An Econometric Analysis of the Volatility Risk Premium*

Jianqing Fan

jqfan@princeton.edu

Michael B. Imerman

mimerman@princeton.edu

Wei Dai

weidai@princeton.edu

July 2012

JEL Classification: C01, C58, G12

* Department of Operations Research & Financial Engineering, Princeton University, Princeton, NJ 08544. Email: mimerman@princeton.edu; Phone: (609) 258-1889.

An Econometric Analysis of the Volatility Risk Premium

ABSTRACT

This paper examines the volatility risk premium, defined as the difference between expected future volatility under the risk-neutral measure and the expectation under the physical measure. This risk premium represents the price of volatility risk in financial markets. In order to compute this risk premium, one must be able to accurately measure volatility under both measures. It is standard to use the VIX volatility index is used to proxy the expectation under the risk-neutral measure. Estimation under the physical measure is less straightforward. Using ultra-high-frequency transaction data on SPDR, the S&P500 ETF, we implement a novel approach for estimating integrated volatility on the frequency domain which allows us to isolate the biases from microstructure noise from the true volatility process. Once we compute the volatility risk premium, we perform a comprehensive econometric analysis of its determinants before, during, and after the Financial Crisis. The results indicate that the volatility risk premium is very sensitive to liquidity and distress in the financial system. We are able to link this statistical relationship to a recent literature in financial economics dealing with intermediation and liquidity provision in the options markets.

1 Introduction

The past decade has seen unprecedented swings in volatility. Consequently market participants are no longer just concerned with the level of volatility, but also the risk surrounding future levels of volatility. The desire to price and hedge volatility risk has lead to new advances in financial econometrics and risk management. One way to measure and study how volatility risk is priced in the market is to analyze the *volatility risk premium.* The volatility risk premium is defined as the amount by which some estimate of market-implied volatility exceeds (or falls short of) some measure of realized volatility. A more technical definition is the difference between expected future volatility under the physical measure and the risk-neutral measure.

In general, the volatility risk premium, as defined above, is thought to be negative which implies that the expected volatility under the pricing measure is greater than under the physical measure. The negative volatility risk premium indicates that investors are willing to pay a premium to hedge volatility risk; it further signifies that investors are risk averse.

Our contributions are twofold. First we use ultra-high-frequency data to estimate integrated volatility. It is widely known that ultra-high-frequency financial data is contaminated by microstructure noise. To address the microstructure noise we use a novel method based on the Fourier transform, which allows us to work on the on the frequency domain rather than time scale. This has the added benefit that the microstructure noise can be autocorrelated and so we need not restrict ourselves to the case where microstructure noise is independent over time. We then filter out the noise component to de-bias the volatility estimator. Simulations (in Appendix B) show that, under our method, sampling at higher frequencies allows for the most precise estimation of integrated volatility. Furthermore our method performs better than naïve subsampling rules that are typically used in high-frequency studies.

Then, after estimating the daily and monthly integrated volatility, we compare these measurements to the VIX volatility index and construct a time series of the volatility risk premium before, during, and after the recent Financial Crisis. We then use traditional risk factors, along with other financial market and macroeconomic variables including liquidity measures to try to understand the determinants of the volatility risk premium. An interesting result we find is that, for a period of time during the Financial Crisis, the volatility risk premium turns *positive*. While this may seem paradoxical, we are able to explain and interpret these results within the context of some recent literature linking the price of volatility in the markets to macroeconomic factors, liquidity, and financial intermediation.

Specifically, we find that the TED spread– a measure of liquidity and confidence in the financial sector – is best able to explain the positive spike in the volatility risk premium. We also find that open interest – a proxy for option demand – is a significant explanatory variable. Furthermore, the onset of the Financial Crisis appears to mark a persistent structural break in the relationship between TED spread and volatility risk premium as well as put option open interest and the volatility risk premium which is evidence of strong nonlinear effects of these variables with the price of volatility risk. We also find that term structure variables and credit spreads are significant in explaining the volatility risk premium, while traditional financial risk factors (Fama-French factors) have very little explanatory power.

The remainder of this paper is structured as follows. The next section, Section 2, reviews the volatility risk premium. Then, in Section 3, we discuss estimating integrated volatility using high-frequency data. Section 4 contains our empirical analysis including data and econometric specifications. In Section 5 we present our results with discussion about the economic interpretations and implications. Section 6 concludes. We have two Appendices: Appendix A covers the technical details on the Fourier transform method that we use to address microstructure noise in our estimation of integrated volatility with the ultra-high-frequency data; Appendix B presents some simulations that demonstrate the benefit of using ultra-high-frequency data in estimating integrated volatility and the extent to which our method performs better.

