MACRO STRATEGY

Functional PCA for Implied Volatility Surface Prediction

May 19, 2020

Abstract

The volatility surface of an option changes over time, and values exist as discrete points on a grid, but the property of smoothness across points on the surface is evident. Common approaches to predicting points on an implied volatility surface include generalized autoregressive conditional heteroskedasticity (GARCH) and Vector Autoregression (VAR) models. These models involve modelling each point on the grid at a time and thus don't incorporate the smoothness typical to a volatility surface. To incorporate this effect, we propose to use a combination of functional data analysis and nonlinear regression modelling for predicting implied volatility while respecting the functional nature of the surface.

- Implied Volatility surfaces can be represented in significantly lower dimensions with little loss of information.
- The projection of implied surfaces to lower dimensions via functional principal component analysis (FPCA) respects the geometric nature of the surfaces while maintaining a minimal loss of information.
- The projection of implied volatility surfaces to lower dimensions via FPCA provides a framework for forecasting values of implied volatility by forecasting values of principal component scores.
- There is evidence of the predictability of the principal component scores of implied volatility surfaces.
- Using a standard forecasting method (VAR) for forecasting implied volatility results in lower out-of-sample error when using principal component scores as features when compared to the same method using all points on the surface.

MetLife Investment Management

1. Introduction

This is the second paper of our FX Implied Vol Forecasting series. In this paper, we implemented Functional Principle Component Analysis (FPCA) to reduce volatility surface dimensions, which helps to increase the accuracy of forecasting models. A volatility surface is a two-dimensional object representing the implied volatility (IV) of an option over a grid of deltas and expiries. The volatility surface of an option changes over time, and values exist as discrete points on a grid, but the property of smoothness across points on the surface is evident. Common approaches to predicting points on an implied volatility surface include GARCH models as discussed by Gong et. al [1] and VAR models as discussed by Ryu and Lee [2]. These models involve modelling each point on the grid at a time and thus don't incorporate the smoothness associated to a volatility surface. To incorporate this effect, we use a combination of functional data analysis and nonlinear regression modelling for predicting implied volatility while respecting the nature of the surface.

2. Previous Work

Standard Principal Component Analysis (PCA) for analysing volatility surfaces has been explored [3]. However, PCA analysis of time series data requires stationary (i.e., mean and variance are constant over time), which is not present in the values of implied volatility. Fengler et al [4], use a PCA approach to perform a joint PCA on implied volatility surfaces of different maturity buckets. The application of PCA to functional data for analysis of volatility surfaces is explored by Cont and Fonseca [5]. The authors give a complete illustration of numerical implementation that was greatly influential to this study. Perhaps most relevant to this study, however, is a study performed by Fengler et. al [6]. They also propose a FPCA approach for modelling the implied volatility surface, but also argue that VAR estimation of the factor loadings of the components is asymptotically equivalent to the estimation based on the unobserved coefficients of a basis expansion of the surface. The contribution of this study is in exploring the predictability of this structural model of IV surfaces, the incorporation of informative economic indicators in predictive models based on framework, and the assessment of nonlinear modelling of the factor loadings of the functional principal components.

3. Methodology

3.1. Data Description

The response data (source: Bloomberg) used in this study consisted of 75 points of the implied volatility surface for eight foreign exchange currencies from January 1, 2014 till June 11, 2019. The 75 points correspond to five values of delta: 10 Call, 25 Call, At the Money, 25 Put and 10 Put, and fifteen values of expiry: 1, 2, and 3 weeks, 1, 2, 3, 4, 6, and 9 months, 1, 1.5, 2, 3, 4 and 5 years out.

3.2. Smoothing the Surface

To represent a discrete grid of observed data points such as the one shown in Figure 1 in functional form such that we can perform FPCA, we fit a basis expansion.

