PARTICLE FILTERING FOR RISK MANAGEMENT AND REPLICATION OF A RISK APPETITE INDICATOR

by

Eric Lanoix
Master of Engineering (Mechanical), McGill University, 1999
Bachelor of Engineering (Honours Mechanical), McGill University, 1997

and

Jing Yuan
Bachelor of Management, Hebei University of Economics and Business, 2012

PROJECT SUBMITTED IN PARTIAL FULFILLMENT OF
THE REQUIREMENTS FOR THE DEGREE OF

MASTER OF SCIENCE IN FINANCE

In the Master of Science in Finance Program
of the
Faculty
of
Business Administration

SIMON FRASER UNIVERSITY

Fall 2013

Approval

Name: **Eric Lanoix, Jing Yuan**

Degree: **Master of Science in Finance**

Title of Project: **Particle Filtering for Risk Management and Replication of a Risk Appetite Indicator**

Supervisory Committee:

_____

Dr. Andrey Pavlov
Senior Supervisor
Professor of Finance

_____

Dr. Peter Klein
Second Reader
Professor of Finance

Date Approved: _____

Abstract

The authors propose a new approach to estimating stochastic volatility parameters. Traditional methods maximize the conditional likelihood. The proposed model optimizes two criteria: the deviation of the observed residuals PDF from the theoretical PDF and the in-sample predictive power of the volatility estimate. The resulting model yields better results than GARCH and than Harvey et al.'s stochastic volatility model. Two more applications of this innovation are also examined. First, volatility fit residuals for two assets are combined to estimate dynamic correlation. The model aptly estimates dynamic correlation when it is significant – though with some lag. Second, these models are used to replicate the RBC Risk Appetite Indicator. Results show that even though the authors are missing 40% of the inputs to the risk indicator, the replication strategy adequately replicates the indicator. We expect these three models to be of significant use to the SIAS team.

Keywords: Stochastic Volatility; Risk Appetite Indicator; Particle Filter; Dynamic Correlation; Residual Distribution Matching.

Dedication

Eric Lanoix's dedication:

A tous ceux qui ont relevé le grand defi, au milieu de leur carrière, de changer leur vie afin d'être plus heureux et d'assurer un meilleur avenir pour leurs enfants; parmi eux mes parents, Julien et Germaine Lanoix.

This work is dedicated to everyone who, in the middle of their career, challenged themselves to do what they truly wanted and went for it – and among them my father and mother, Julien and Germaine Lanoix.

I also dedicate my work to the SFU Student Investment Advisory Service (SIAS) mentors, advisors, clients, supporters, and team members – past and present. Thank you all for this tremendous learning opportunity. I tip my hat to all 2012-2013 SIAS team members for that great ride we shared on the markets. I hope these models can help future M. Sc Finance students provide value added for the Fund.

Acknowledgements

The authors wish to express their gratitude to the following individuals who contributed (from close or from afar) to the success of this research project: Andrey Pavlov for overall thesis project supervision, for suggestions, and for his flexibility; Peter Klein for secondary review and for suggesting a project on particle filtering in the first place; Anton Theunissen for innumerable discussions on finance in general, and for throwing the occasional intellectual *curve ball* that kept us thinking of how we could improve both the project topic and focus; Derek Yee for suggesting the construction or replication of a risk appetite indicator, Avi Bick for helping us understand option pricing theory and objective transition probabilities[1] far beyond the context of masters' level derivatives courses; Suzanne Yim and Brian Chand for their support; Carlos da Costa for his seasoned "*this is how the real world works*" advice; Jeongeun Kim for making the source code of her thesis (which was used for particle filtering and smoothing) available on the web; and to Eric Lascelles (chief economist at RBC) publishing the RBC Risk Appetite Indicator and describing it in enough detail to enable us to replicate it.

[1] This was the original topic of this thesis, but it eventually morphed into a particle filter idea…

Table of Contents

Glossary

AR	Autoregressive Model
ARCH	Autoregressive Conditional Heteroskedasticity
BE	Bank of England
BIS	Bank for International Settlements
BofA	Bank of America
CBOE	Chicago Board Options Exchange
CDS	Credit Default Swap
CSFB	Credit Suisse First Boston Risk-Appetite Index
EGARCH	Exponential General Autoregressive Conditional Heteroskedastic
EKF	Extended Kalman Filter
EMA	Exponential Moving Average
FRED	Federal Reserve Economic Data
GARCH	Generalized Auto Regressive Conditional Heteroskedasticity
GRAI	Global Risk-Appetite Index
IGARCH	Integrated Generalized Autoregressive Conditional Heteroskedasticity
IMF	International Monetary Fund
PDF	Probability distribution function
PF	Particle Filter
QE	Quantitative Easing

RAI	Risk Appetite Indicator
RBC	Royal Bank of Canada
RDM	Residuals Distribution Matching
SIAS	Student Investment Advisory Service
SIS	Sequential Importance Sampling
SV	Stochastic Volatility
TGARCH	Threshold GARCH model
UKF	Unscented Kalman Filter
VIX	CBOE Volatility Index

1: Introduction

"No pain no Spain! So no risk, no return…" - unattributed

1.1 Motivation and Research Objectives

Given that expected return and risk are the two factors that drive financial markets, the authors attempt to quantify risk in a way that is as accurate and efficient as possible. This research project has two main objectives:

1. Implement a Stochastic Volatility (SV) model using particle filtering, as well as a simple dynamic correlation model. These models have applications in risk management, option pricing (including the pricing of spread options), as well as capital markets expectations. In addition, the output of these models serves as input to replicate the RBC RAI;
2. Replicate the RBC Risk Appetite Indicator (RAI). This model has applications in risk management and capital markets expectations. As RBC publishes their RAI infrequently, index replication can provide more up-to-date situational awareness regarding the overall level of risk in financial markets.

Upon completion of this project, the models or data will be supplied to the SIAS team upon request.

1.2 Literature Review

1.2.1 Filtering

Estimation constitutes a fundamental problem of econometrics and engineering. It consists of determining the value of parameters and states (a.k.a. hidden variables) that optimally fit a set of noisy measurements. To compute these estimates, one uses a *best guess* of the dynamics of the system. This is known as the *state propagation* equation.

For example, to estimate the motion of a car, one would use Newton's Second Law as the state propagation equation. In finance, the state propagation equation for a volatility estimator could be, for example, a GARCH(1,1) process [Bollerslev, 1986]:

$$\sigma_t^2 = \alpha_0 + \alpha a_{t-1}^2 + \beta \sigma_{t-1}^2 \quad (1)$$

which can be rewritten as:

$$x_t = \alpha_0 + \alpha a_{t-1}^2 + \beta x_{t-1} \quad (2)$$

where x_t is the variance (square of volatility) at time t. In the literature, x_t is the standard designation given to the state variable. But as is often the case in most applications, the state variable is not directly observable; it can only be inferred from movements in underlying asset prices. To remedy this problem, estimation formulations also rely on a so-called *measurement equation*, which describes (in this example) how volatility influences asset prices. For the GARCH(1,1) process, the measurement equation takes the following form:

$$a_t = \sigma_t \epsilon_t \quad (3)$$

or

$$y_t = \sqrt{x_t} \epsilon_t \quad (4)$$

where the asset log-returns are denoted by y_t, which is the standard way of denoting measurements at time t. Together, the state propagation and the measurement equations (equations (2) and (4) above) are referred to the *state space equations* of the system. This notation is used throughout this thesis.

The estimate of the states is updated over time as noisy measurements become available. Optimal state estimation has become a field of its own. For a complete primer on the topic, the interested reader is referred to Simon [2006].

Famous for guiding Apollo spaceships to the Moon and back, the Kalman filter [Kalman, 1960] has found applications in finance. It is the optimal linear filter, regardless of the noise distribution[2]. However, linear filters do not handle financial time series well because of severe non-linearities. To mitigate this shortcoming, a few refinements have been made to Kalman filter over the years.

The Extended Kalman Filter (EKF) [Anderson and Moore, 1979] uses a first-order truncated

[2] Some publications maintain that the Kalman filter is the optimal linear filter for systems with Gaussian noise. This is incorrect [Simon 2006]. For any given data time series, some non-linear filters may perform better than the Kalman filter, but this filter is the optimal linear filter for any system.

Taylor's series expansion about the current state to linearize non-linear systems. It is the most popular filter for non-linear systems [Simon, 2006]. Though more robust than the Kalman filter against non-linearities, the EKF tends to underestimate the covariance of the states [van der Merwe et. al, 2000(a)].

The Unscented Kalman Filter (UKF) [Julier and Uhlmann, 1997] is based on a mapping, the unscented transformation. It generally handles non-linear systems better than the EKF. In particular, state covariance matrices computed using the UKF are much closer to the truth than those computed using the EKF. However, the UKF does not perform well for arbitrary (non-Gaussian) noise distributions [van der Merwe et. al. 2000(b)], such as those routinely encountered in finance.

Simultaneous robustness against non-linearity and arbitrary noise distributions requires Monte Carlo simulation. The Particle Filter (PF) [Gordon et. al., 1993] is a Bayesian approach that simulates the path of a large number of samples (particles, as they are also called). In terms of accuracy and robustness, this type of algorithm dominates Kalman-type filters [Simon, 2006]. However, this comes at the cost of increased computational time.

The basic particle filtering process follows [Simon, 2006]:

1. Generate samples from the probability distribution of the initial state;
2. For each sample (particle), propagate the state of the system (xk) one step forward in time. This yields the a-priori distribution of the particles for that next time step;
3. Using the measurement equation, compute the likelihood of each a-priori particle conditioned on the measurement (yk) at that next time step;
4. Scale the relative likelihoods based on the results of step 3;
5. Compute a-posteriori particles based on the relative scaled likelihoods. These particles follow the PDF of p(xk|yk);
6. Go back to step 2 above and repeat the process until the last measurement is processed.

1.2.2 Stochastic Volatility

Computing asset price volatility has critical applications in finance such as the calculation of value at risk, mean variance optimization, and pricing options[3]. Standard price diffusion theory

[3] The price of long-term options is usually based on the expected long-term volatility, which can be determined using implied volatility [Crack, 2012]. On the other hand, asset volatility is known to have significant auto-correlation. Therefore, shorter-term options can be priced using historical volatility, but one must also account for the jump diffusion premium.

states that returns are log-normally distributed:

$$S_{t+1} = S_t e^{(\mu - \frac{1}{2}\sigma^2)\Delta t + \sigma\epsilon_t\sqrt{\Delta t}} \tag{5}$$

In the above expression, the asset volatility is assumed constant. However, asset price volatility is rarely constant, though it tends to exhibit auto-correlation (days with high volatility tend to follow days of high volatility, and vice-versa [Tsay, 2010]). Since the development of the ARCH model by Nobel laureate Robert Engle [1982], the computation of time-varying (dynamic) volatility has flourished into a rich field of econometrics research. For example, Bollerslev [1986] introduced the popular GARCH model, which is routinely used to forecast volatility and option prices [Chatfield, 2004]. However, GARCH tends to overstate volatility. It also shows a tendency to overstate the duration of high volatility periods [Gonzalez et al., 2002]. To mitigate this problem, researchers have patched the basic GARCH formulation as best as they could, resulting in such formulations as IGARCH, EGARCH, TGARCH, etc. [Tsay, 2010]. One of the major problems with these formulations is the specification (or misspecification) of the state equation for volatility. Revisiting Equation (1), one notices that volatility is a deterministic function; once one knows the volatility at the previous time step and the innovation, then the volatility at the current time step is known with *certainty*. There is no allowance for any diffusion of the volatility process – no error term in the state equation. Such models are called *deterministic* volatility models.

In contrast, *Stochastic* Volatility (SV) models include a diffusion term not only in the asset price process (measurement equation), but also in the volatility process (state equation) [Chatfield, 2004]. The Heston [1993] model is a popular SV model. The present thesis uses another SV formulation [Harvey et al., 1994], which is obtained by shifting the mean drift in equation (5) to zero, by relaxing the constant volatility assumption, and by assuming a time step of 1:

$$y_t = ln\frac{S_{t+1}}{S_t} = \sigma_t\epsilon_t = \epsilon_{1t}e^{(x_t/2)} \tag{6}$$

where the state (x_t) is the log-variance of the asset price, which follows an AR(1) process:

$$x_t = ln(\sigma_t^2) = \gamma + \phi x_{t-1} + \epsilon_{2t} \tag{7}$$

and where ε_{1t} and ε_{2t} are normally distributed random variables with zero mean. The general version of this model allows the error terms to be correlated. This captures the empirically observed *leverage* effect whereby volatility is higher for large price drops than for large price

gains [Yu, 2005]. This research project neglects leverage effects.

1.2.3 Dynamic Correlation

This research project also seeks to determine the time-varying (dynamic) correlation across assets. The ideal way to tackle this would be to derive a multivariate stochastic volatility and correlation formulation. For example, Chib's approach [2006] does accomplish this modeling feat of arms, but at great cost: a very large number of parameters to estimate and slow speed.

We seek a simpler algorithm, which mirrors an approach examined by Engle [2002]. It consists of computing the dynamic correlation using a two-stage divide-and-conquer strategy:

1. Estimate the stochastic volatility of each asset separately using a univariate formulation;
2. Blend the residuals from step 1 to compute the dynamic correlation using an Exponential Moving Average (EMA).