Figure 1: Time Series of Volatility Risk Premium

Appendix

A The Fourier Transform Method

A.1 Fourier Domain Representation

To simplify the notion, the drift will not be considered, as it only accounts for a term of order $\log(\Delta t)\sqrt{\Delta t}$. First define the discrete Fourier transform of the increment process $\Delta U_{t_j} = U_{t_{j+1}} - U_{t_j}$ of a sample from a generic time series $U_{t_j}, j = 1, \cdots, N$,

$$J_k^{(U)} = \sqrt{\frac{1}{N}} \sum_{j=1}^{N} \Delta U_{t_j} e^{-2\pi i t_j f_k}, f_k = \frac{k}{T}, U = X, Y, \epsilon. \tag{1}$$

The first and second order structures of $\{J_k^{(X)}\}_k$ are as follows:

$$\begin{aligned} \mathbb{E}\{J_k^{(X)}\} &= 0, \\ \varphi_{k_1,k_2}^{(X)} &\equiv Cov\{J_{k_1}^{(X)}, J_{k_2}^{(X)}\} \end{aligned} \tag{2}$$

In particular,

$$\varphi_{k,k}^{(X)} = \frac{1}{N} \int_0^T \mathbb{E}\{\sigma_s^2\} ds \tag{3}$$

A.1.1 Debias

Now comes the debias procedure. By calculating the first and second order structures of $\{J_k^{(X)}\}_k$, an oracle shrinkage estimator (in which $\varphi_{k,k}^{(X)}$ is unknown) would be

$$L_k = \frac{\varphi_{k,k}^{(X)}}{\varphi_{k,k}^{(X)} + a^2 |2\sin(\pi f_k \Delta t)|^2)}, \tag{4}$$

which yields an estimator $\widehat{\langle X, X\rangle}_T^{(L_k)} = \sum_{k=0}^{N-1} L_k |J_k^{(Y)}|^2$. To estimate L_k, the Whittle log-likelihood is proposed, such that

$$\widehat{L_k} = \frac{\hat{\sigma}_X^2}{\hat{\sigma_X^2} + \hat{a}^2 |2\sin(\pi f_k \Delta t)|^2}, \tag{5}$$

where $(\hat{\sigma}_X^2, \hat{a}^2)$ comes from maximizing the modified Whittle log-likelihood

$$l(\sigma_X^2, a^2) = -\sum_{k=1}^{N/2-1} \log(\sigma_X^2 + a^2|2\sin(\pi f_k \Delta t)|^2) - \sum_{k=1}^{N/2-1} \frac{|J_k^Y|^2}{\sigma_X^2 + a^2|2\sin(\pi f_k \Delta t)|^2}, \tag{6}$$

which gives a final estimator

$$\widehat{\langle X, X\rangle}_T^{(\widehat{L_k})} = \sum_{k=0}^{N-1} \widehat{L_k} |J_k^{(Y)}|^2. \tag{7}$$

A.1.2 Properties

It was proved that (first is asymptotically unbiasedness, and second is consistency)

$$\mathbb{E}\{\widehat{\langle X, X\rangle}_T^{(\widehat{L_k})}\} \quad = \mathbb{E}\{\textstyle\int_0^T \sigma_t^2 \, dt\} + O(\Delta t^{1/4}), \tag{8}$$

$$\{\widehat{\langle X, X\rangle}_T^{(\widehat{L_k})}\} \quad = \textstyle\int_0^T \sigma_t^2 \, dt + O_p(\Delta t^{1/4}). \tag{9}$$

B Simulations

We simulate data using a Heston (1993) model (following the simulation in Olhede, Sykulski, and Pavliotis (2009)), and compare the performance of Fourier method and naive subsampling at different sampling frequency. The Heston (1993) model is specified as:

$$\begin{aligned} dX_t &= (\mu - \nu_t/2)dt + \sigma_t dB_t, \\ d\nu_t &= \kappa(\alpha - \nu_t)dt + \gamma \nu_t^{1/2} dW_t, \end{aligned}$$

where $\nu_t = \sigma_t^2$. The parameters are set as follows: $\mu = 0.05$, $\kappa = 5$, $\alpha = 0.04$, $\gamma = 0.5$, and the correlation between the two Brownian motions B_t and W_t is $\rho = -0.5$.[1] The initial values are $X_0 = 0$ and $\nu_0 = 0.04$. We take T as one day, and simulate data with $\Delta_t = 0.1$s, which yields a sample path of length $N = 234,000$ in one trading day. We first calculate the underlying true integrated volatility by a Riemann sum approximation of the integral, i.e.: $\frac{T}{N}\sum_{i=1}^{N}\sigma_i^2 = \int_0^T \sigma_t^2 dt$. Then we add iid noise $\epsilon \sim \mathcal{N}(0, \sigma_\epsilon^2)$ to get the observed data $Y_i = X_i + \epsilon_i$, where we set $\sigma_\epsilon = 5 \times 10^{-4}$.