Figure 1: Discrete 75-point IV Surface

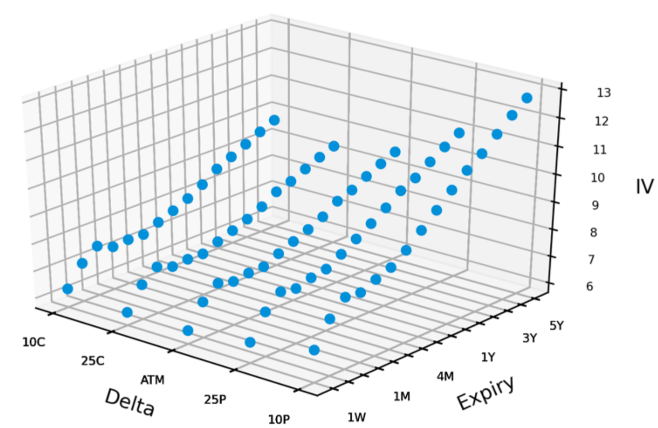

Data Source: Bloomberg

Notate the observed values of IV as Y_{tde} for time point t, delta d, and expiry e, and the unobserved smooth generating function of Y_{tde} as $Y_t(d,e)$. A basis $\{h_1, h_2, \dots, h_k\}$ is specified. Common choices for bases are polynomial, B-splines, and wavelets. In application, we use B-splines. With the model assumption

$$Y_t(d,e) = \sum_{i=1}^{k} c_{t,k} h_k(d,e) + e_t(d,e)$$

the coefficients $c_{t,k}$ are estimated via ordinary least squares. The smoothed surface of a 75-point volatility surface is given in Figure 2.

Figure 2: Fitted Basis Expansion of IV Surface

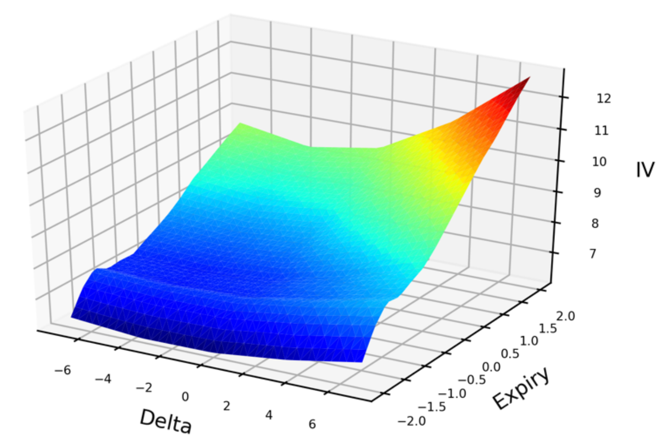

Data Source: Bloomberg

Figure 3 gives a view of the quality of fit of these expansions through the relative root mean square error (RMSE). The average relative RMSE over all the fitted curves is about 0.05%. With this level of error in basis expansion, we can be confident in the functional representation of the volatility surface data.

Figure 3: Relative RMSE (%) of Basis Expansion

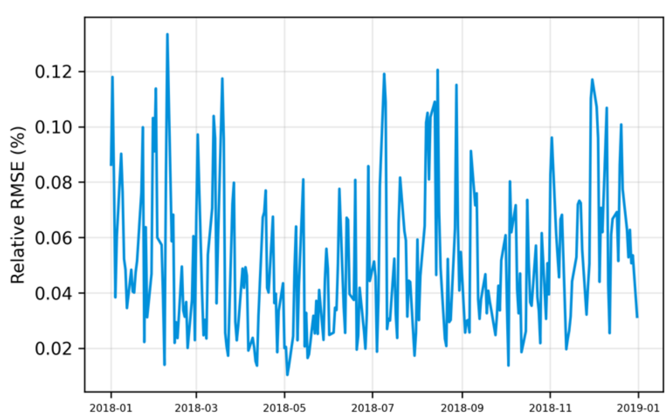

Data Source: MetLife Investment Management (MIM)

3.3. Functional PCA (FPCA)

FPCA is the functional analogue of standard PCA in multivariate statistical analysis and is useful in determining common signals or factors in the dynamics of underlying functions. To implement FPCA, we take the finite Karhunen-Loeve decomposition of the smoothed surfaces for which the kth basis function is the kth eigenfunction of the response, that is, the function that depicts the dominant mode of variation in the response. Notating the smoothed surfaces as $y_t(d, e)$ and concatenating the delta and expiry dimensions to obtain smooth functions $y_t(s)$, we obtain a mean function μ, K eigenfunctions $\phi_1, \phi_2, \dots, \phi_K$, each with corresponding eigenvalues $\lambda_1, \lambda_2, \dots, \lambda_K$, and principal component scores

$$\xi_{k,t} = \int \left(y_t(s) - \mu(s)\right)\phi_k(s)\, ds$$

In practice, these values are evaluated at points on the surface so that

$$\begin{aligned} \xi_{k,t} &= < y_{t,s} - \mu_s, \phi_{k,s} > \\ &= \sum_k (y_{t,s} - \mu_s)\phi_{k.s} \end{aligned}$$