As will be shown later, if one neglects cross-volatility effects, this requires the simultaneous optimization of three parameters using one set of measurements (asset returns). Opting for the simultaneous and dynamic estimation of the volatility and correlation of two assets would have brought about the dreaded *curse of dimensionality* - with at least nine parameters to estimate simultaneously using only two sets of measurements (returns of two securities). In light of this, one can see why the simpler approach is preferable.

1.2.4 Risk Appetite Indicators

Risk appetite reflects investors' willingness to hold risky assets. Changes in risk appetite can explain movements in capital markets that seem unrelated to the flow of economic and political news. A rising risk appetite implies that investors are willing to hold riskier assets. This in turn drives stock prices up.

Several indices have been created to measure market risk appetite. Tarashev, Tsatsaronis, and Karampatos [2003] developed a Risk Appetite Index (RAI) at the Bank for International Settlements (BIS). This indicator is computed as follows:

1. Use a GARCH model and historical data to predict the statistical distribution of future asset returns;
2. For the same assets, use option prices with different strikes to determine the implied volatilities and draw the volatility smile;

3. The volatility smile is mapped into a *subjective* probability distribution of future payoffs. The value of the index is the ratio of the left tails of the two distributions (i.e., the ratio of the statistical downside risk to the subjective downside risk).

The Gai and Vause [2004] RAI, developed at the Bank of England (BE), is similar to the BIS method, but uses the ratio of the full distributions instead.

The Kumar and Persaud [2002] Global Risk-Appetite Index (GRAI) is constructed by first ranking assets by their riskiness (variance of past returns), and then by their excess returns (difference between futures and spot prices measured at a single point in time). This indicator is used as an input for both the IMF and JPMorgan RAIs.

The Credit Suisse First Boston Risk-Appetite Index (CSFB) [Wilmot, Mielczarski, and Sweeney 2004] is similar to the GRAI in that it compares risk (past volatility) and excess returns across assets. The value of the CSFB on a given day is the slope of the cross-sectional linear regression of risk and excess returns - the higher the slope, the higher the risk appetite.

Credit Suisse Global Risk Appetite Index is based on the relationship of excess returns and volatility of safe assets and risky assets. This RAI examines equities and bonds of developed and emerging markets [Lascelles, 2013].

The Royal Bank of Canada (RBC) Risk Appetite Index combines 46 different inputs measuring overall market price, volatility, correlation, flow, sentiment surveys, and third-party risk indices [Lascelles, 2013]. The index weighs all inputs equally.

1.3 Innovation of this Research

This research innovates in the way Stochastic Volatility (SV) parameters are estimated. Though the first two steps of the stochastic volatility estimation process (particle filtering and particle smoothing) are identical to Kim's model [2005, 2008], the final step in the process (parameter estimation) is done completely differently. While Kim uses the traditional approach of conditional likelihood maximization, this research proposes a new technique called Residuals Distribution Matching (RDM). It uses a standard Nelder-Mead optimization algorithm[4] to find the combination of parameters that:

1. Results in a residuals distribution with the desired shape;
2. Selects estimates with superior in-sample volatility prediction power.

[4] fminsearch function in MATLAB.

The final result is an estimator that:

1. Appears to outperform the traditional GARCH volatility estimator;
2. Has better volatility prediction capability than another popular SV model: the Harvey et al. model [1994];
3. Is simpler to derive and use than Kim's multi-step likelihood estimator;
4. Has better empirical convergence than Kim's likelihood estimator.

1.4 Thesis Structure

Section 2 discusses the main theoretical aspects of this project. Section 3 presents and discusses the major results of the research. It also validates the results of this research against those of previous studies. Finally, Section 4 presents the conclusions of this research and outlines potential future work.

2: Theoretical Considerations

"All models are wrong, but some are useful" – George Box (1919-2013)

2.1 Particle Filtering

The Particle Filter (PF) is a probability-based (Bayesian) estimator based on Monte Carlo methods [Simon 2006]. Figure 2.1 below shows an example of a filtered system using 100 particles (or Monte Carlo samples).

Figure 2.1 Evolution of 100 Particles or Monte Carlo Samples to Estimate Volatility

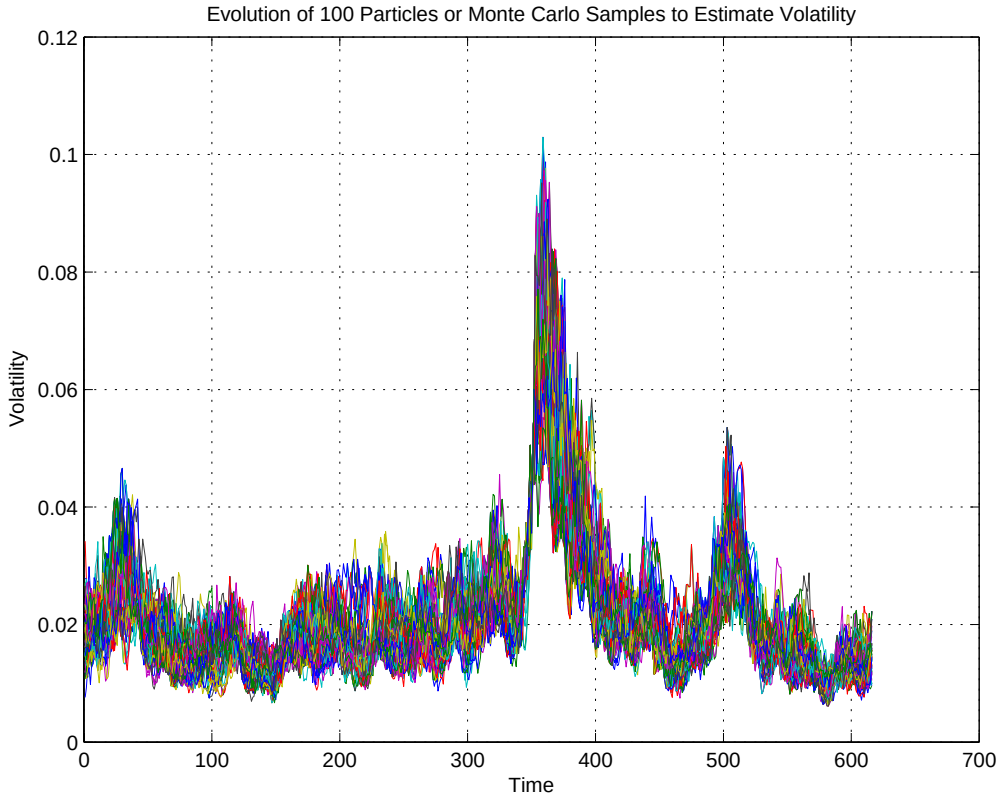

Since the particle filter is nothing but a Monte Carlo simulation, the estimated volatility at each time step is the mean value of all particles at that time step. In addition, one can obtain an estimate of the volatility estimation error by computing the error statistics of the particles. The

95% confidence interval, for example, is computed by finding the 2.5th and 97.5th percentiles of the particles.

Originally imagined by Metropolis[5], the PF is based on the simple idea that as the number of trials of an experiment approaches infinity, the probability of an event taking place converges to the ratio of number of occurrences of that event divided by the number of trials. With the advent of modern computing, Metropolis' idea has spawned a family of filters that can outperform any Kalman-type filter at the cost of extra computation.

The first PF model was put forth by Gordon, Salmond, and Smith [1993]. Known as Sequential Importance Sampling (SIS), this original formulation can at times suffer from sample impoverishment. Given certain conditions, all Monte Carlo samples may end up with zero weight in the estimate, except for a single sample [Simon, 2006]. Several mitigation strategies have been devised over the last 20 years.

This thesis uses two particle methods (an SIS filter and an SIS smoother) to estimate the stochastic volatility of assets over time. Section 2.2.6 discusses how degeneracy issues are handled.

2.2 Stochastic Volatility Estimation

2.2.1 State-Space Equations

The measurement and state propagation equations above ((6) and (7)) constitute the starting point of the stochastic volatility estimator. The proposed model mainly follows the treatment of Kim [2005, 2008], who slightly modified the formulation of Harvey et al. [1994]. Kim substitutes x for $(x-\gamma)$ in both equations and defines β as $\exp(-\gamma)$. She obtains:

$$y_t = ln\frac{S_{t+1}}{S_t} = \epsilon_{1t}\beta e^{(x_t/2)} \tag{8}$$

$$x_t = \phi x_{t-1} + \epsilon_{2t} \tag{9}$$

But Equation (8) still has a multiplicative error term, which is rather difficult to handle in filtering. One can transform the measurement equation so that the error term becomes additive by

[5] of Metropolis-Hastings fame.

squaring both sides and taking the logarithm of both sides [Shumway and Stoffer, 2006]. The resulting state space-formulation is

$$y_t = ln((r_t - \bar{r}_t)^2) = ln(\beta^2) + x_t + ln(\epsilon_{1t}^2) \quad (10)$$

$$x_t = \phi x_{t-1} + \epsilon_{2t} \quad (11)$$

where the left-hand side of the measurement equation (10) is now the log of the square of the zero-mean log-returns[6], and the error term in the measurement equation is the log of the square of a Gaussian random variable, which is a $ln(\chi_1^2)$ random variable [Shumway and Stoffer, 2006]. On the other hand, the noise term in the state propagation equation (11) remains Gaussian. Note that in the above derivations, y_t and x_t do not necessarily always represent the same quantities. For example, y_t denotes the log-returns in Equation (6). On the other hand, it represents the log of the square of the log-returns in Equation (10). Think of y_t and x_t as the measurement and the state, which change as we rearrange the state-space equations from an intractable form to a shape that can reasonably be tackled. There remains only one problem: the $ln(\chi_1^2)$ distribution does not have zero mean, which is a desirable feature of an error term. To remedy this issue, the following substitution is made:

$$\alpha = ln(\beta^2) + E(ln(\epsilon_t^2)) \quad (12)$$

where $E(ln(\varepsilon_t^2))$ is the expected value of a $ln(\chi_1^2)$ random variable, which is -1.27. The final version of the state-space description of the stochastic volatility model is

$$y_t = x_t + \alpha + v_t \quad (13)$$

$$x_t = \phi x_{t-1} + \epsilon_{2t} \quad (14)$$

where the error term in the measurement equation (v_t) follows a zero-mean $ln(\chi_1^2)$ distribution, ε_{2t} is a Gaussian error term with unknown but constant variance. This system has three unknown parameters: α, ϕ, and the variance of the ε_{2t} error term (henceforth denoted by Q). The value of these parameters must somehow be estimated in an optimal fashion.

Shumway and Stoffer [2006] propose modeling the observation noise as a mixture of two Gaussian distributions. This model requires the simultaneous estimation of seven parameters

[6] See Equations (5.6) and (5.7) in Kim's thesis [2005].

using only one set of measurements. Kim [2005] estimates the parameters of the US/GBP exchange rate volatility using both the three-parameters model derived above and the more complex seven-parameters model. She finds that both models give similar results. In light of this, the authors elect to stick to the simpler model.

2.2.2 Summary of Stochastic Volatility Determination Process

Figure 2.2 below shows the overall flow diagram of the stochastic volatility determination process. The details of that process are explained in the following sections.

Figure 2.2 Flow Chart of the Stochastic Volatility Determination Process

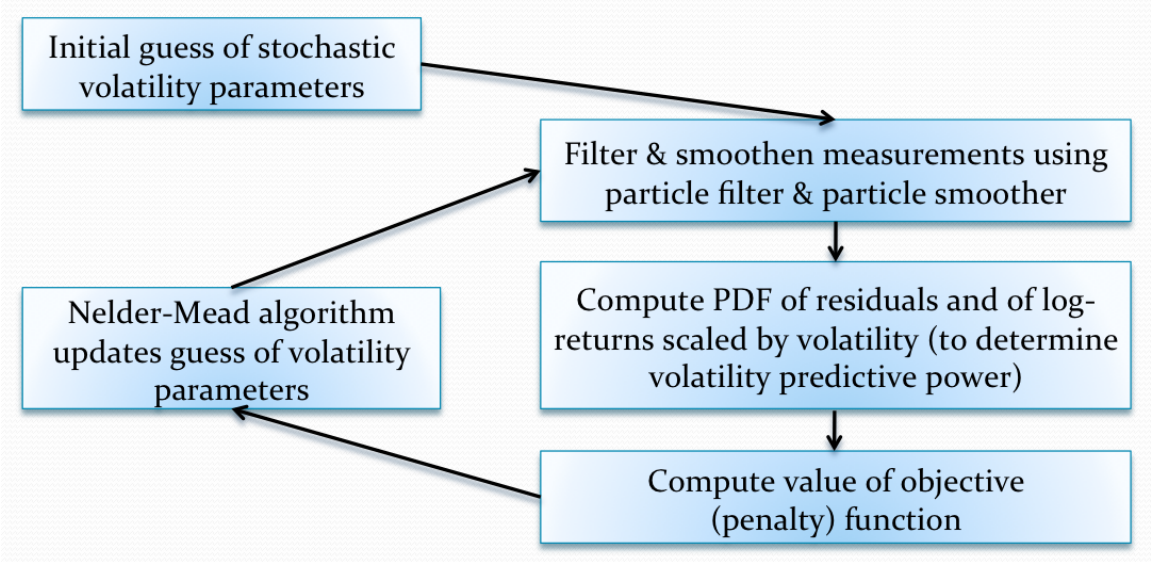

2.2.3 Parameter Estimation Process

2.2.3.1 Kim's Approach

Kim [2005] uses the conditional likelihood maximization approach to estimate the three unknown parameters (α, ϕ, Q). This algorithm is similar in spirit to the basic algorithm for estimating a ARCH or a GARCH process [Tsay, 2010]. Kim's Ph.D. thesis actually includes the source code she uses for her calculations. During initial testing, the authors found instances where Kim's likelihood maximization algorithm required baby-sitting because the program had to be

interrupted and restarted several times as the accuracy of the solution improved (see Section 2.4.3 in Kim's thesis [2005]). In addition, Kim's algorithm did not always deliver satisfactory results[7].

