a prerequisite for risk managers to generate accurate estimation for the back-testing of risk models, trading strategies or the calculation of Value-at-Risk. In this chapter, an iterative PCA-based data imputation algorithm is proposed for handling missing values in financial time series. The designed backfilling algorithm generates good results for both synthetic data and real data, covering equity, currency rates and FX asset classes. Further the proposed model outperforms two other commonly used approaches for data imputation. It is worth noting that the performance of the model depends on the fraction of the missing data and the noise of the data set. The higher the fraction of missing values, naturally the poorer is the quality for the prediction of missing values. To reduce the impact of noise, first we can exclude less relevant time series through correlation based filtering, and second we can apply statistical tools, for example, an iterative regularized PCA imputation algorithm to overcome overfitting. We can also apply cross-validation to select the optimal number of components in the applied PCA.

One of the potential limitations of the proposed model is that it relies on the correlation structure of the multivariate time series and hence may fail to appropriately backfill the missing data that are actually a result of outliers. However, to the best of our knowledge, we are not aware of any backfilling method that can provide appropriate and reliable estimates for missing financial data that were the result of outliers. Thus, it can be considered as particularly difficult to generate imputation methods for missing outlier data, unless substantial additional information about the nature of the missing observations and reasons why they are missing is available.

Another concern that is probably relevant for all imputation methods are the economic consequences of errors as a result of missing value imputation. Due to the existence of an imputation error and bias – the latter measured by the average deviation of the estimated values from the (unobservable) true values – there is almost certainly an economic cost involved with backfilling missing data. The actual economic costs of this error and bias due to imputation are very difficult to estimate and depend on what the data will be used for. Possible use of the data could be made in model estimation, pricing, risk reporting, creating hedging strategies, etc. However, overall the general rule is that the smaller the imputation error and the bias, the lower will be the economic costs of backfilling missing data.

Therefore, with regards to the issue, the proposed model can be considered as a great improvement in comparison to the rather too simplistic approaches currently being used by many practitioners. For example, approaches like proxying missing values using another time series by discretion, taking the average of the most adjacent available data for the missing value, etc., are still widely being used in practice. In contrast, the proposed method makes use of the embedded correlation structure of multivariate time series and infers the missing value through PCA. Therefore we believe the improvement on having more appropriate estimates for missing values in comparison to many simplistic approaches used by practitioners will provide also advantages with regards to economic costs.

Future work can be extended to research on backfilling large portions of missing data, which are commonly observed in financial time series, particularly for asset classes related to credit products. While the proposed model shows superior performance in comparison to the benchmark competitors further empirical studies, using data for different asset classes could be applied in future research.

Further, the development of an automated algorithm for the proposed backfilling framework would also be highly desirable. Such an automated algorithm could also be integrated with identifying meaningless or unlikely data points in financial time series, such as, e.g., negative FX spot rates, negative FX option volatility, and a sudden spike of FX volatility in a normal market environment.

References

Alexander, C., 2001. Market models: A guide to financial data analysis. John Wiley & Sons, Chichester.

Alexander, C., 2009. Market Risk Analysis, Volume II: Practical Financial Econometrics. John Wiley & Sons, Chichester.

Alvera-Azcarate, A., Barth, A., Rixen, M., Beckers, J. M., 2005.Reconstruction of incomplete oceanographic data sets using empirical orthogonal functions: applications to the Adriatic Sea surface temperature, Ocean Modelling, 9, 325–346.

Arteaga, F., Ferrer, A., 2002. Dealing with missing data in MSPC: several methods, different interpretations, some examples. Journal of Chemometrics 2002; 16:408–418.

Beckers, J., Rixen, M., 2003. EOF calculations and data filling from incomplete oceanographic data sets. Journal of Atmospheric and Oceanic Technology 20, 1839-1856.

Breiman, L., 2001. Random forest. Machine Learning, 45, 5-32.

Grung, B., Manne, R., 1998. Missing values in principal component analysis. Chemometrics and Intelligent Laboratory Systems 42(1-2):125–139.

Hellem,T., Dysvik,B., and Jonassen,I., 2004. LSimpute: accurate estimation of missing values in microarray data with least squares methods, Nucleic Acids Research, 32(1).

Hochberg, Y., Tamhane, A.C., 1987. Multiple Comparison Procedures. Wiley, Hoboken, NJ.

Hollander, M., Wolfe, D.A., 1999. Nonparametric Statistical Methods. Wiley, Hoboken, NJ.

Hron, K., Templ, M., Filzmoser, P., 2010. Imputation of missing values for compositional data using classical and robust methods. Computational Statistics & Data Analysis 54 (12), 3095-3107.

Jolliffe, I.T., 2002. Principal component analysis. Springer.

Jorion, P., 2000. Value at risk: the new benchmark for controlling derivatives risk. McGraw- Hill, New York.

Josse, J., Pages, J., Husson F., 2011. Multiple imputation in principal component analysis. Advances in Data Analysis and Classification 5, 231-246.

Josse, J., Husson, F., 2012. Selecting the number of components in principal component analysis using cross-validation approximations. Computational Statistics & Data Analysis 56(6), 1869-1879.