The first three eigenfunctions of the AUDUSD volatility surface computed over the period of November 1, 2014 through June 11, 2019 are shown in Figure 4, and the associated components are shown in Figure 5. The first 95% of the functional variation in the IV surfaces of all pairs is explained by the first two components, and 99% is explained by the first three.

3.4. Prediction of Functional Principal Components

We use functional principal components (PCs) both as a target for volatility prediction as well as lagged principal components as features. To predict future values of PCs, we first select k, the number of principal components to be included in the prediction. K can be chosen to ensure a certain proportion of explained functional variance. Using an appropriate multivariate model, the h-step out predicted components $\hat{\xi}_{k,t+h}$ are then predicted. Then the predicted components are then mapped to predicted surface points as

$$\hat{y}_{t,d,e} = \hat{\mu} + \sum_{k=1}^{K} \hat{\xi}_{k,t}\phi_{k,d,e}$$

where $\hat{y}_{tde}$ is the predicted surface point at time t, delta d, and expiry e, ϕ_k is the kth eigenfunction, and $\hat{\xi}_k$ is the kth PC. Using PC scores as a target can be seen both as response dimension reduction and as a technique for respecting the inherent geometry of the implied volatility surfaces. As a dimension reduction technique, we find FPCA is highly effective in increasing degrees of freedom in any regression model in that it reduces the response dimension from 75 to 3 with only a 1% reduction in explained variance in the response, thus greatly reducing the risk of overfitting and allowing for greater signal sensitivity. As a feature, the lagged scores prove to be informative for all models for all targets and are used ubiquitously in the solution. We found that predicting functional PCs resulted in lower out-of-sample prediction error rates after mapping forecasted components to surface predictions than forecasting individual expiry/delta combinations. This result was seen from using both linear and nonlinear models for component prediction.