In light of all this, we propose a different algorithm that requires no baby-sitting and yields good results for all seven data series analyzed so far[8]. We call this proposed algorithm the Residual Distribution Matching (RDM) algorithm.

2.2.3.2 Residuals Distribution Matching (RDM)

This algorithm is the innovation of this research project. It estimates the optimal value of the system parameters iteratively as follows:

1. The user specifies an initial guess for the system parameters. In particular, we find that an initial guess of $[\phi, Q, \alpha] = [0.95, 0.05, -8]$ works for all seven times series analysed so far;
2. The algorithm uses a particle filter and a particle smoother to compute the volatility for this guess. That part of the algorithm is identical to Kim's algorithm. Refer to Section 2.1.2 in Kim's thesis [2005] for more information;
3. The time-varying volatility signal is compared to the actual asset returns to generate a time-series of residuals. This is done using Equation (6) above;
4. The empirical PDF of these residuals is compared to the PDF of the ideal distribution of residuals: the $\ln(\chi 12)$ distribution (see Figure 2.3 below for a typical example);
5. The algorithm computes the area of the empirical residuals PDF falling outside the $\ln(\chi 12)$ distribution. Ideally, this area would be zero or near zero. Using the area outside the theoretical PDF as an objective function is preferable to comparing the mean and variance of the empirical PDF to that of the theoretical PDF because area is much less sensitive to outliers. The algorithm also adds a small penalty for cases where the estimate has poor in-sample one-step-ahead volatility prediction capability. This ensures both a good fit and reasonable in-sample volatility prediction capability;
6. The whole algorithm is wrapped in a Nelder-Mead optimization algorithm9. This is the well-known fminsearch function in MATLAB. The reader should remain mindful that the

[7] During initial tests carried out by the authors of this study.

[8] The RDM algorithm was successfully tested on the following data series: DEX Universe Bond Index, spot Gold price, spot Brent oil price, S&P 500 index, 10-yr U.S. Treasury swap rate, S&P TSX Composite index, US Dollar index.

Nelder-Mead algorithm does not guarantee convergence any more than the conditional likelihood maximization algorithm does.

Figure 2.3 Theoretical & Empirical Distribution of Residuals for the Optimal SV Estimate for SPX

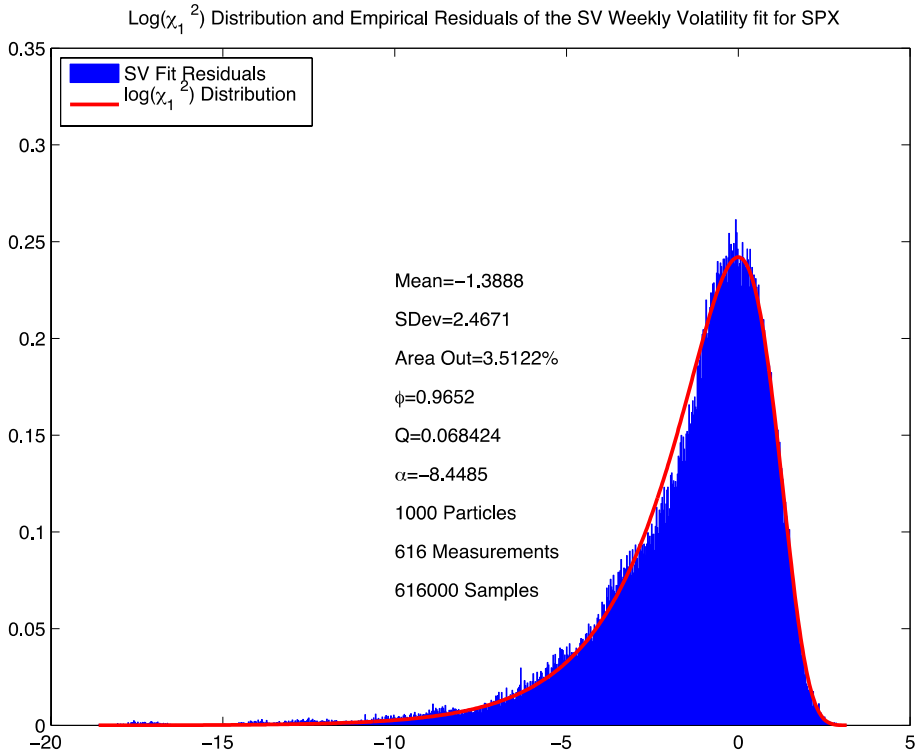

In essence, whether one maximizes conditional likelihood or seeks to match the empirical residuals PDF to the theoretical residuals PDF, one should get the same result, as long as the model is correctly specified, that is, as long as it truly represents reality. Initial testing reveals that the RDM algorithm provides better robustness than the conditional likelihood maximization model. In addition, it requires no *baby-sitting* and contains a penalty feature that rewards good in-sample volatility prediction.

[9] In initial testing, the Nelder-Mead method yielded much better results than the steepest ascent method.

2.2.4 Data Considerations

2.2.4.1 Data Start and End

Unless otherwise mentioned, all data series used in this research are weekly[10] asset prices/returns starting on January 1st 2002, and ending on November 1st, 2013 for a total of 616 measurements. 2002 was selected as a starting point for two reasons:

1. There is no need to deal with missing data issues resulting from the financial markets shutdown that followed the tragic events of September 11th, 2001.
2. The FRED database begins tracking the 10-yr interest rate swap in 2001, and the 10-yr swap is an input of the volatility/correlation component of the Risk Appetite Indicator.

2.2.5 Singularity Avoidance

The proposed model has a singularity that must be protected against. In Equation (10) above, if the log-return between two data points matches the average log-return, then the argument of the log function vanishes, which results in a singularity. To avoid this situation, the data points near the singularity are slightly adjusted to ensure that the system remains well behaved. This is done by adjusting these vulnerable data points so that the return is some random value very close, but not identical to the average log-return. This enables the algorithm to perform as expected, but creates a *volcanoe*-shaped scaled log-return distribution (as shown in Figure 2.4 below). This shape is created by the scarcity of returns near 0 (about the mean log-return), and the resulting abundance of log-returns just around the mean. As shown empirically in Section 3.2, this singularity protection process does not negatively impact the results. This makes theoretical sense because what really drives volatility processes is not samples where returns are near the average, but samples that show large gains or losses.

[10] Last trading day of the week.

Figure 2.4 Empirical Predictive PDF of Log-Returns for the S&P 500 Index (Scaled by SV)

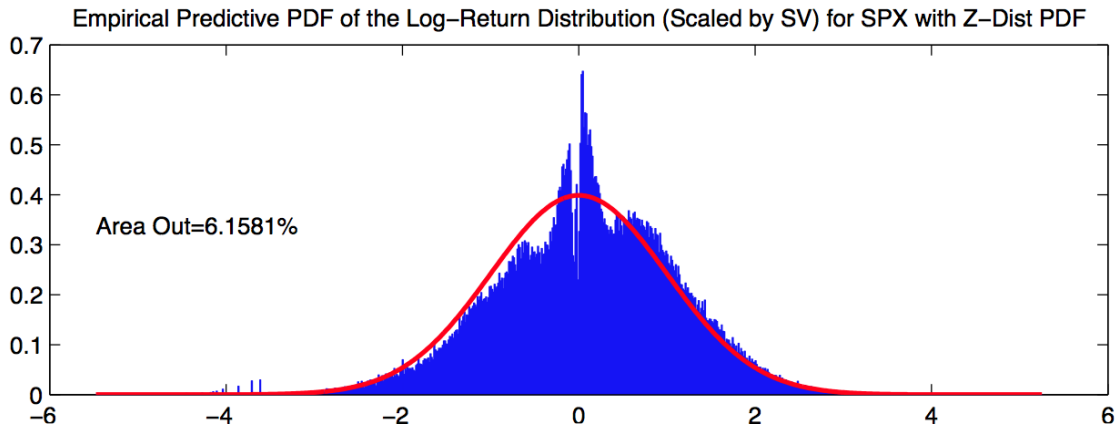

2.2.6 Particle Filter Degeneracy Issues

The SIS particle filter has well-documented sample impoverishment (degeneracy) problems [Simon, 2006]. To ensure that a given set of simulation results is not plagued by particle degeneracy, the authors visually inspect the 95% confidence interval of the volatility estimate for possible narrowing or abnormally high jitteriness. Such behaviour (shown in Figure 2.5 below for a simulation of Gold using only 10 particles) betrays particle filter degeneracy [Doucet and Johansen, 2009]. But as seen in Figure 2.6 and Figure 2.7 below, 250 particles yield satisfactory results.

Figure 2.5 Gold Volatility Estimate and 95% Confidence Interval (with 10 Particles)

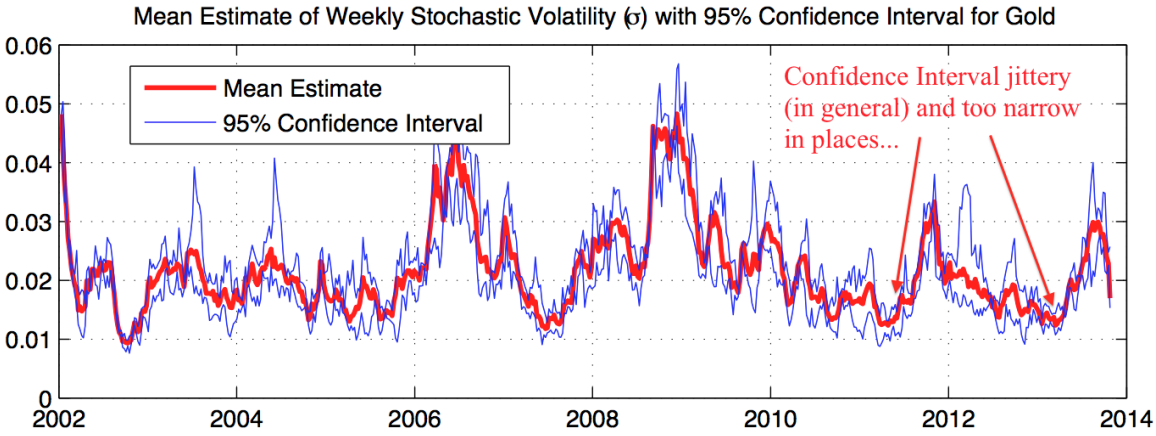

Figure 2.6 Gold Volatility Estimate and 95% Confidence Interval (with 1000 Particles)

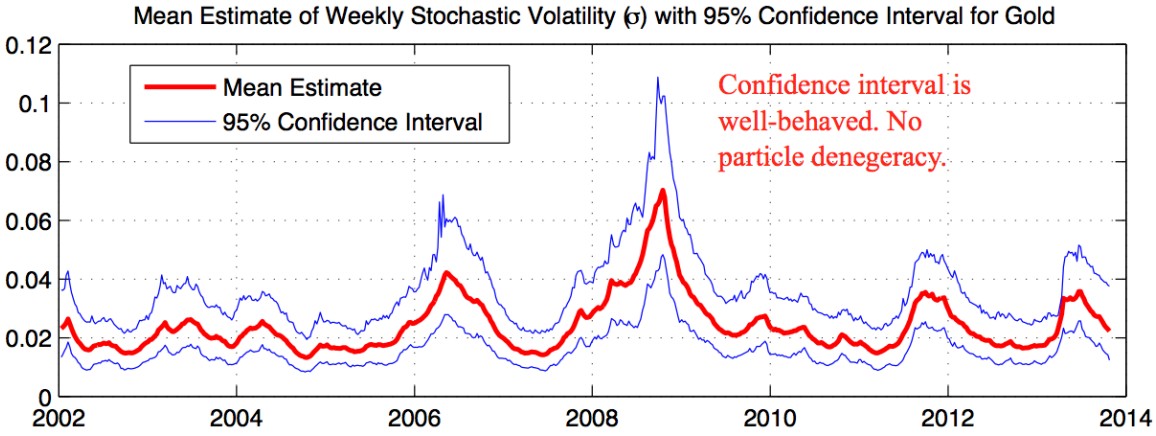

Figure 2.7 Gold Volatility Estimate and 95% Confidence Interval (with 250 Particles)

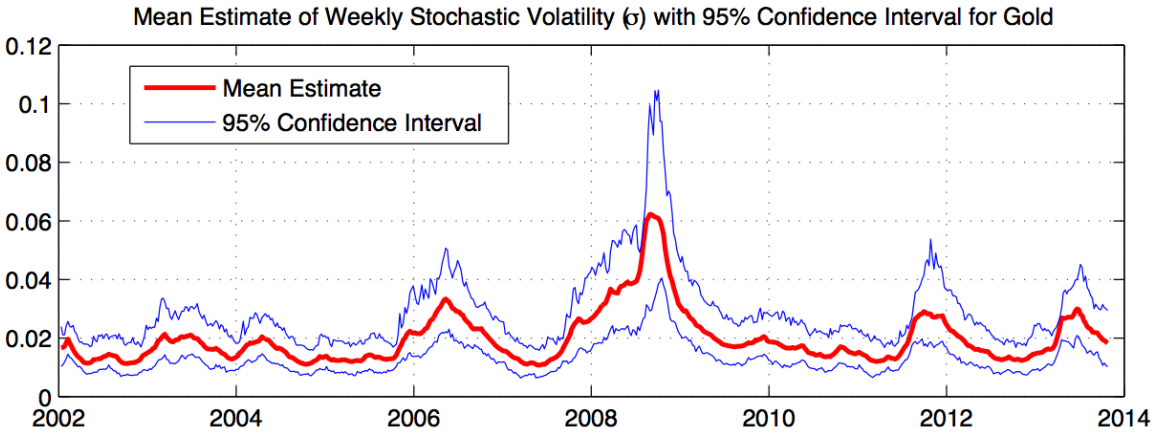

It would be possible to eliminate the prospect of particle filter degeneracy entirely by implementing a more robust resampling algorithm - Markov Chain Monte Carlo resampling [van der Merwe et. al. 2000(b)]. But for now, a visual inspection of the 95% confidence interval suffices.