Karelmo, K., 2010. Interest Rate Term Structure Modelling in the Presence of Missing Data. Master Thesis.

Kim, J., Pachepsky, Y.A., 2010. Reconstructing missing daily precipitation data using regression trees and artificial neural networks for SWT streamflow simulation. Journal of Hydrology, 394, 305-314.

Kondrashov, D., Ghil, M., 2006. Spatio-temporal filling of missing points in geophysical data sets. Nonlinear Processes in Geophysics, 13(2), 151-159.

Little, R.J.A., Rubin, D.B., 1987. Statistical Analysis with Missing Data. New York: Wiley.

Mailleta, B., Merlinb, P., 2009. Hedge Fund Time-series Completion and Scenario Generation for Robust Asset Allocation and Risk Measurement. Association Française de Science Economique Congress.

Marivate, V. N., Nelwamondo, F. V. , Marwala, T.,2007. Autoencoder, principal component analysis and support vector regression for data imputation. Proceedings of the International Federation of Automation and Control Conference 2008, 2007.

Minsky, B., Obradovic, M., Tang, Q., Thapar, R., 2010. Applying a Global Optimisation Algorithm to Fund of Hedge Funds Portfolio Optimisation. Proceedings of the World Congress on Engineering 2010, London, U.K.

Nourani, V., Mogaddam, A.A., Nadiri, A.O., 2008. An ANN-based model for spatiotemporal groundwater level forecasting. Hydrological Processes, 22, 5054-5066.

O'Kelly,M., Ratitch, B., 2014. Clinical Trials with Missing Data: A Guide for Practitioners, John Wiley & Sons, Chichester.

Rodwell, L., Lee, K.J., Romaniuk, H., Carlin, J.B., 2014. Comparison of methods for imputing limited-range variables: a simulation study, Medical Research Methodology, 14:57.

Rustum, R., Adeloye, A.J., 2007. Replacing outliers and missing values from activated sludge data using Kohonen Self-Organizing Map. Journal of Environmental Engineering, 14, 853-871.

Schafer, J.L., 1997. Analysis of Incomplete Multivariate Data. London: Chapman and Hall.

Schneider, T., 2001. Analysis of incomplete climate data: Estimation of mean values and covariance matrices and imputation of missing values. Journal of Climate, 14, 853–871.

Stanimirova, I., Daszykowski, M., Walczak, B., 2007. Dealing with missing values and outliers in principal component analysis. Talanta, 72(1), 172–178.

Xia, Y., Fabian, P., Stohl, A., Winterhalter, M., 1999. Forest climatology: estimation of missing values for Bavaria, Germany. Agricultural and Forest Meteorology, 96,131-144.

Xia, Y., Fabian, P., Winterhalter, M., Zhao, M., 2001. Forest climatology: estimation and use of daily climatological data for Bavaria, Germany. Agricultural and Forest Meteorology, 106, 87–103.