Figure 4: First Three Eigensurfaces

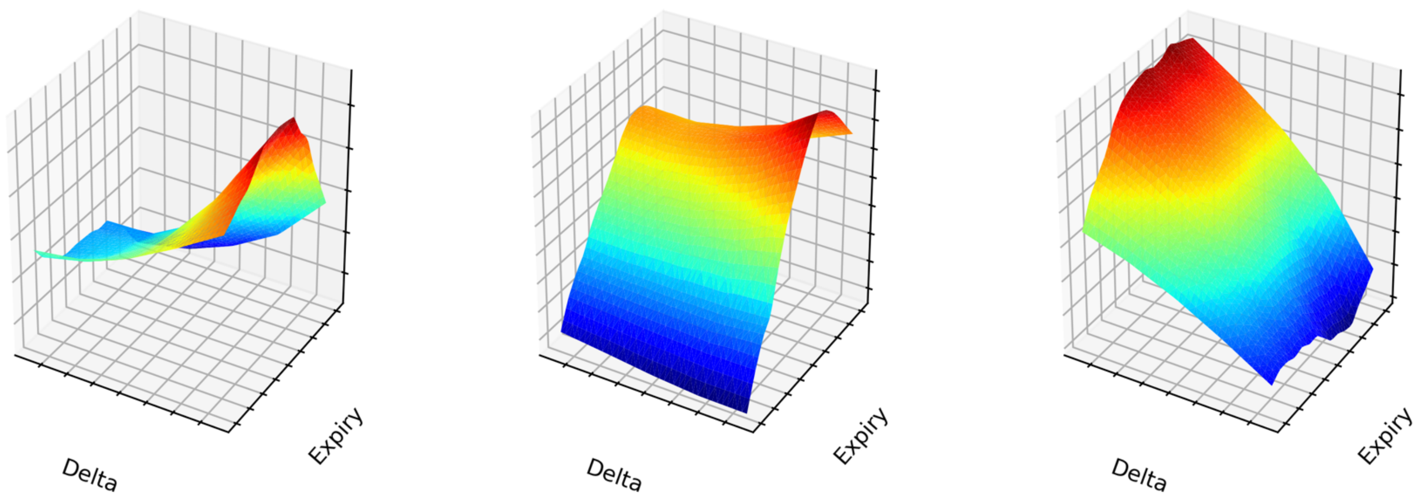

Data Source: MIM

Figure 5: First Three Principal Component Scores

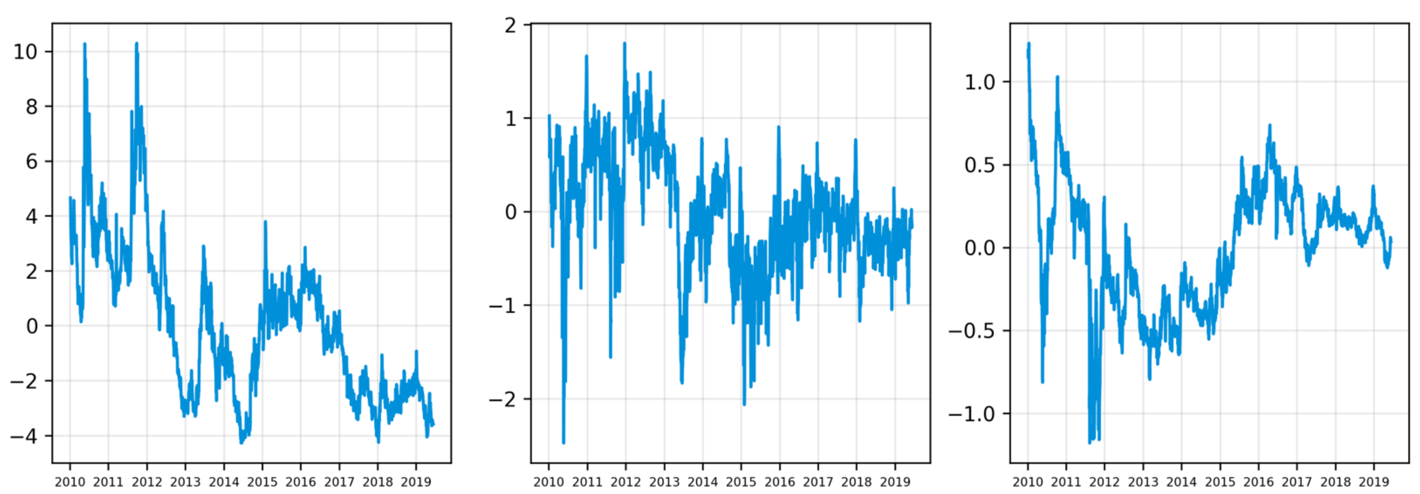

Data Source: MIM

4. Out-of-Sample Performance

We now observe the out-of-sample performance of multivariate models using functional PCs as a response versus using discrete points as a response. To establish a baseline model, we generate IV level forecasts from an order-5 Vector-Autoregression (VAR) model [1,2] using the (differenced) discrete IV surface points as a response.

The time frame of our dataset was January 1, 2014 to June 11, 2019. We took all data until January 1, 2018 as training data and iteratively fit the models and made one-day-out predictions over the remaining holdout data. In order to most cleraly show the impact of using functional PCs as responses compared to using discrete points, we use the same model, the order-5 VAR to predict the (differenced) functional PC values, and map the predicted PC values back surface points before calculating RMSE. The errors and the actuals against the predicted values for the 3M ATM IV for the AUDUSD are given in Figures 6 and 7, and the RMSE's of the one day out out-of-sample predictions across all values of expiries, deltas, and pairs are given in Figure 8. It is clear that the FPCA VAR model dominates the standard VAR model across all points on the surface.