2.2.7 Why not use a Kalman Filter?

As mentioned in the literature review section, the Kalman filter is the best possible linear filter for any noise distribution. So why not use it to filter the stochastic volatility system, which is linear (see Equations (13) and (14))? It turns out that though the model is linear, the true stochastic volatility process may not be. Hence, a non-linear formulation like the particle filter used here is better suited.

2.3 Dynamic Correlation

As mentioned in the literature review section, we seek a simple divide-and-conquer approach to determining dynamic cross-asset correlation. The selected strategy uses a two-stage process:

1. Estimate the SV of each asset separately using a univariate formulation;

2. Blend the residuals from step 1 to compute the dynamic correlation using an Exponential Moving Average (EMA).

The RiskMetrics group suggested the EMA approach [Engle, 2002]. Their model has a smoothing factor *(λ)* of 0.94 and it is regularly used in the industry. The model uses a recursive definition of the correlation matrix q_{ij}:

$$q_{ij,t} = (1-\lambda)(\epsilon_{i,t-1}\epsilon_{j,t-1}) + \lambda q_{ij,t-1} \tag{15}$$

from which one obtains the dynamic correlation as follows:

$$\rho_t = \frac{q_{ij,t}}{\sqrt{q_{ii,t} * q_{jj,t}}} \tag{16}$$

Like every other EMA, this filter has a lag. In this case, the lag is 17 time steps[11].

2.4 Risk Appetite Indicator

The authors seek to replicate the RBC Risk Appetite Index (RAI), as it is only published intermittently, which is not convenient for investors.

2.4.1 Data Inputs in the RBC Risk Appetite Index

The RBC RAI has 46 inputs [Lascelles, 2013]. Table 1 below shows all RAI inputs in column 1. Column 2 shows which inputs were located and used in the replicated RAI: "√" implies that the input was found, while "×" indicates that the input was not found. Column 3 indicates the *polarity* of each indicator: "*1*" implies that a high value of the input betrays a risk loving attitude, while "*-1*" implies that a high value of the inputs denotes a risk averse sentiment. Finally, columns 4 and 5 show the source of the data inputs that were located.

[11] Filter lag = $(1-\lambda)^{-1}$ = $(1-0.94)^{-1}$=16.67~17

Table 1 RBC RAI Inputs, List of Inputs Obtained by the Authors, and References

RBC RAI Inputs	Included in Replicated RAI	Risk Loving/ Averse	Data Source	Reference Link
PRICE				
U.S. 10-Year Treasury Real Yield	√	1	FRED ST. LOUIS FED	http://research.stlouisfed.org/fred2/search?st=10-Year+Treasury+Inflation-Indexed+Security%2C+Constant+Maturity
U.S. Corporate Bond Spread	√	−1	FRED ST. LOUIS FED	http://research.stlouisfed.org/fred2/series/BAMLC0A0CM
U.S. High Yield Bond Spread	√	−1	FRED ST. LOUIS FED	http://research.stlouisfed.org/fred2/series/BAMLH0A0HYM2/
Global Corporate Bond Spread	×	N/A		
Global High Yield Bond Spread	×	N/A		
Emerging Market Sovereign Bond Spreads	×	N/A		
Sovereign Credit Default Swap (CDS) Spreads CDS spreads	×	N/A		
Price-to-Earnings (P/E) Ratio P/E ratio of S&P 500 Index.	√	1	Bloomberg	Bloomberg
Put-Call Ratio Ratio on S&P 500 Index.	√	−1	Bloomberg	Bloomberg
VOLATILITY				
Euro-U.S. Dollar (EURUSD) Option Volatility 1 Year	×	N/A		
U.S. Swaption Volatility 1Y 5Y Normalized	×	N/A		
CBOE Volatility Index (VIX)	√	−1	Bloomberg	Bloomberg
Consensus Economics Standard Deviations of GDP Forecasts – U.S.	×	N/A		
Consensus Economics Standard Deviations of GDP Forecasts – Eurozone	×	N/A		
Consensus Economics Standard Deviations of GDP Forecasts – U.K.	×	N/A		
Consensus Economics Standard Deviations of GDP Forecasts – Japan	×	N/A		
Survey of Professional Forecasters Real GDP Forecast Dispersion	×	N/A		
Economic Policy Uncertainty Index	×	N/A		
CORRELATION				
Cross-Asset Class Correlation	√	−1	Dynamic Correlation	Dynamic correlation estimated in this thesis
Correlation of S&P 500 Stocks and S&P 500 Index	×	N/A		

FLOW				
Net Speculative Long Positions in U.S. Treasury	×	N/A		
U.S. Corporate Bond Issuance	√	1	Sifma	http://www.sifma.org/research/statistics.aspx
Eurozone Corporate Bond Issuance	×	N/A		
U.S. Net Equity Issuance	√	1	Sifma	http://www.sifma.org/research/statistics.aspx
Eurozone Net Equity Issuance	×	N/A		
U.S. Mutual Fund Flows Into Equities	√	1	ICI	http://www.ici.org/research/stats
U.S. Mutual Fund Flows Into Bonds	√	1	ICI	http://www.ici.org/research/stats
U.S. Mutual Fund Flows Into Emerging Markets	×	N/A		
Credit Growth – U.S.	×	N/A		
Credit Growth – Eurozone	√	1	MD Briefing	http://mdbriefing.com/eurozone-credit.shtml
SURVEY				
Thomson Reuters/University of Michigan Consumer Sentiment Index	√	1	Bloomberg	Bloomberg
European Commission Economic Sentiment Indicator –EU	√	1	Bloomberg	Bloomberg
European Commission Economic Sentiment Indicator – Eurozone	√	1	Bloomberg	Bloomberg
Japanese Consumer Confidence Index	√	1	Bloomberg	Bloomberg
National Association of Home Builders Housing Market Index	√	1	Bloomberg	Bloomberg
Market Vane Bullish Consensus Stock Index	×	N/A		
Rasmussen Investor Index	√	1	Rasmussen	http://www.rasmussenreports.com/public_content/business/indexes/rasmussen_consumer_index/rasmussen_consumer_index
THIRD-PARTY RISK INDICES				
Bloomberg Financial Conditions Index – U.S.	√	1	Bloomberg	
Bloomberg Financial Conditions Index – Europe	√	1	Bloomberg	Bloomberg
BofA Merrill Lynch Financial Stress Index	√	−1	Bloomberg	Bloomberg
Kansas City Financial Stress Index	√	−1	Bank of Kansas City	http://www.kansascityfed.org/research/indicatorsdata/kcfsi/
St. Louis Federal Reserve Bank Financial Stress Index	√	−1	FRED ST. LOUIS FED	http://research.stlouisfed.org/fred2/series/STLFSI
Citigroup Macro Risk	×	N/A		

Index				
Westpac Risk Aversion Index	√	−1	Bloomberg	Bloomberg
Credit Suisse Global Risk Appetite Index	√	1	Credit Suisse	https://doc.research-and-analytics.csfb.com/docView?language=ENG&format=PDF&source_id=csplusresearchcp&document_id=805297150&serialid=DNuqDu3QYtx6%2BfzZW95CNdT3bFZg7fizc7vlTMbbhUc%3D
State Street Investor Confidence Index	√	1	Bloomberg	Bloomberg

2.4.2 Complementary Data Inputs Required to Replicate the RBC RAI

As seen in the above table, 24 out of 35 inputs for the price, flow, survey and third-party sub-indices were found. On the other hand, few volatility and correlation inputs were located. The following explains how the models implemented for this thesis fill this gap.

For volatility, the authors use the SV model to construct an in-house composite market volatility index combining the S&P 500 index, gold, Brent oil, the 10-yr US-treasury floating swap, and the US dollar index. This is in addition to the implied volatility VIX index data. Hence, two volatility inputs are used to construct the volatility sub-index.

For correlation, the authors reproduce the composite correlation index in the RBC RAI using the dynamic correlation model created in this thesis. It combines the pair-wise correlations among the S&P 500 index, gold, Brent oil, the 10-yr US-treasury floating swap, and the US dollar index (10 correlation estimates in total). These estimates are combined into a composite market correlation index. In the end, this composite correlation sub-index looks very similar to the correlation composite published by RBC.

Figure 2.8 below shows the relative weight of each of the six sub-indices in the RBC RAI.
Table 2 also shows this information, as well as the number of inputs that the authors use to replicate the index.

Figure 2.8 Information Content of each Sub-Index in the RBC RAI

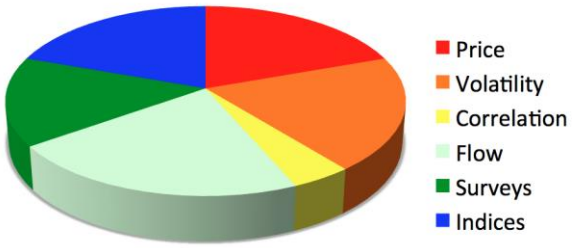

Table 2 Sub-Indices in the RBC RAI along with number of Inputs in each Sub-Index.

Sub-Indices	RBC RAI	Replicated RAI
PRICE	9	5
VOLATILITY	9	2
CORRELATION	2	1
FLOW	10	5
SURVEY	7	6
THIRD-PARTY RISK INDICES	9	8
Total	46	27

2.4.3 Filtering High Frequency Data

The put/call ratio and the Westpac Risk Aversion Index are particularly noisy. To compensate for this jitteriness, the put/call is filtered through a 5-day EMA, while the Westpac index is filtered through a 10-day EMA.

2.4.4 Overview of Risk Appetite Indicator Replication Process

Figure 2.9 below shows a flowchart of the RAI replication process. The process is detailed in the following sub-sections.

Figure 2.9 Flowchart of the RAI Replication Process

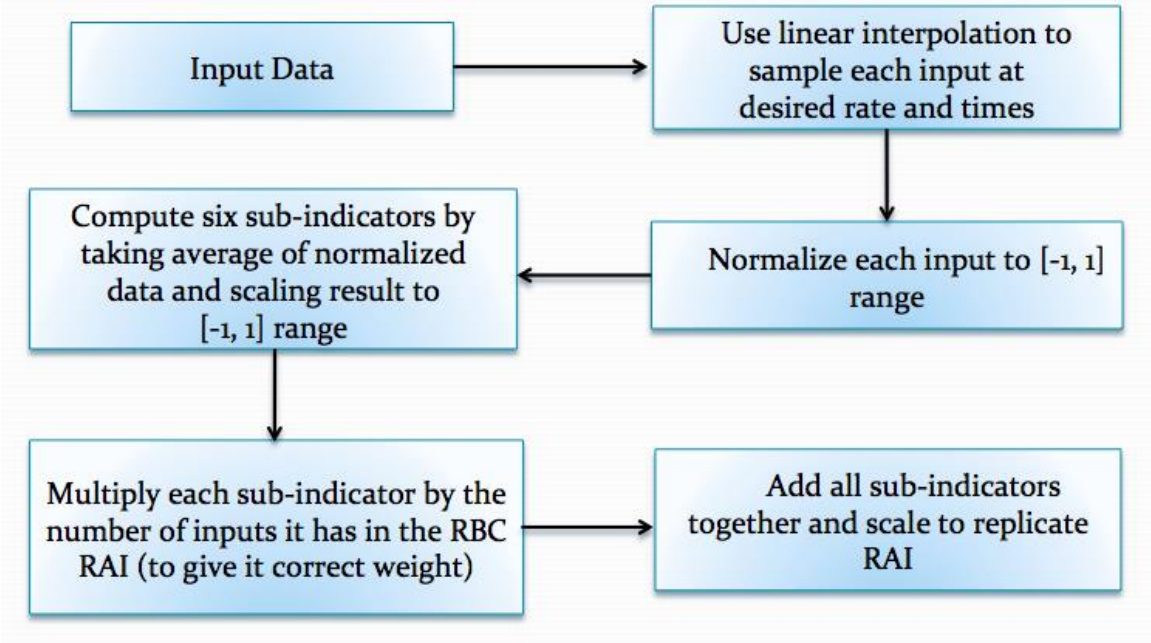

2.4.5 Linear Interpolation for Time Sampling

Given that the various inputs are provided at different frequencies, the RAI replication strategy uses linear interpolation to sample each input at the desired frequency – every five days in this case. Hence, regardless of how much or how little data is available for a given input, the replicated RAI always uses the data in the most meaningful manner possible.