Appendix A: Plots for CNY rates imputation

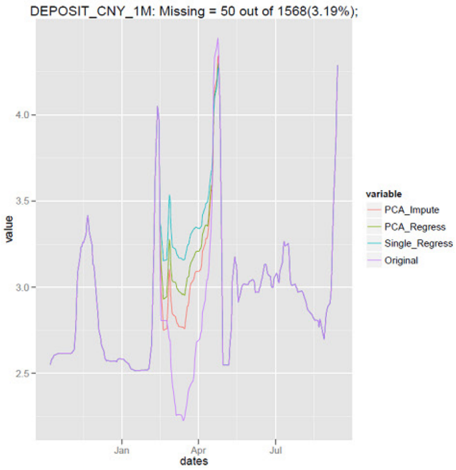

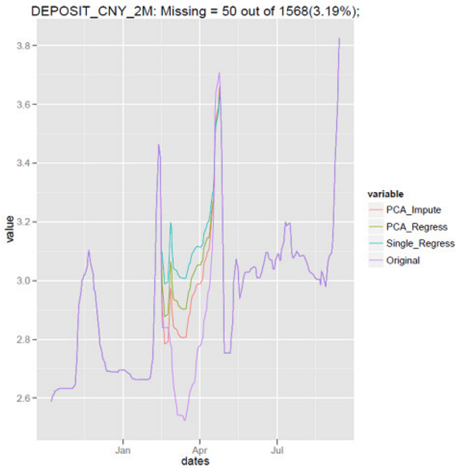

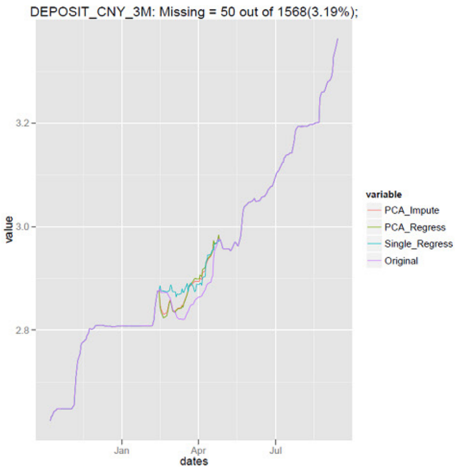

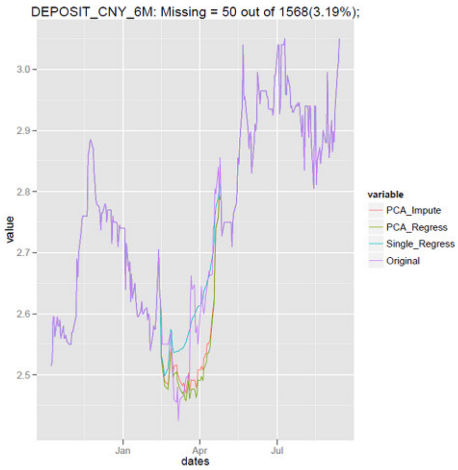

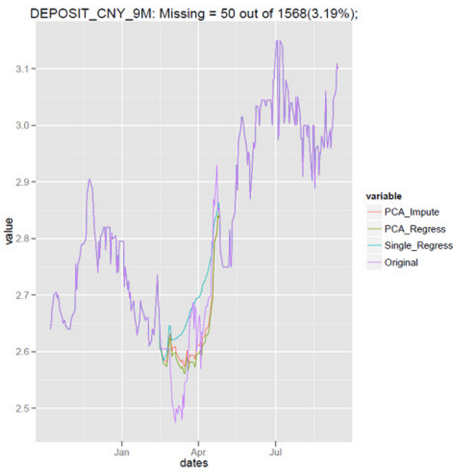

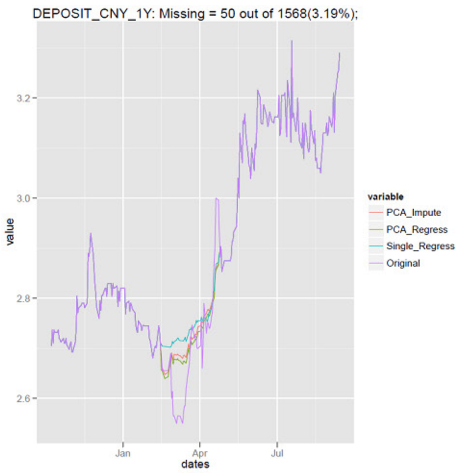

Note that purple line is the time series of original CNY deposit and swap rates at various tenors at length of 1568, in which 50 consecutive points (accounting for 3.19%) are removed for imputation. The lines in red, green and light blue represent the backfilled results using the proposed Iterative PCA Imputation and two other competing methods (PC Regression and Univariate Linear Regression) respectively.