Figure 6: One day out predictions and actuals for the AUDUSD 3M ATM Implied Volatility

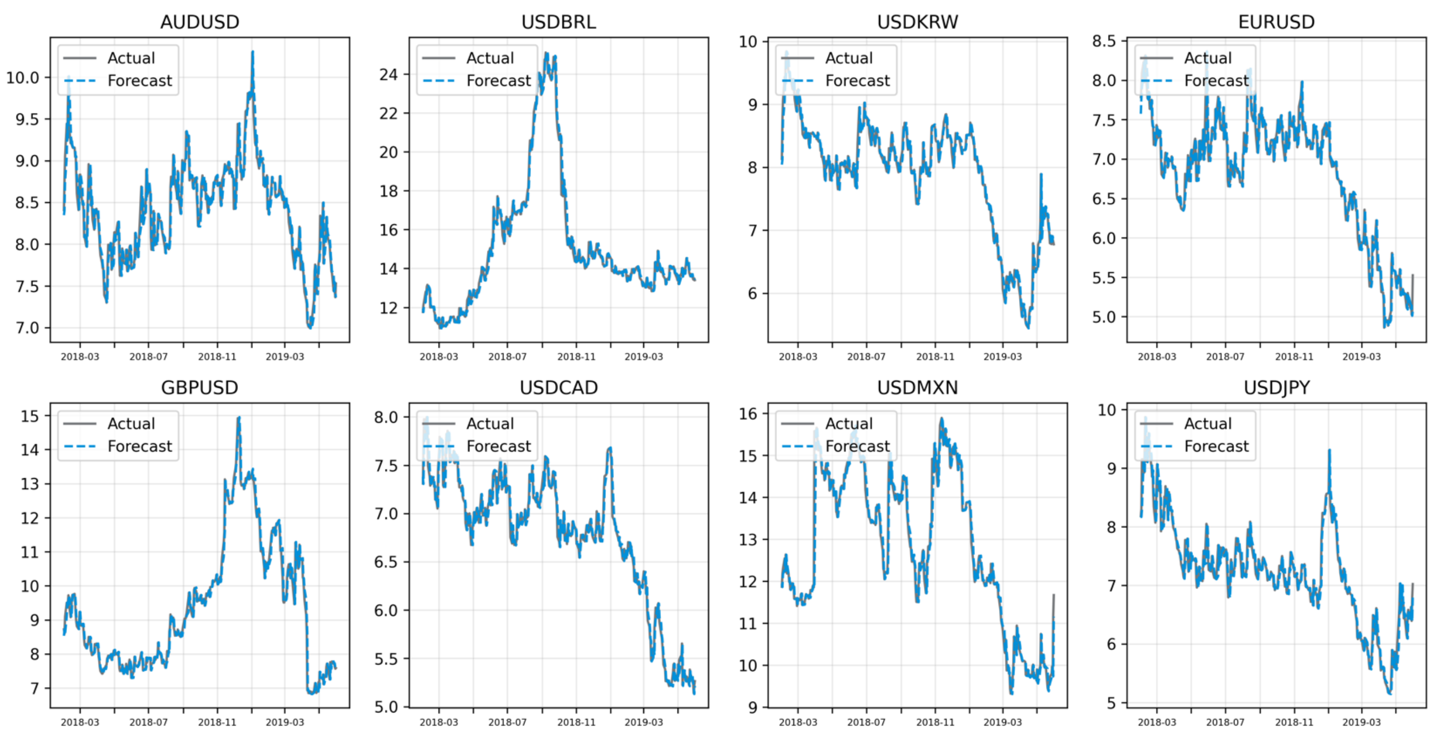

Data Source: MIM

Figure 7: Out of Sample Errors for the one day out predictions of the AUDUSD 3M ATM Implied Volatility