We estimate the integrated volatility using two methods, the Fourier method and the naive subsampling, which yields $< X, X >_T^{Fourier}$ and $< X, X >_T^{subsampling}$. We calculate the RMSE (root-mean-square error) of the estimates to the truth over 200 simulated sample paths. To further illustrate the effect of high frequency data, we evaluate two methods from $\Delta_t = 1$s up to $\Delta_t = 600$s. Below is a figure showing the RMSE of the Fourier method and the naive subsampling against decreasing sampling frequencies.

The takeaway of this figure is two folds. First, the Fourier method can effectively filter the microstructure noise, and works better than naive sampling method (we didn't implement other more sophisticated methods for comparison, as the simulation is not to illustrate Fourier method is superior, but rather to justify the use of high frequency data). Second, if we are able to filter the microstructure noise, higher frequency gives us a better estimate as we are able to utilize more data (hence more information).

[1] These are the same as those used in the Olhede, Sykulski, and Pavliotis (2009) simulations.

References

Heston, S. L., 1993, “A closed-form solution for options with stochastic volatility with applications to bond and currency options,” *Review of Financial Studies*, 6(2), 327–343.

Olhede, S., A. Sykulski, and G. Pavliotis, 2009, “Frequency Domain Estimation of Integrated Volatility for Ito Processes in the Presence of Market-Microstructure Noise,” *SIAM - Multiscale Modeling and Simulation*, 8(2), 393–427.

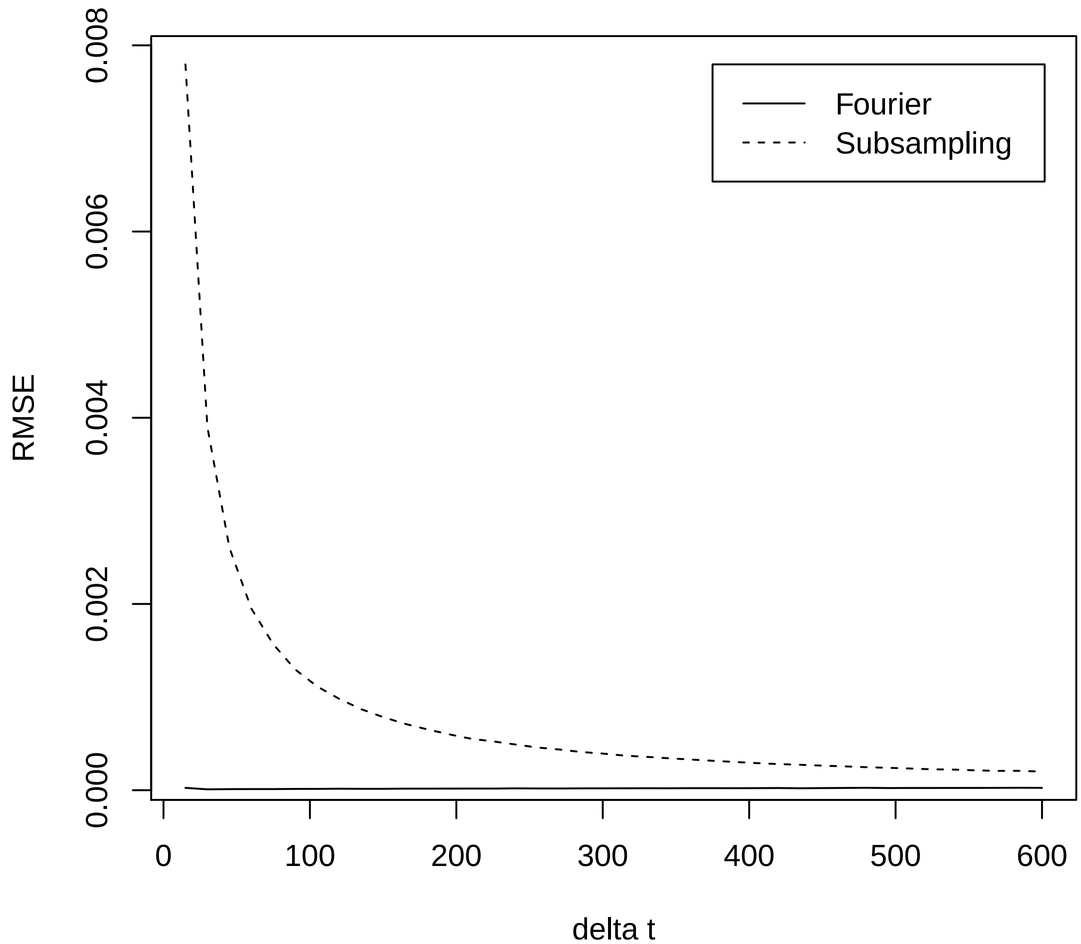

Figure 2: RMSE of Two Methods