Appendix B: Plots for Heng Seng Index option volatility surface imputation

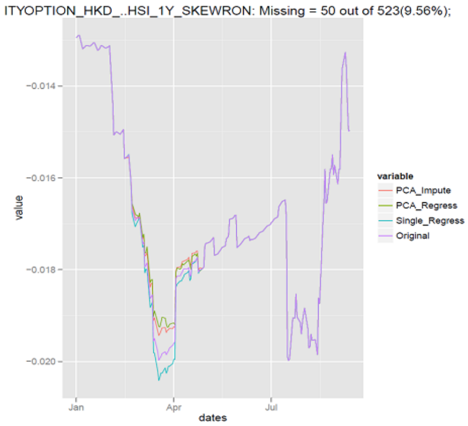

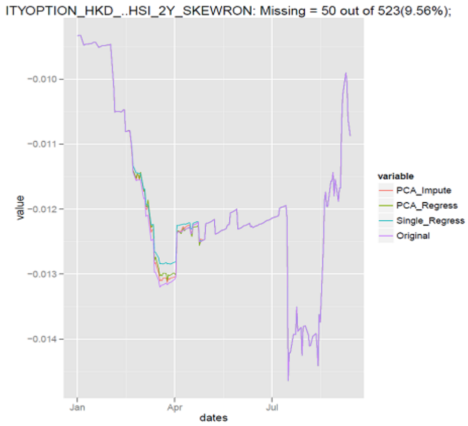

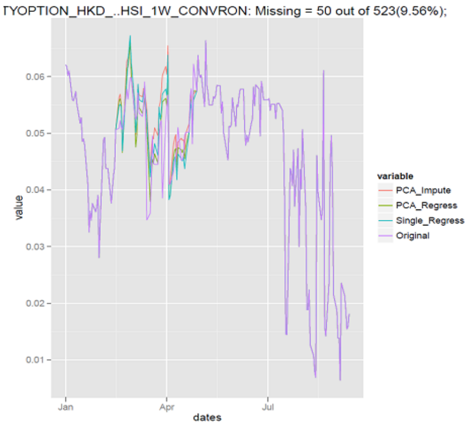

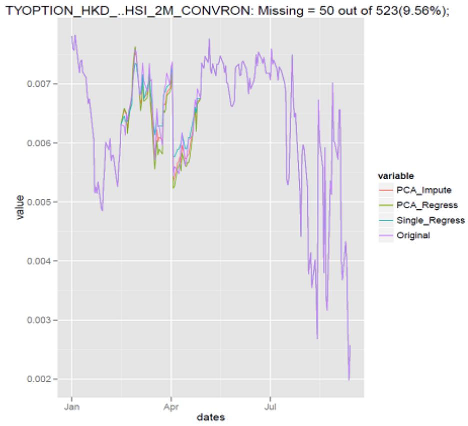

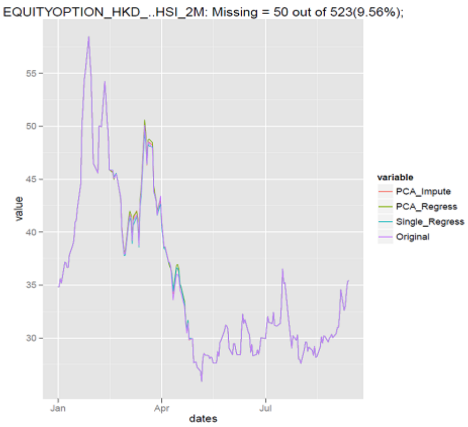

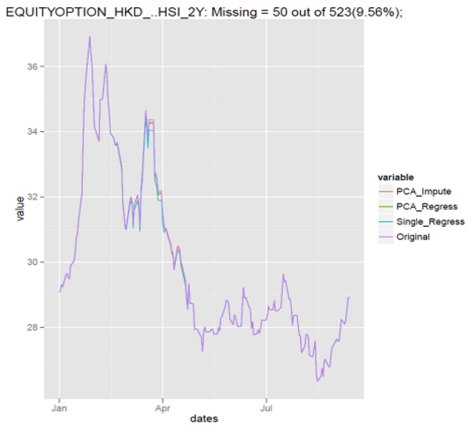

Note that purple line is the time series of original Hong Kong Heng Seng Index Option volatility surface, expressing in terms of ATM volatility, skewness (SKEWRON) and convexity (CONVRON), at length of 523, in which 50 consecutive points (accounting for 9.56%) are removed for imputation. The lines in red, green and light blue represent the backfilled results using the proposed Iterative PCA Imputation and two other competing methods (PC Regression and Univariate Linear Regression) respectively.