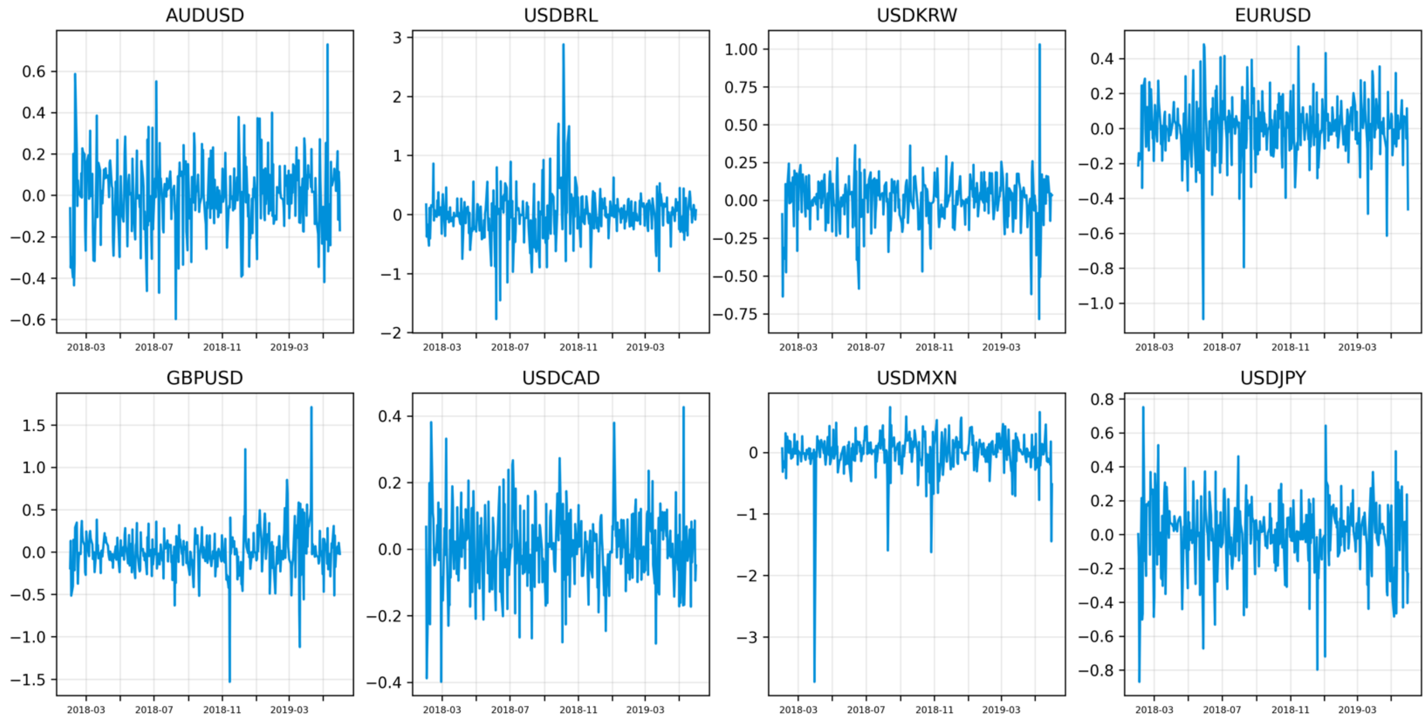

Data Source: MIM

Figure 8: Out of Sample RMSE

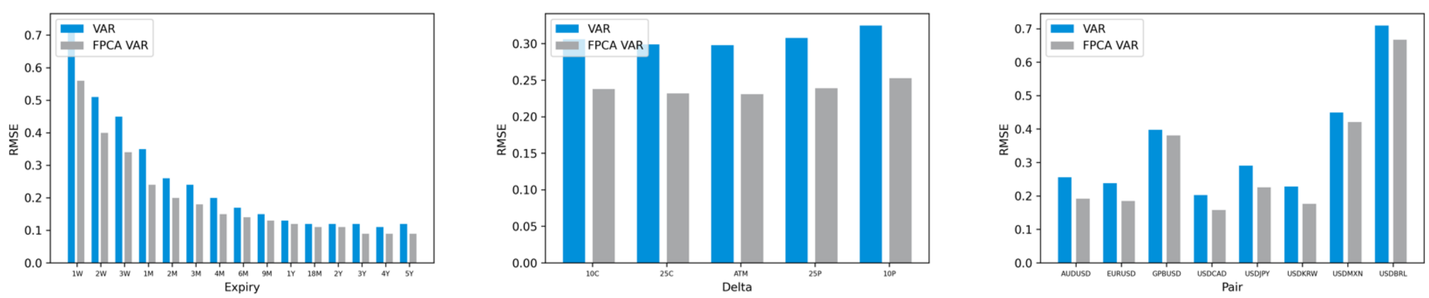

Data Source: MIM

5. Conclusions

Here we have performed a study using the implied volatility surfaces of eight foreign exchange currencies to develop and compare predictive models of IV surfaces using FPCA. We found clear evidence of the predictability of this dimension reduced target as well as evidence that the proposed models outperform a standard discrete point forecast.

References

[1] Gong, H., A. Thavaneswaran, and J. Singh. “A BLACK-SCHOLES MODEL WITH GARCH VOLATILITY.” Mathematical Scientist 35.1 (2010).