2.4.6 Input Scaling

Given that not all RAI input have the same scale, it is necessary to normalize all inputs to give them the correct weight in the final index. For example, a P/E ratio of 40 does not at all mean the same thing as a corporate bond spread of 40bps. Hence, all inputs are normalized to a range between -1 and +1, where -1 is the most risk averse value encountered in the dataset and +1 is the most risk loving value in the dataset. Figure 2.10 below shows an example of this process:

Figure 2.10 Normalized Values of Inputs for the Price Sub-Indicator of the RAI

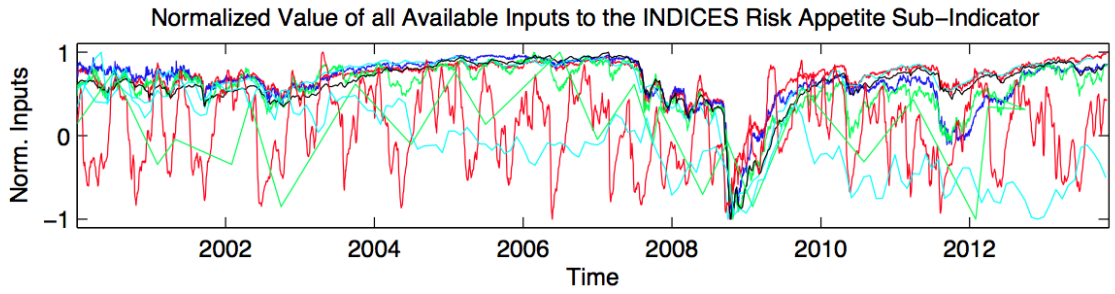

2.4.7 RAI Replication Strategy

Given that 40% of the inputs of the RBC RAI could not be obtained, the authors attempt to best replicate the RBC RAI. Lascelles [2013] states that each of the 46 inputs in the RBC RAI is given equal weight. Hence, all normalized inputs within a sub-indicator (say price) are added and then normalized within the -1 to +1 range as in Figure 2.11 below:

Figure 2.11 Normalized Indices Sub-Indicator of the Replicated RAI

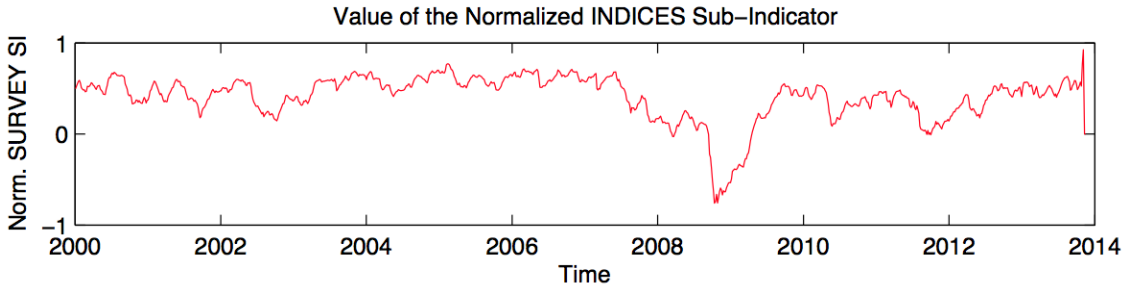

Each of the replicated sub-indicators is scaled by the number of inputs in the sub-indicator. For example, the replicated price sub-indicator is multiplied by 9 because the RBC RAI contains 9 price inputs. This strategy is particularly useful to ensure that the replicated volatility sub-index (which only has two inputs as opposed to nine inputs in the RBC RAI) does not ruin the accuracy of the replicated RAI.

All scaled sub-indicators are then added together and scaled to a range similar to that of the RBC RAI. Reasonable thresholds for risk aversion, risk loving, risk seeking, risk neutral, and risk reluctant are added (though they do not match the values in the RBC RAI).

2.5 How the Volatility/Correlation Estimators Relate to RAI

So far, the reader might be tempted to think that this research project is really two separate projects: 1. volatility/correlation estimation, 2. Risk Appetite Indicator; but that is not quite the case. It turns out that getting good dynamic correlation data free of charge is quite difficult.

As mentioned in the last section, the RAI requires a fair amount of volatility and correlation data (11 inputs out of 46). For the volatility part of the indicator, volatility estimation provides 1 out of 9 inputs for the RAI. But since only two inputs are available (in-house stochastic volatility indicator and the VIX index), each input has a weight equivalent to 4.5 inputs in the RBC RAI. For correlation, the in-house correlation estimator provides 1 out of 2 inputs for the RAI, which has a weight equal to 2 inputs in the RBC RAI. In total, volatility/correlation estimation gives 6.5/46=14% of the total information content of the replicated RAI, which is non-negligible (see Figure 2.12).

Figure 2.12 Information Content in Replicated RAI from Volatility/Correlation Models in this Project

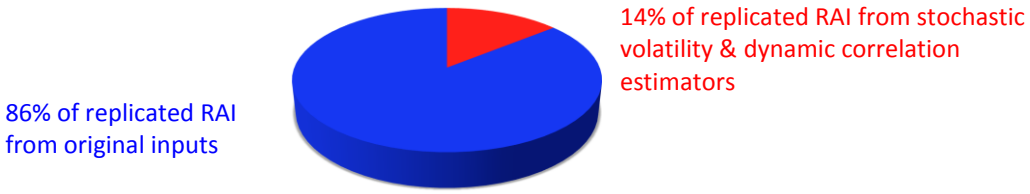

Due to the significant influence of in-house volatility/correlation estimation on the replicated RAI, the authors undertook both projects simultaneously.

3: Results and Discussion

"To err is human, but to really screw up you need a computer" – popular expression

This section not only presents the results obtained from the proposed models, but also verifies/validates them against theoretical results and/or the findings of previous research campaigns. The authors also seek to determine where the proposed models perform better or worse than other models.

3.1 Computational Speed

The MATLAB source code originally used by Kim [2005] for her Ph.D. thesis is optimized for speed through vectorization (replacement of for loops with vector operations), and other improvements. The updated code runs about 33% faster than the original code written by Kim. The final version of the filtering/smoothing code, along with the Nelder-Mead algorithm for optimization takes about two hours12 to find the optimal volatility estimate for a given asset; with 616 measurements and 1000 particles. Reducing the number of particle to 250 while keeping the dataset intact requires 25 minutes of computational time. Speed-up13 would be possible with the replacement of some of the code by .mex files, but this is left as future work.

3.2 Stochastic Volatility Estimator

3.2.1 Volatility of the US Dollar vs British Pound Exchange Rate

Figure 3.1 shows the estimated stochastic and GARCH volatility of the US dollar / British pound exchange rate between October 1st, 1981 and June 28th, 1986. On this plot, the noisy green signal corresponds to the log-returns, the blue curve represents the GARCH estimate, and the red curve denotes the Stochastic Volatility (SV) estimate. One immediately notices that the GARCH and SV estimates generally track each other.

[12] Using a MacBook Pro laptop computer with a 2.6GHz Intel Core i7 processor.
[13] Perhaps by as much as 80%-90%.

Figure 3.1 Stochastic and GARCH Volatility Estimates for US Dollar vs British Pound Exchange Rate

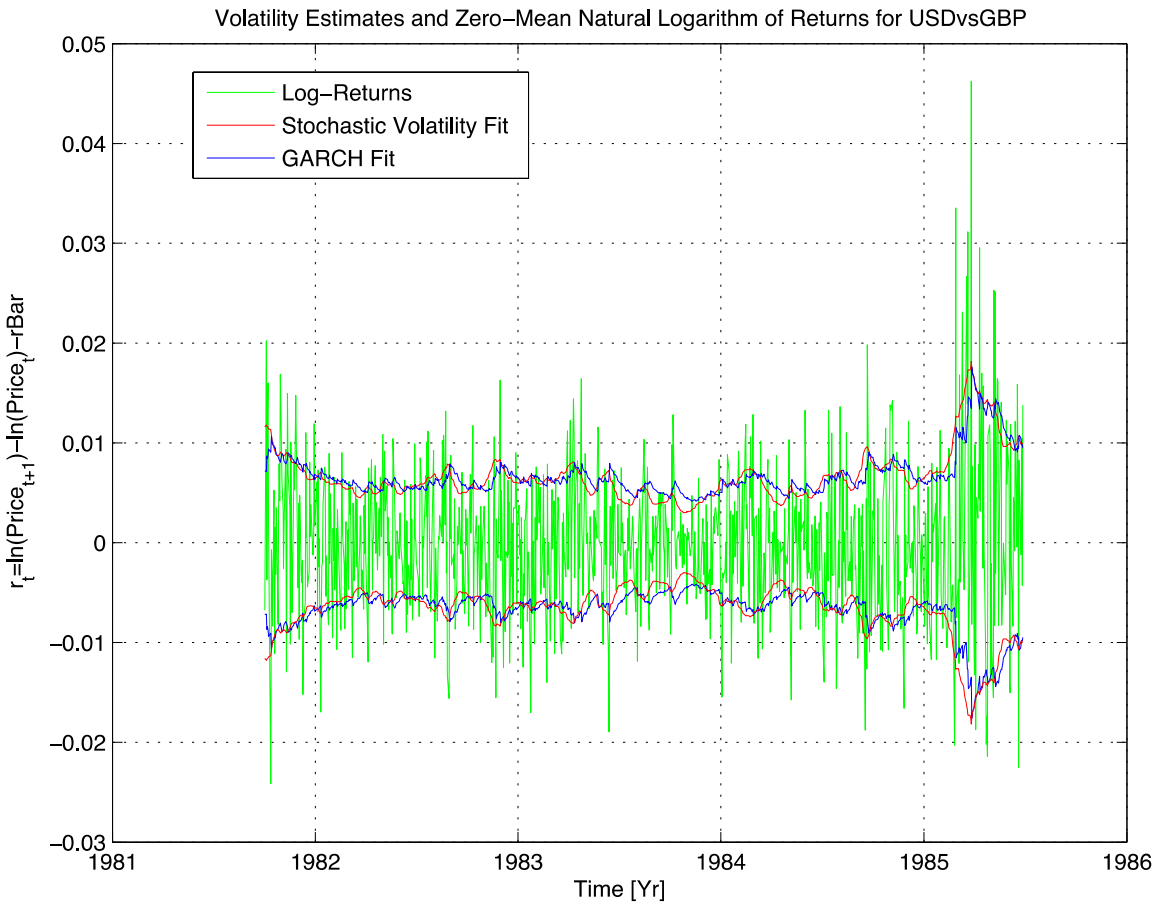

But a closer look at the volatility estimates (see Figure 3.2) shows that the stochastic estimate appears to perform better the GARCH estimate in that it:

1. Does not lag the measurements;
2. Quickly responds to sudden decreases in volatility.

Though not detectable in this dataset, the GARCH model also tends to over-estimate volatility in when shocks occur. This is indeed the case in almost all the datasets with large volatility shocks that are analysed.

Figure 3.2 Zoomed-in Portion of

Figure 3.1

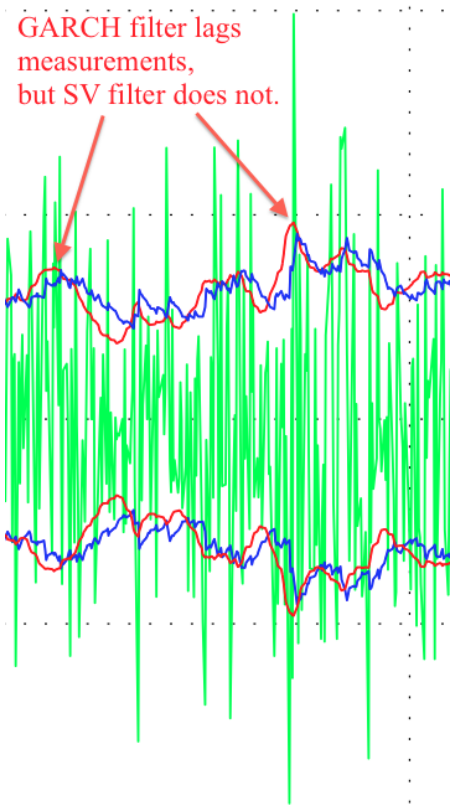

This US Dollar / British Pound dataset has become something of a benchmark case in the SV literature. Harvey et al. [1994] also estimate the volatility of this dataset using a quasi-maximum likelihood method, based on a Kalman filter. Recall that the Kalman filter is not as suitable for non-linear systems (such as most financial time series) as particle filters, which are specifically designed to tackle non-linear problems. The volatility estimates obtained by Harvey et al. [1994] (see Figure 3.3 below) are smoother than a GARCH fit, but appear to have a significant number of outliers near the end of the dataset, where there is a volatility shock. In addition, the Harvey et al. model appears to miss the period of low volatility between times 515 and 545 in Figure 3.3. Hence, it appears that the proposed model has better in-sample volatility prediction capability than the well-known Harvey et al. model14.