Appendix C: Plots for EURGBP FX option volatility surface imputation

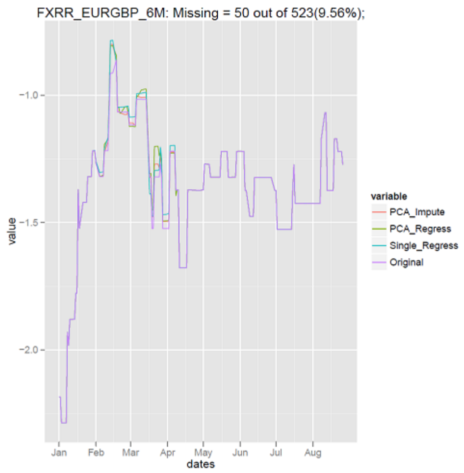

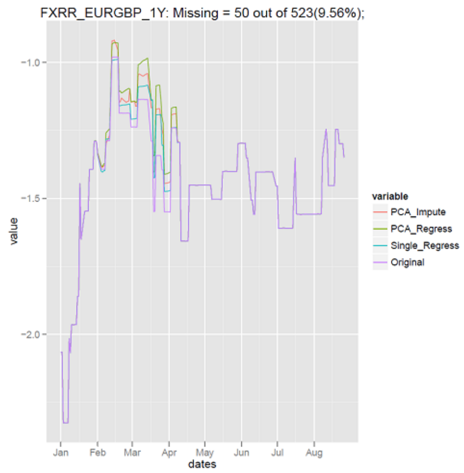

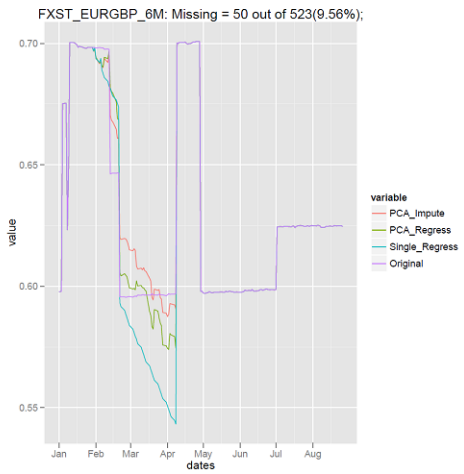

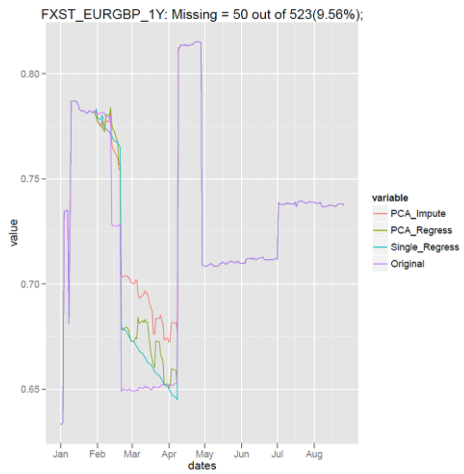

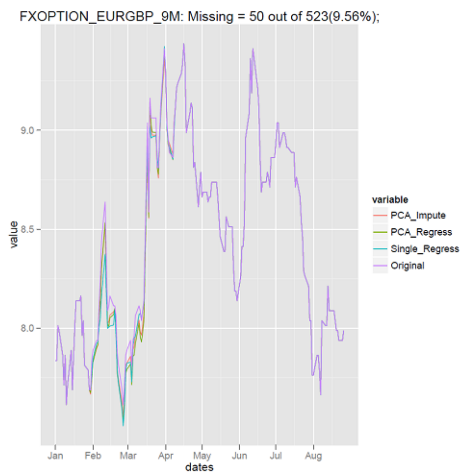

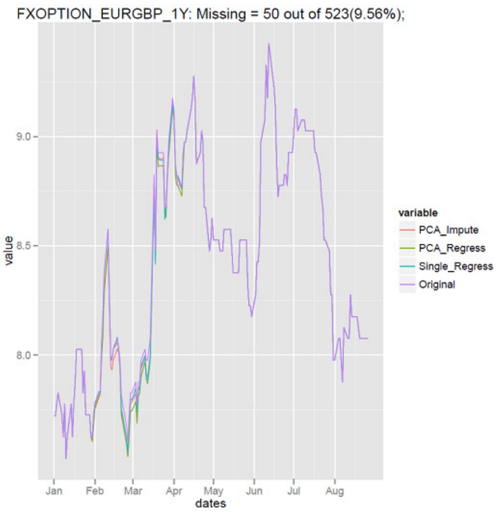

Note that purple line is the time series of original EURGBP FX Option volatility surface, expressing in terms of at-the-money volatility (FXOPTION), the risk reversal (FXRR) and the straddle (FXST), at length of 523, in which 50 consecutive points (accounting for 9.56%) are removed for imputation. The lines in red, green and light blue represent the backfilled results using the proposed Iterative PCA Imputation and two other competing methods (PC Regression and Univariate Linear Regression) respectively.

6. Summary and Conclusions

This PhD thesis addresses the several key topics in financial risk management and fund management. In particular, three research areas are selected for this study: risk management for Asia-focused hedge funds, fee structure in fund management, and ensuring the quality of historical data for financial risk management. The three topics are closely related. Risk management for hedge funds and manger's fee structures have attracted more attention from both academia and industry in the aftermath of the global financial crisis (GFC). One reason for this is that during the GFC, many managed funds posted dramatic losses. Even hedge funds, which claim to be capable of achieving a positive return on investment regardless of the market situation, also suffered severely from the GFC. With poor performance during the GFC, the manager's fee structure has come under scrutiny as well.

Therefore, risk management for managed funds, particularly for hedge funds and manager's fee structure, deserve to be well studied. In risk management, backfilling missing historical data plays a crucial role, as incomplete historical data adversely affects the accuracy and the reliability of key risk figures, such as back-testing results, Value-at-Risk, etc. Three chapters are thus devoted to shed light on these topics. This section aims to summarise and highlight the major contributions of each of the chapters in the thesis.

In lights of the above research fields worthy to explore, we firstly address topics in risk management for managed fund, with a specific focus on Asian hedged funds, in Chapter 3, which is devoted to identifying the risk factors contributing to the Asian hedge fund performance and proposing a measure to quantify the market risk of hedge funds. Next, in Chapter 4, we investigate the fee structure of managed funds with the aid of agency theory and shed light on actual contracts between investors, financial planners, licensees and product providers in Australia. Finally, in Chapter 5, we contribute to the literature of risk management by developing a robust method to backfill the missing financial data.

Chapter 3, titled "*Style Analysis and Value-at-Risk of Asia-Focused Hedge Funds*", identifies style factors for Asia-focused hedge funds represented by the HFRI Emerging Market-Asia exclude Japan index. We make use of the style analysis framework initially suggested by Agarwal and Naik (2000) and Dor et al. (2003). We employ a two-step procedure proposed by Lobosco and Dibartolomeo (1997) to test for the significance of the considered style factors. A rolling window style analysis provides further insight into the dynamic structure of style factor weights and risk exposure. This is one of the first empirical studies to apply these techniques with a particular focus on the Asian hedge fund industry.

The empirical results show that the most significant equity factors relating to the HFRI Emerging Market-Asia exclude Japan index are emerging equity markets, especially emerging markets in Asia. The two factors representing global and Asian emerging markets together account for a weight of approximately 45% on average. Risk exposures are consistent with the investment objectives of the hedge fund strategy. With respect to the fund's exposure to bond markets, we find that Asia-focused hedge funds indicate positive exposure to cash and high credit rating bonds but negative exposure to world government and emerging market bonds. In general, these fixed income factors account for a weight of 45%. The rolling window style analysis captures the hedge fund managers' style drift in responding to dynamic trading and changing market situations. For both static and rolling period style analysis, our model provides a high explanatory power for returns of the hedge fund index.

We further conduct an extensive analysis with respect to the ability of the models to provide appropriate forecasts for volatility and Value-at-Risk of the index. We use identified factors and factor weights of the rolling window style analysis in combination with a multivariate GARCH, moving average or exponentially weighted moving average (EWMA) model. The results are also compared to an approach that applies a univariate EWMA and GARCH model directly to the index returns. With respect to volatility forecasting, the models are compared based on a set of different loss functions. We find that none of the models performs best for all of the considered loss functions or significantly outperforms all of the other models. Nonetheless, the best results are obtained for three of the considered models: the EWMA and GARCH model using the actually observed returns of the hedge fund index as well as a model using the estimated style factor weights in combination with an EWMA scheme for the volatility.