[2] Lee, Bong Soo, and Doojin Ryu. “Stock returns and implied volatility: A new VAR approach.” Economics: The Open-Access, Open-Assessment E-Journal 7.2013-3 (2013): 1-20.

[3] Skiadopoulos, George, Stewart Hodges, and Les Clewlow. “The dynamics of the S&P 500 implied volatility surface.” Review of derivatives research 3.3 (2000): 263-282.

[4] Fengler, Matthias R., Wolfgang K. Härdle, and Christophe Villa. “The dynamics of implied volatilities: A common principal components approach.” Review of Derivatives Research 6.3 (2003): 179-202.

[5] Cont, Rama, and José Da Fonseca. “Dynamics of implied volatility surfaces.” Quantitative finance 2.1 (2002): 45-60.

[6] Fengler, Matthias R., Wolfgang K. Härdle, and Enno Mammen. “A semiparametric factor model for implied volatility surface dynamics.” Journal of Financial Econometrics 5.2 (2007): 189-218.

[7] Ramsay, James O., and Bernard W. Silverman. Applied functional data analysis: methods and case studies. Springer, 2007

Authors

Caleb Weaver, Ben Perryman, Jun Jiang, Alex Chien, Erez Meoded, Al Kirk

About MetLife Investment Management

MetLife Investment Management (MIM),[1] MetLife, Inc.’s (MetLife’s) institutional investment management business, serves institutional investors by combining a client-centric approach with deep and long-established asset class expertise. Focused on managing Public Fixed Income, Private Capital and Real Estate assets, we aim to deliver strong, risk-adjusted returns by building tailored portfolio solutions. We listen first, strategize second, and collaborate constantly as we strive to meet clients’ long-term investment objectives. Leveraging the broader resources and 150-year history of the MetLife enterprise helps provide us with deep expertise in navigating ever changing markets. We are institutional, but far from typical.

For more information, visit: **metlife.com/investmentmanagement**

[1] MetLife Investment Management (“MIM”) is MetLife, Inc.’s institutional management business and the marketing name for the following affiliates that provide investment management services to MetLife’s general account, separate accounts and/or unaffiliated/third party investors: Metropolitan Life Insurance Company, MetLife Investment Management, LLC, MetLife Investment Management Limited, MetLife Investments Limited, MetLife Investments Asia Limited, MetLife Latin America Asesorias e Inversiones Limitada, MetLife Asset Management Corp. (Japan), and MIM I LLC.

Disclosure

This material is intended for institutional investor and professional investor use only. Not for use with retail public.

This document has been prepared by MetLife Investment Management ("MIM") solely for informational purposes and does not constitute a recommendation regarding any investments or the provision of any investment advice, or constitute or form part of any advertisement of, offer for sale or subscription of, solicitation or invitation of any offer or recommendation to purchase or subscribe for any securities or investment advisory services. The views expressed herein are solely those of MIM and do not necessarily reflect, nor are they necessarily consistent with, the views held by, or the forecasts utilized by, the entities within the MetLife enterprise that provide insurance products, annuities and employee benefit programs. The information and opinions presented or contained in this document are provided as the date it was written. It should be understood that subsequent developments may materially affect the information contained in this document, which none of MIM, its affiliates, advisors or representatives are under an obligation to update, revise or affirm. It is not MIM's intention to provide, and you may not rely on this document as providing, a recommendation with respect to any particular investment strategy or investment. The information provided herein is neither tax nor legal advice. Investors should speak to their tax professional for specific information regarding their tax situation. Investment involves risk including possible loss of principal. Affiliates of MIM may perform services for, solicit business from, hold long or short positions in, or otherwise be interested in the investments (including derivatives) of any company mentioned herein. This document may contain forward-looking statements, as well as predictions, projections and forecasts of the economy or economic trends of the markets, which are not necessarily indicative of the future. Any or all forward-looking statements, as well as those included in any other material discussed at the presentation, may turn out to be wrong.

L0520004078[exp0522][All States], L0520004017[exp0522][All States], L0520004011[exp1020][All States], L0520004009[exp0522][All States]

MetLife Investment Management

One MetLife Way | Whippany, New Jersey 07981