[14] One may notice that the log-returns in the present analysis slightly differ from those in Harvey et al. [1994]. This is likely due to the number of significant digits in the two datasets. This effect does not invalidate the results of this comparison.

Figure 3.3 Stochastic Volatility Estimates for USD/GBP Obtained by Harvey et al. [1994]

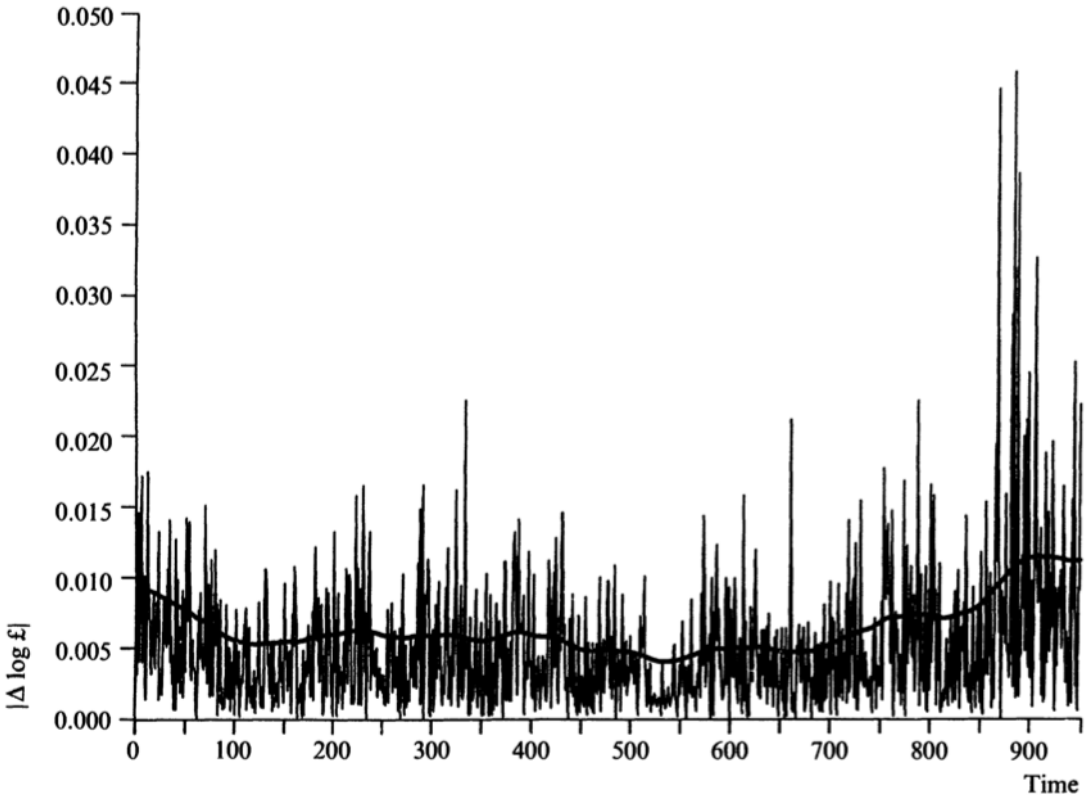

As for Kim [2005, 2008], she does not provide plots of the estimated SV, only of the model parameters. Therefore, one cannot directly compare her volatility filtering results to those obtained with the proposed model.

Another way to judge the quality of a volatility fit consists of normalizing the current measurement by the previous volatility estimate. This gives a measure of in-sample volatility-prediction power. Under ideal circumstances, the resulting PDF would be Gaussian. All other factors held equal, the RDM algorithm tends to select solutions that have good in-sample volatility prediction power. The SV computed here does indeed have better predictive power than the GARCH filter for all seven assets studied. In other words, the empirical PDF of the log-return distribution (scaled by volatility) obtained using the SV fit more closely resembles a Gaussian PDF than it does for the GARCH fit. An example of this is shown below for the US Dollar Vs. British Pound exchange rate (Figure 3.4 below):

Figure 3.4 Volatility Prediction Power for GARCH and SV Fits for USD/GBP Exchange Rate

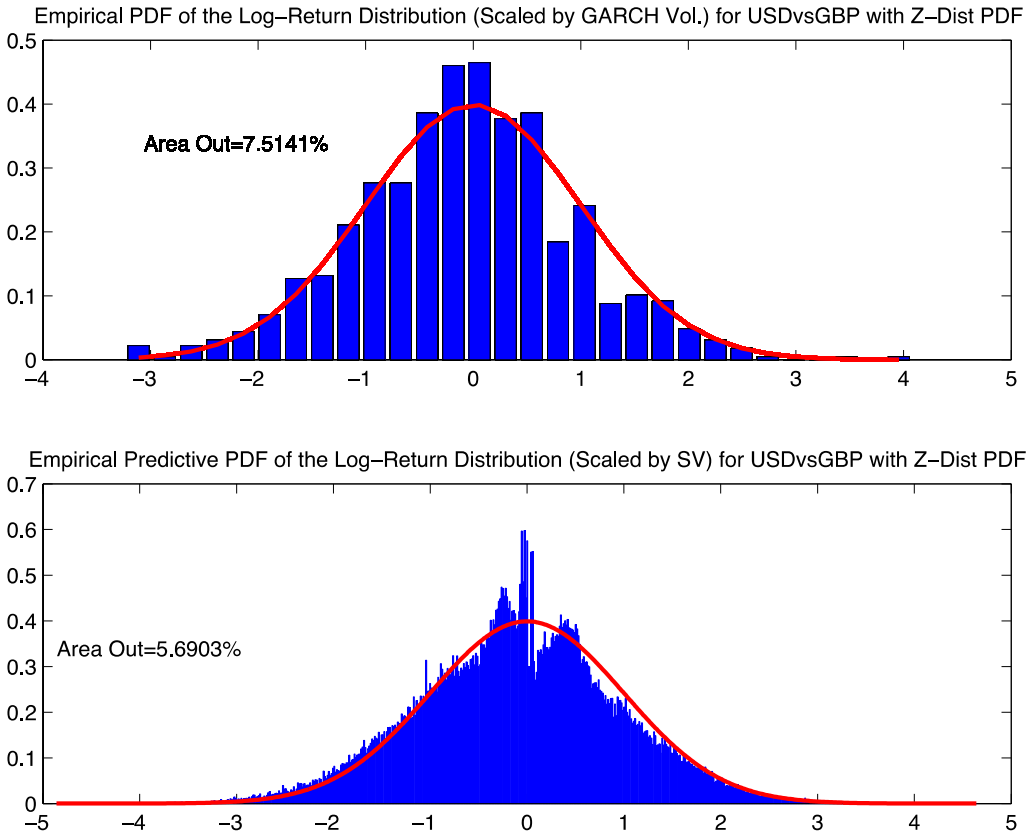

Figure 3.5 and Figure 3.6 below show the SV and GARCH volatility estimates for the S&P 500 index and for gold:

Figure 3.5 *GARCH & SV Fits for the S&P 500 Index*

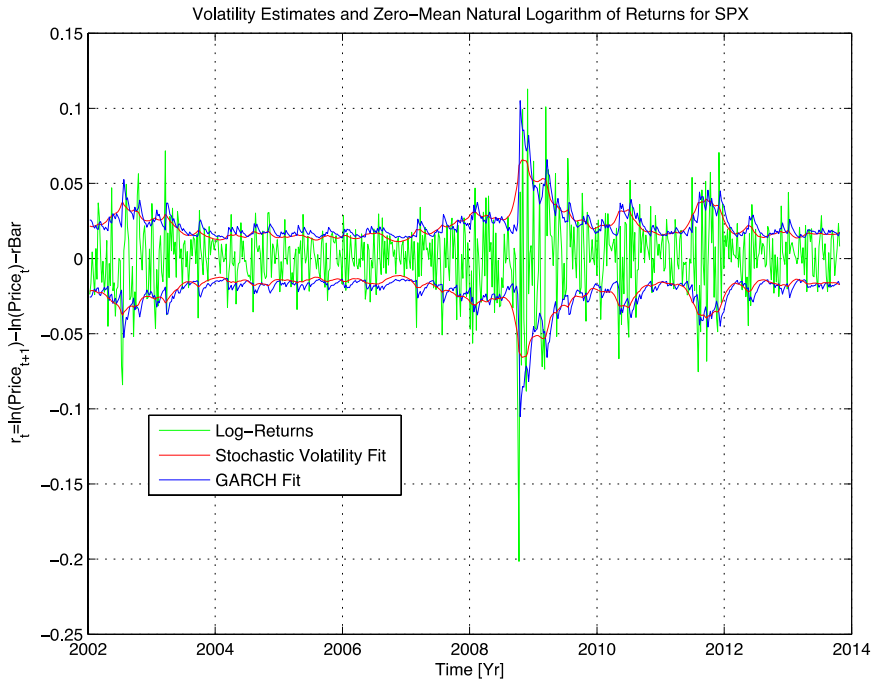

Figure 3.6 *GARCH & SV Fits for Gold*

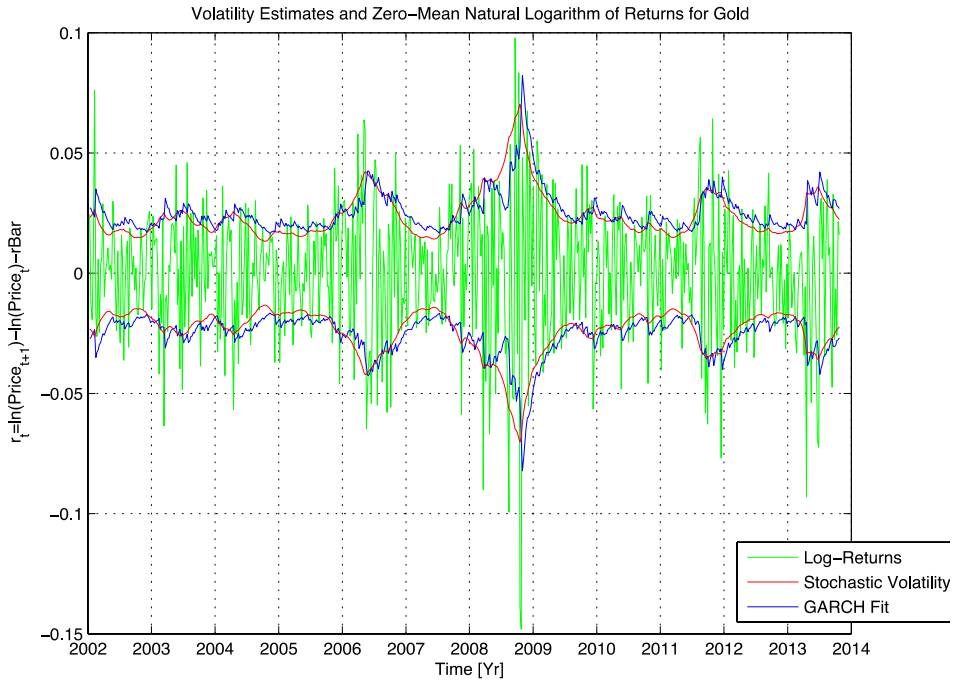

As seen in the above volatility plots, the SV estimate does not peak as high as the GARCH estimate, which is desirable [Gonzalez et al., 2002]. In fact, from all the relevant criteria identified (smoothness, absence of lag, quick response to decreases in volatility, reasonable as opposed to over-reaction to volatility spikes, and volatility prediction power) we conclude that the SV algorithm proposed here appears to outperform both the GARCH and the Harvey et al. models. This result does not come as a surprize because the present model has a stochastic term for volatility (which the GARCH model lacks), uses a particle filter (while the model developed by Harvey et al. uses a Kalman filter), and has a bias for estimates with better volatility prediction power. In conclusion, the proposed SV model is verified and valid. However, such improved results come at the cost of much lower computational efficiency than the GARCH model.

3.3 Dynamic Correlation Estimator

3.3.1 Verification of Correlation Estimator using Step Function

Two artificial datasets are generated to verify the dynamic correlation model. The first one is a simple step function, which instantly jumps to a different value every 400 seconds.

Figure 3.7 True Simulated Correlation for Step Function Case

The proposed dynamic correlation model is used to estimate the correlation estimation error using this known step function profile. The results are shown in Figure 3.8 below for the residuals computed with both the GARCH and SV fits.