In a second step, based on hypothesis tests for unconditional and conditional coverage, we further evaluate the performance of the considered models with respect to VaR estimation. We also apply different assumptions for the return distribution. Finally, the magnitude of the observed VaR exceptions is compared to those implied by the estimated VaR models. Our results indicate that the accuracy of the VaR models is dominated by their ability to capture the tail distribution of the hedge fund returns. Moreover, the performance of the models in VaR prediction seems to be more heavily dependent on the distributional assumption for the returns than on the chosen approach for volatility modelling: all models assuming a Student t distribution for the returns of the hedge fund index are significantly better than their counterparts assuming a Gaussian distribution. Overall, the best models for VaR estimation are a GARCH BEKK model based on the underlying style factors and a GARCH model that is based on the hedge fund returns only. Our findings further suggest that, in VaR forecasting, all parametric models outperform a simple historical simulation approach being purely based on past return observations. Finally, all of the considered VaR models perform reasonably well in forecasting the magnitude of the loss, conditional on a VaR exception.

Overall, our findings suggest that style analysis in combination with an appropriate parametric model for the identified factors provides an appropriate quantification of the risk for the considered Asian hedge fund index. We also find that multivariate models based on identified style factors and style weights significantly outperform a historical simulation approach with respect to volatility or VaR forecasting. On the other hand, our analysis indicates that they do not necessarily outperform simpler models like a univariate GARCH or EWMA model being directly applied to the hedge fund return series. However, they provide important insights on the exposures and investment style of a fund and indicate how fund returns can be replicated by observable market factors. In a time-varying setting, style analysis also provides information on how the weights of the different style factors potentially change through time as a reaction to different market conditions. Finally, style analysis might be useful for risk management when only a short period of observations is available for the fund itself while the identified style factors provide a much longer history that can be employed for estimating VaR or other risk measures. Therefore, we believe that style analysis approach should also be of particular help when individual hedge funds with a short track record are analysed and the use of hedge fund returns only for risk analysis will fail due to the lack of historical data. This issue should be thoroughly investigated in future research.

Chapter 4, titled "*Agency Theory and Financial Planning Practice*" augments the model of Dybvig et al. (2010) with a generalised log utility function, in conjunction with other theoretical contributions, to shed light on actual contracts between investors, financial planners, licensees and product providers in Australia. To our best knowledge, it is the first attempt to do so.

From our results, the optimal contract between an investor and a fund manager whose effort level cannot be verified by the investor first carves out the total protected wealth of the investor and the manager. It then subjects the remaining wealth of the investor to a fee structure with a flat component and two asset-based components. One asset fee is a standard proportional fee on fund earnings. The other is a symmetrical fulcrum-style performance fee.

To the extent that investors on the cusp of retirement have concerns about their wealth falling short of some predetermined value, they should simply allocate their wealth partly to safe interest-bearing assets before contracting with the active manager/financial planner. The amount of money exposed to risk by an active manager should be less than the entire wealth that can be invested by a client, especially in the case of investors on the cusp of retirement.

In practice, however, planners are entrusted with the bulk of the superannuation balances of their clients and derive most of their income from asset-based fees. By contrast, our theory suggests that asset-based fees on actively managed funds should include a fulcrum component, contrary to current practice.

Chapter 5, titled "*Backfilling Financial Data with an Iterative PCA-based Imputation*" proposes an iterative PCA-based data imputation for handling missing values in financial time series. The designed backfilling algorithm generates satisfactory results for both synthetic data and real data, covering equity, rates and FX asset classes. Our proposed model outperforms two other commonly used approaches for data imputation. It is worth noting that the performance of the model depends on the fraction of the missing data and the noise of the data set. The higher the fraction of missing value, the poorer the quality of prediction of missing values. To reduce the impact of noise, first we can exclude less relevant time series from PCA through correlation based filtering, and second we can apply statistical tools, for example, iterative regularised PCA imputation, to overcome over-fitting, and cross-validation to select the optimal number of components in PCA. Future work can be extended to research on backfilling large portions of missing data, which are commonly observed in financial time series, particularly for credit asset class.

7. References

Ackermann, C., McEnally, R., Ravenscraft, D., 1999. The performance of hedge funds: risk, return, and incentives. Journal of Finance 54(3), 833-874.

Agarwal, V., Naik, N.Y., 2000. Generalised style analysis of hedge funds. Journal of Asset Management 1, 93-109.

Agarwal, V., Naik, N.Y., 2004. Risks and portfolio decisions involving hedge funds. Review of Financial Studies 17(1), 63-69.

Alexander, C., 2001. Market models: A guide to financial data analysis. John Wiley & Sons, Chichester.

Alexander, C., 2009. Market Risk Analysis, Volume II: Practical Financial Econometrics. John Wiley & Sons, Chichester.

Amenc, N., Martellini, L., Vaissié, M., 2002. Benefits and risks of alternative investment strategies. Journal of Asset Management 4(2), 96-118.