Figure 3.8 Correlation Estimates for Step Function Verification Case

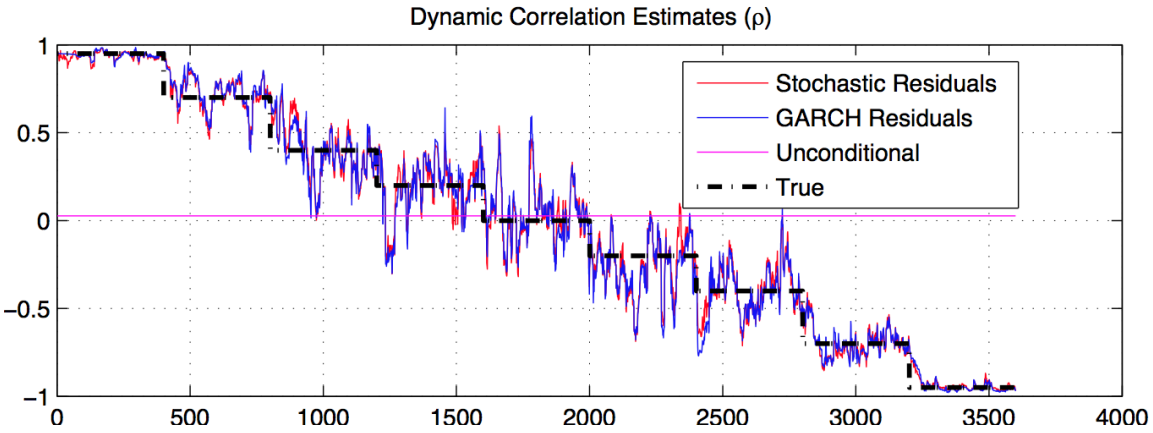

This SV fit uses only 25 particles, yet yields results very similar to the GARCH fit. The simulation is repeated with a larger number of particles, but yields a very similar result. Hence, what drives the accuracy of the correlation estimate is not the quality of the residuals (input) passed from the volatility filter, but the quality of the correlation filter itself.

As seen in the above plot, the RiskMetrics Exponential Moving Average (EMA) is much better at estimating dynamic correlation when it is significantly different from zero. So this estimator could estimate or predict the active risk of a passive or semi-active portfolio against its benchmark because:

1. In such cases, the correlation between the portfolio and the benchmark is high;
2. Estimates of asset volatility with in-sample predictive power are available;
3. Asset volatility and cross-asset correlation suffice to compute forward-looking active risk (in a Gaussian log-returns world).

Note how the unconditional correlation estimate (in magenta on the above plot), which is computed by taking the standard deviation of the residuals, completely fails to capture the time-varying nature of this correlation process.

By computing the difference between the true and estimated correlation, one obtains the estimation error. But to gain some confidence in this estimate and ensure that the model does not fall victim to over-fitting, one must find out whether the error estimate provides reasonable results for a completely different correlation process. This is done below.

3.3.2 Validation of Correlation Estimator using Mountain Function

The second artificial correlation process follows a mountain function:

Figure 3.9 True Simulated Correlation for Mountain Case

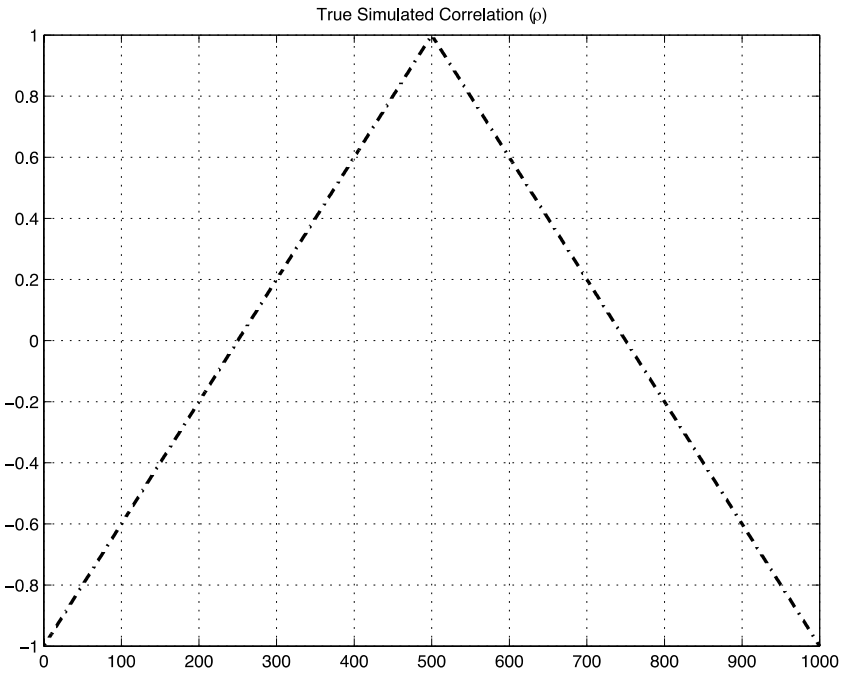

This artificial process is simulated and the data is passed to the volatility and correlation estimators. The estimated correlation and the 95% confidence interval are show below:

Figure 3.10 Dynamic Correlation Estimate, along with True Correlation and Confidence Interval

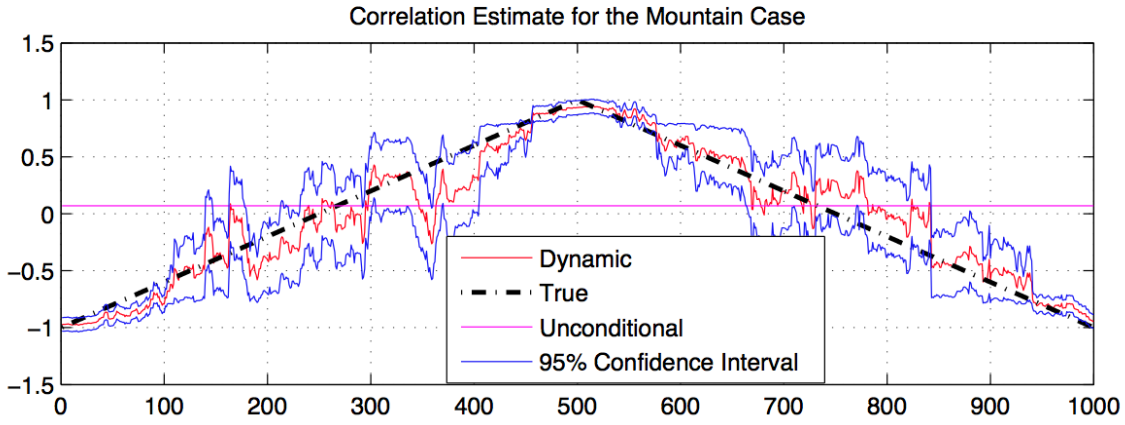

As seen in the above plot, the EMA lags the true signal slightly (17 time steps to be exact). This is from the very nature of the EMA. The estimate (in red) is below the true correlation (in black) as the correlation rises; and the hangs above the true correlation as it decreases. On the other hand, the lag is not big enough to throw the 95% confidence interval outside the true signal too often. The only real problem is that the confidence interval is very large near zero correlation. In fact, the 95% confidence interval near 0 is +/-0.34. Given this caveat, the dynamic correlation model is considered verified.

As for computational time, the correlation estimate is obtained instantly because no heavy-duty filter is used in the correlation estimation process. The residuals are passed to a simple EMA filter, which takes no time at all. Hence, to compute the cross-asset correlation of 10 assets, for example, one would need to run SV filter 10 times (once for each asset). But once volatility is estimated, the 45 possible cross-asset correlations are obtained at negligible computational cost.

To summarize, this dynamic correlation algorithm is extremely rapid, yields reasonable estimates for large correlations (both positive and negative), but does not perform well for low correlations. This shortcoming is not due to the quality of the residuals data from the volatility estimation process, it is rather a true shortcoming of the correlation EMA. It might be possible to fix this by implementing a particle filter or a double-EMA, but this is left as future work.

3.4 Cross-Asset Correlations of Selected Securities

3.4.1 S&P 500 Vs. S&P TSX Stock Indices

We now turn to the analysis of the dynamic correlation of real-world securities. To begin, Figure 3.11 shows the correlation between the S&P 500 and S&P TSX indices. As expected, the correlation between the two stock indices is quite high (unconditional/static correlation around 0.8). However, the dynamic correlation has dropped significantly (beyond the 95% confidence interval) since the beginning of 2013. This is likely due to the relative overweigh position in materials stocks in the TSX and the relative overweigh position in tech stocks in the S&P 500 index. Also, one notices a sharp rise in correlation in late 2008. This is consistent with the financial wisdom that cross-asset correlation increases in times of crisis.

Figure 3.11 Correlation between the S&P 500 and the S&P TSX Stock Indices

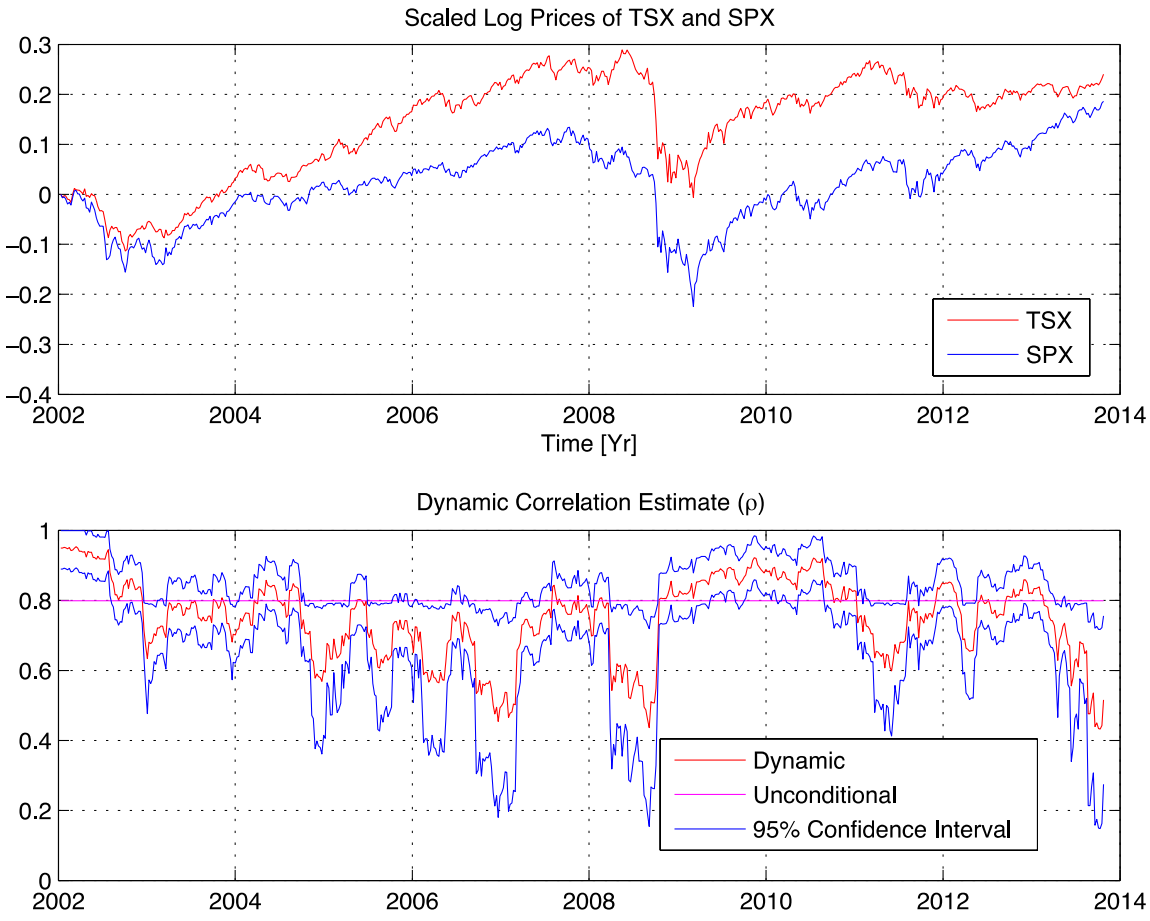

3.4.2 S&P 500 Index Vs. US Dollar Index

Figure 3.12 below shows the correlation between the S&P 500 index and the US dollar index. Common financial wisdom dictates that when stock markets suffer heavy losses, investors seek refuge in the US dollar. This is certainly reflected in the unconditional correlation of about -0.35 between the two securities. The correlation was particularly negative and severe (dynamic correlation hovering around -0.5) between 2009 and 2013. This may well have been due to the Quantitative Easing (QE) campaign pursued by the US Federal Reserve. As the central bank pumped capital into bond markets (lowering yields), investors turned to stocks, but lost confidence in the strength of the US dollar. As seen on the right-hand side of Figure 3.12, the data appears to say that this effect has become muted since early 2013. However, it is likely that the eventual tapering of QE would cause a resumption of negative correlation because QE tapering would likely cause a stock correction and an increase in the value of the US dollar index.