Australian Prudential Regulation Authority 2012, Statistics: Annual Superannuation Bulletin, viewed June 2013, http://www.apra.gov.au/Super/Publications/Pages/annual-superannuation-publication.aspx.

Australian Securities and Investments Commission 2013, Regulatory Guide 246: Conflicted Remuneration, viewed June 2013, <http://www.asic.gov.au/asic/pdflib.nsf/LookupByFileName/rg246-published-4-March-2013-B.pdf/$file/rg246-published-4-March-2013-B.pdf.>

Bateman, H., Kingston, G., and Thorp, S. 2007, 'Financial Engineering for Australian Annuitants', in H. Bateman (ed.), Retirement and Scary Markets, Edward Elgar, Cheltenham, pp. 124-148.

Bauwens, L., Laurent, S., Rombouts, J.V.K., 2006. Multivariate GARCH models: a survey. Journal of Applied Econometrics 21, 79-109.

Bengen, W. 2001, 'Conserving Client Portfolios during Retirement, Part IV', Journal of Financial Planning (May issue), pp. 110-119, viewed June 2013, <http://www.fpanet.org/>.

Berkowitz, J., 2001. Testing density forecasts, with applications to risk management. Journal of Business and Economic Statistics 19(4), 465-474.

Bollen, N.P.B., Robert, E.W., 2008. Hedge fund risk dynamics: Implications for performance appraisal. Journal of Finance 64, 985-1035.

Bollerslev, T., Engle, R.F., Nelson, D.B., 1994. ARCH models, in Engle, R.F. and McFadden, D., eds. The Handbook of Econometrics Volume 4, 2959-3038. Amsterdam: North-Holland.

Bollerslev, T., Ghysels, E., 1996. Periodic autoregressive conditional heteroskedasticity. Journal of Business and Economic Statistics 14, 139-184.

Brown, S.J., Goetzmann, W.N., Ibbotson, R.G., 1999. Offshore hedge funds: survival and performance 1989-1995. Journal of Business 72(1), 91-118.

Christoffersen, P.F., 1998. Evaluating interval forecasts. International Economic Review 39, 841-862.

Christoffersen, P.F., Diebold, F., 2000. How relevant is volatility forecasting for financial risk management. Review of Economics and Statistics 82, 12–22.

Cochrane, J. 2012, ‘Hedge Funds’ (slides), Booth School of Business, University of Chicago, viewed June 2013, <http://faculty.chicagobooth.edu/john.cochrane/.>

Crnkovic, C., Drachman, J., 1996. A universal tool to discriminate among risk measurement techniques. Risk 9, 138–143.

Diebold, F.X., Mariano, R., 1995. Comparing predictive accuracy. Journal of Business and Economic Statistics 13, 253-264.

Ding, J., Kingston, G. and Purcal, S. forthcoming, ‘Dynamic Asset Allocation when Bequests are Luxury Goods’, Journal of Economic Dynamics and Control.

Dor, A.B., Jagannathan, R., Meier, I., 2003. Understanding mutual funds and hedge fund styles using return-based style analysis. Journal of Investment Management 1(1), 97-137.

Duffie, D., Pan, J., 1997. An overview of value at risk. Journal of Derivatives (Spring), 7-49.

Dybvig, P., Farnsworth, H., and Carpenter, J. 2010, ‘Portfolio Performance and Agency’, Review of Financial Studies, vol. 23, pp. 1-23.

Engle, R.F., Kroner, K.F., 1995. Multivariate simultaneous generalised GARCH. Econometric Theory 11, 122-150.

Financial Planning Association 2008, 'Simplifying Statements of Advice: FPA Example SOA', viewed June 2013, <http://www.scribd.com/doc/60713243/SOA-Transition-to-Retirement>.

Financial Planning Association 2009, 'Superannuation Review: Phase Two – Operation and Efficiency', viewed June 2013, <survey.fpa.asn.au/files/Policy09CooperPhase2submission.pdf>.

Fung, W., Hsieh, D.A., 1997. Empirical characteristics of dynamic trading strategies: the case of hedge funds. The Review of Financial Studies 10(2), 275-302.

Fung, W., Hsieh, D.A., 2000. Performance characteristics of hedge funds and commodity funds: natural vs. spurious biases. Journal of Quantitative and Financial Analysis 35(3), 291-307.

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

Fung, W., Hsieh, D.A., 2002. The Risk in fixed-income hedge fund styles. Journal of Fixed Income 12(2), 16-27.

Fung, W., Hsieh, D.A., 2003. The risk in equity long/short hedge funds. Working Paper, London Business School and Duke University.

Fung, W., Hsieh, D.A., 2004. Hedge fund benchmarks: a risk based approach. Financial Analyst Journal 60, 65-80.

Gibson, M.S., Boyer, B.H., 1998. Evaluating forecasts of correlation using option pricing. Journal of Derivatives, Winter, 18-38.

Glosten, L.R., Jagannathan, R., 1994. A contingent claim approach to performance evaluation. Journal of Empirical Finance 1, 133-160.

Grossman, S., and Stiglitz, J. 1976, 'Information and Competitive Price Systems', American Economic Review vol. 66, pp. 246–253.

Grung, B., Manne, R., 1998. Missing values in principal component analysis. Chemometrics and Intelligent Laboratory System 42(1-2):125-139.

Harmer, J. 2009, Pension Review Report, Commonwealth of Australia, Canberra.

Hull, J., White, A., 1998. Value at risk when daily changes in market variables are not normally distributed. Journal of Derivatives 5(3), 9-19.

Jolliffe, I.T., 2002. Principal component analysis. Springer.

Jorion, P., 2000. Value at risk: the new benchmark for controlling derivatives risk. McGraw- Hill, New York.

Jorion, P., 2001. Value-at-risk: the new benchmark for controlling market risk. Chicago: McGraw-Hill.

Josse, J., Husson, F., 2012. Selecting the number of components in principal component analysis using cross-validation approximations. Computational Statistics & Data Analysis 56(6), 1869-1879.

Josse, J., Pages, J., Husson F., 2011. Multiple imputation in principal component analysis. Advances in Data Analysis and Classification 5, 231-246.

JP Morgan, 1996. RiskMetrics technical document.

Karelmo, K., 2010. Interest Rate Term Structure Modelling in the Presence of Missing Data. Master Thesis.

Kingston, G. 2009, 'Financial Plans for Baby Boomers: How much Risk?', Economic Papers, vol. 28, pp. 65-74.

Kondrashov, D., Ghil, M., 2006. Spatio-temporal filling of missing points in geophysical data sets. Nonlinear Processes in Geophysics, 13(2), 151-159.

Kupiec, P.H., 1995. Techniques for verifying the accuracy of risk measurement models. Journal of Derivatives, winter, 73-84.

Liang, B., 1999. On the performance of hedge funds. Financial Analysts Journal 55(4), 72-85.

Little, R.J.A., Rubin, D.B., 1987. Statistical Analysis with Missing Data. New York: Wiley.

Lobosco, A., Dibartolomeo, D., 1997. Approximating the confidence intervals for Sharpe style weights. Financial Analysts Journal 53(4), 80-85.

Lopez, J.A., Walter, C.A., 2000. Evaluating covariance matrix forecasts in a Value-at-Risk framework. Federal Reserve Bank of San Francisco, working paper.

Loudon, G., Okunev, J., White, D., 2006. Hedge fund risk factors and the value at risk of fixed income trading strategies. Journal of Fixed Income 16(2), 46-61.

Mailleta, B., Merlinb, P., 2009. Hedge Fund Time-series Completion and Scenarii Generation for Robust Asset Allocation and Risk Measurement. Association Française de Science Economique Congress.

Minsky, B., Obradovic, M., Tang, Q., Thapar, R., 2010. Applying a Global Optimisation Algorithm to Fund of Hedge Funds Portfolio Optimisation. Proceedings of the World Congress on Engineering 2010, London, U.K.

Mitchell, M., Pulvino, T., 2001. Characteristics of risk and return in risk arbitrage. Journal of Finance 56 (6), 2135 – 2175.

Mitchell, O., and Utkus, S. 2012, 'Target-Date Funds in 401(k) Retirement Plans', working paper, Wharton School, University of Pennsylvania.

Rainmaker Group 2013, 'Financial Planning Report', viewed June 2013, <http://www.rainmaker.com.au>.

Rosenblatt, M., 1952. Remarks on a multivariate transformation. Annals of Mathematical Statistics 23, 470–472.

Ross, S. 1973, 'The Economic Theory of Agency: The Principal's Problem', American Economic Review vol. 62, pp. 134-139.

Russell Investments 2008, 'The Russell 10/30/60 Retirement Rule', viewed June 2013, <http://www.russell.com/CA_FP/PDF/Invest4Retir_Investr_Final.pdf>.

Schafer, J.L., 1997. Analysis of Incomplete Multivariate Data. London: Chapman and Hall.

Schneider, T., 2001. Analysis of incomplete climate data: Estimation of mean values and covariance matrices and imputation of missing values. J. Climate, 14, 853–871.

Sharpe, W.F., 1992. Asset allocation: management style and performance measurement. Journal of Portfolio Management 18(2), 7-19.

Stanimirova, I., Daszykowski, M., Walczak, B., 2007. Dealing with missing values and outliers in principal component analysis. Talanta, 72(1), 172–178.

Stoughton, N., Wu, Y. and Zechner, J. 2011, 'Intermediated Investment Management', Journal of Finance, vol. 66, pp. 947-980.

Teo, M., 2009. The Geography of hedge funds. The Review of Financial Studies

22(9), 3531-3561

Wachter, J. and Yogo, M. 2010, ‘Why do Household Portfolio Shares Rise in Wealth?’, Review of Financial Studies, vol. 23, pp. 3929-3965.

Whitelaw, T., O’Donnell, K. and Gray, P. 2011a, ‘Best Practices in Managed Fund Performance Fees’, unpublished paper, Morningstar.

Whitelaw, T., O’Donnell, K. and Gray, P. 2011b, ‘Active Share: The Activeness of Large-Cap Australian Share Fund Managers’, unpublished paper, Morningstar.