Figure 3.12 Correlation between the S&P 500 and the US Dollar Indices

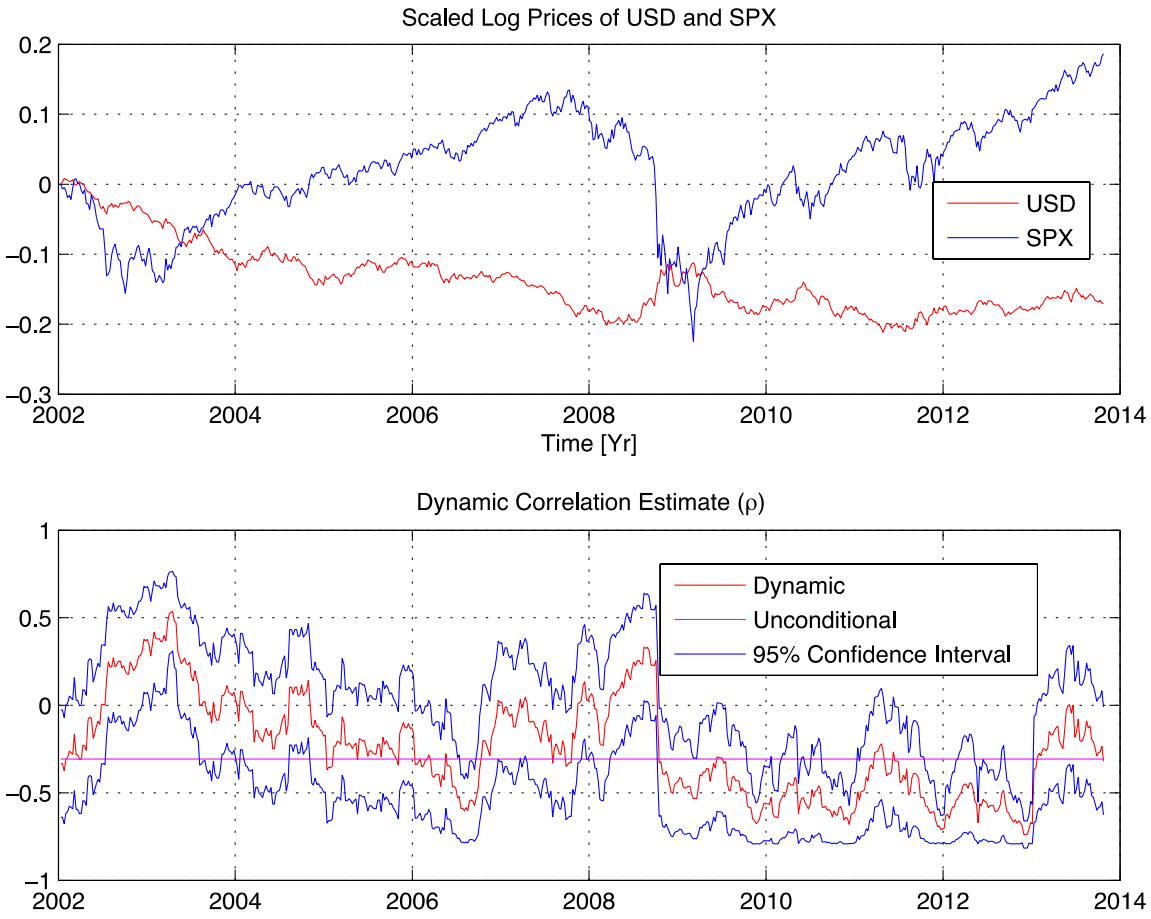

3.5 Risk Appetite Estimator

Figure 3.13 below shows the original RBC RAI, as well as the replicated RAI. The replication strategy is discussed in the theory section.

Figure 3.13: RBC and Replicated RBC Risk Appetite Indicators

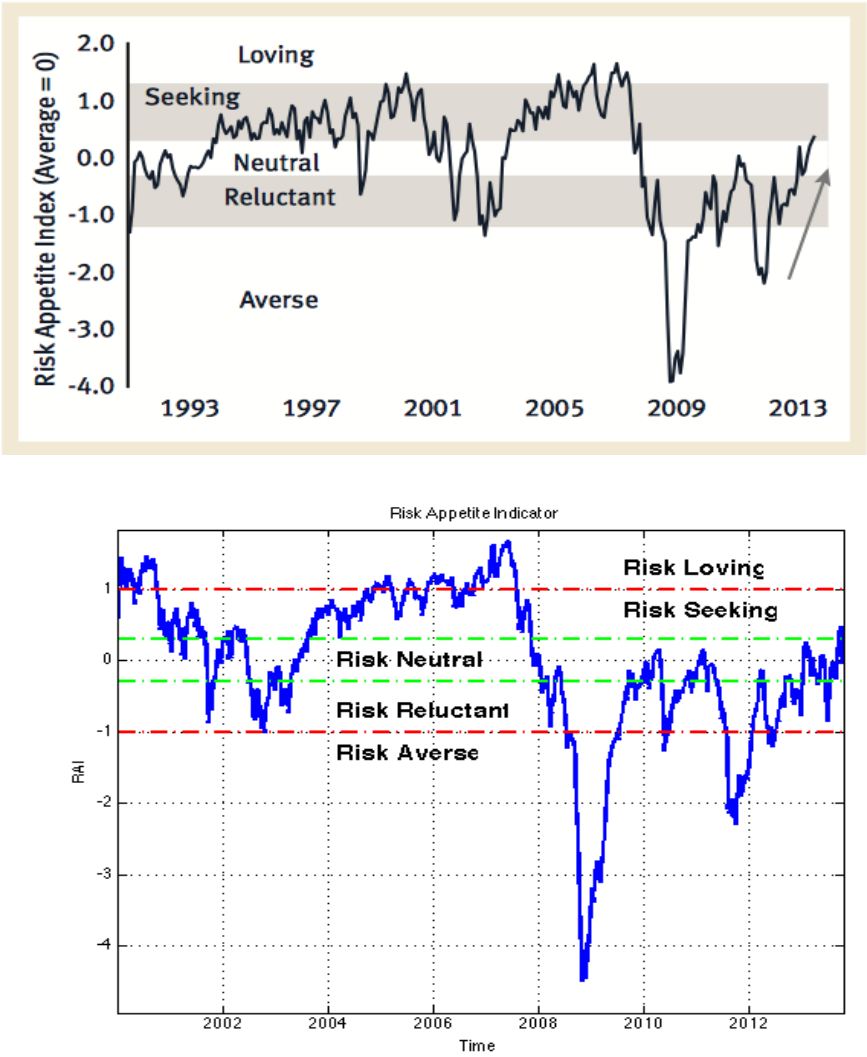

From the above plot, the replicated RAI closely resembles the RBC RAI from 2002 to until now. It shows the revived risk appetite from its lowest historical level in late 2008 to its highest level since the financial crisis. Furthermore, both the RBC and the replicated RAI show that the current appetite level is barely in risk seeking territory, meaning that the previous extreme risk aversion seems to come to the end. Risk appetite is now near its historically average.

The replicated RAI is considered validated.

4: Conclusions

4.1 Stochastic Volatility Estimation

Using the Residual Distribution Matching (RDM) method, the authors estimate the stochastic volatility of assets. This process differs from the traditional conditional likelihood maximization method. In addition, RDM gives more value to estimates with better in-sample volatility prediction power. The proposed model appears to outperform the famous GARCH model. Compared to GARCH, RDM:

1. Is smoother*;*
2. Does not lag the measurements;
3. Quickly responds to sudden decreases in volatility;
4. Provides superior in-sample volatility prediction capability.

In addition, the output of the proposed model is compared to that of a well-known stochastic model [Harvey et al., 1994]. Results hint that the proposed model provides better in-sample volatility prediction than the Harvey et al. model during periods of extreme and quiescent volatility.

However, the RDM method is slower than GARCH because it is based on Monte Carlo simulation (particle filter).

4.2 Dynamic Correlation

Using the residuals from the stochastic volatility fit, the model uses a simple exponential moving average filter to estimate dynamic (time-varying) correlation. Artificial simulation is then used to determine estimation error. Results show that this model (first proposed by the RiskMetrics group):

1. Provides good estimates of the dynamic cross-asset correlation when it is high;
2. Is not ideally suited for low-correlation environments;
3. Has a lag of 17 time steps.

It might be possible to improve this model by implementing a double exponential moving average filter or a particle filter.

4.3 Risk Appetite Indicator

Although 40% of the inputs of the RBC RAI could not be located, the authors successfully replicate the RAI. This indicator will provide SIAS managers with a quick overview of the overall level of risk in markets.

4.4 Future Work

The following is a compilation of the actions that could be pursued to improve the quality of the models:

1. Speed up model execution by replacing some of the computationally slow .m files by much faster .mex files. Computational speed could potentially be improved by as much as 90%;
2. Improve the dynamics correlation model, possibly through the use of a particle filter or of a double EMA filter;
3. Replace the SIS particle filter by a Markov Chain Monte Carlo particle filter to improve efficiency by reducing the number of required particles while staying away from filter degeneracy problems.

In addition, the SV and dynamic correlations could be used to:

1. Predict the active risk of passive and semi-active portfolios (see Section 3.3.1);
2. Price relatively short-maturity options or spread options, especially those with high correlation.

Finally, it would be possible to create an even faster (though somewhat less accurate) SV model by:

1. Doing away with the particle filter, which is the reason why the SV model is so slow in the first place;
2. Using a Nelder-Mead optimizer to optimize the three SV parameters in Section 2.2.1 using a objective function penalizing:
    a. Inadequate in-sample volatility prediction power;

b. Residuals distribution that is not close to the ideal residuals distribution;

c. High amount of noise in estimated volatility signal.

This new model would run a little slower than a GARCH estimator, but much faster than the SV model used here.

As for the RAI, it would be desirable to locate some of the missing inputs and include them in the replicated version of the RAI. In addition, the authors were forced to eye-ball some of the data from plots because numerical data was simply not available for some inputs. It would be desirable to obtain actual numerical data for these inputs.

Reference List

Anderson, R., and Moore, J., Optimal Filtering, Prentice-Hall, 1979.

Bollerslev, T., Generalized Autoregressive Conditional Heteroskedasticity, Journal of Economics, Vol. 31, pp.307-327, 1986.

Chatfield, C., The Analysis of Time Series: an Introduction, Chapman and Hall CRC, p. 229, 2004.

Chib, S., Nardari, F., Shephard, N., "Analysis of High Dimensional Multivariate Stochastic Volatility Models', Journal of Economics, Vol. 134, pp.341-371, 2006.

Crack, T.F., Heard on the Street: Quantitative Questions from Wall Street Interviews, 13th ed., p.161, 2012.

Doucet, A., Johansen, A., A Tutorial on Particle Filtering and Smoothing: Fifteen Years Later, The Oxford Handbook of Nonlinear Filtering, Ed. By Crisan, D., and Rozovskii, B., Section 8.2, 2009.

Engle, R.F., "Autoregressive Conditional Heteroscedasticity with Estimates of the Variance of United Kingdom Inflations", Econometrica, Vol. 50, pp.987-1007, 1982.

Engle, R.F., ""Dynamic Conditional Correlation: A Simple Class of Multivariate Generalized Autoregressive Conditional Heteroskedasticity Models", Journal of Business and Economic Statistics, Vol. 20, No. 3, pp. 338-350, 2002.

Gonzalez, E.A., Rodriguez, F.F., Rodriguez, J.P., Volatility Bias in the GARCH Model: a Simulation Study", working document number 2002-02, 2002.

Gordon, N., Salmond, D., Smith, A., "Novel Approach to Non-Linear/Non- Gaussian Bayesian State Estimation," IEEE-Proceedings-F, Vol. 140 No. 2, pp. 107-113, April 1993.

Harvey, A., Ruiz, E., Shephard, N., "Multivariate Stochastic Variance Models"

Heston, S., "A Closed-Form Solution for Options with Stochastic Volatility with Applications to Bond and Currency Options", The Review of Financial Studies, Vol. 6, No. 2, pp.327-343, 1993.

Illing, M., Aaron, M., "A Brief Survey of Risk-Appetite Indexes", Financial System Review, p.37-43, 2012.

Julier, S.J., and Uhlmann, J.K., "A New Extension of the Kalman Filter to Non-Linear Systems," Proceedings of AeroSense: The 11th International Symposium on Aerospace/Defence Sensing, Simulation and Control, Orlando, Florida, Vol. Multi Sensor Fusion, Tracking and Resource Management II, 1997.

Kalman, R., "A New Approach to Linear Filtering and Prediction Problems," ASME Journal of Basic Engineering, Vol. 82, pp., 35-45, March 1960.

Kim, J., Stoffer, D. S., "Fitting Stochastic Volatility Models in the Presence of Irregular Sampling Via Particle Methods and the EM Algorithm", Journal of Time Series Analysis, Vol. 29, No. 5, 2008, pp.811-833.

Kim, J., “Parameter Estimation in Stochastic Volatility Models with Missing Data using Particle Methods and the EM Algorithm”, Ph.D. dissertation, University of Pittsburgh, 2005.

Lascelles, E., “Risk Appetite Renaissance”, Economic Compass, RBC Global Asset Management, http://media.rbcgam.com/pdf/economic-compass/rbcgam-economic-compass-risk-appetite-201303.pdf, (accessed on November 23rd, 2013), 2013.

Shumway, R., Stoffer, D.S., Time Series Analysis and its Applications with R Examples, 2nd Ed., Springer, New York NY, p.388, 2006.

Simon, D., Optimal State Estimation, Wiley, pp.462, 469-70, 2006.

Tsay, R.S., Analysis of Financial time Series, 3rd ed., Wiley, p.111, 120-121, 140-153, 2010.

van der Merwe, R., Doucet, A., De Freitas, N., Wan, E., ”The Unscented Particle Filter,” Proceedings of Neural and Information Processing Systems, Vol. 13, 2000(a).

van der Merwe, R., Doucet, A., De Freitas, N., Wan, E., ”The Unscented Particle Filter,” Technical Report CUED/F-INFENG/TR 380, Cambridge University Engineering Department, 2000(b).

Yu, J., “On Leverage in Stochastic Volatility Model”, Journal of Economics, Vol. 127, pp.165-178, 2005.