ARE STOCK RETURN DYNAMICS TRULY EXPLOSIVE OR MERELY CONDITIONALLY LEPTOKURTIC? A BRIEF CASE STUDY ON THE IMPORTANCE OF DISTRIBUTIONAL ASSUMPTIONS IN ECONOMETRIC MODELING (ABRIDGED)

Peter A. Ammermann, College of Business Administration, California State University – Long Beach, 1250 Bellflower Boulevard, Long Beach, CA, 90840, 562-985-7526, peter.ammermann@csulb.edu

ABSTRACT

This paper uses the estimation of the Self-Excited MultiFractal (SEMF) model as a case study to explore the impact of distributional assumptions on the model-fitting process. It is found that incorporating reasonable distributional assumptions including a non-zero mean and the leptokurtic Student's t distribution can have a substantial impact on the estimation results that are obtained for the model's parameters and can mean the difference between parameter estimates that imply unstable and potentially explosive volatility dynamics versus ones that describe more reasonable and realistic dynamics for the returns. The form of the model specification can also have a substantial impact on the estimated value for the mean for the process.

INTRODUCTION

Efforts to develop statistical models of stock and other financial returns have been undertaken for well over a century, since at least the completion in 1900 of Louis Bachelier's dissertation entitled "Théorie de la spéculation", in which he modeled stock returns as a "random walk", which, in discrete time, can be described as:

$$R_t = \mu + \varepsilon_t, \tag{1}$$

where the ε_t are presumed to be a series of iid standard Normal variables. Unfortunately, while this model can be useful for conceptual purposes, it is sorely inadequate for statistical modeling and risk assessment purposes, because the statistical characteristics exhibited by real-world stock returns are much too complicated to be adequately captured by such a simple model.

Among the statistical characteristics generally found in financial returns that complicate their adequate modeling are a broad range of features including:

- fat tails in the distributions of financial returns,
- a lack of autocorrelation within the returns themselves, but
- long memory and clustering within the volatility of the returns,
- multifractal properties,
- time reversal asymmetry and the "leverage effect," and
- "bubbles" and "crashes."

The development of the Generalized AutoRegressive Conditional Heteroskedasticity (GARCH) family of time-series statistical models [1,2,3] was a step in the right direction in improving upon the simple random walk model, but such models can still account for only a portion of the variety of stock return characteristics listed above. More recently, and especially over the past decade, an alternative that has shown promise and gained increasing attention are multifractal models, which are defined by the scaling properties of their moments, and a variety of such models have been developed, including the Multifractal Model of Asset Returns (MMAR), the Poisson Multifractal Model (PMM), the Markov-Switching Multifractal (MSM) model, the Multifractal Random Walk (MRW) model, and Quasi-Multifractal (QMF) model [4,5,6,7,8,9,10,11,12,13], and, most recently, the Self-Excited Multifractal (SEMF) model [14,15,16].

THE SELF-EXCITED MULTIFRACTAL MODEL

The SEMF model incorporates a conditional volatility with the following form:

$$\ln \sigma_t = \ln \sigma_0 - \frac{1}{\sigma_0} \sum_{\tau=0}^{t-1} h_{t-\tau-1} R_\tau. \tag{2}$$

The SEMF model can readily be viewed as a hybrid between the other multifractal stochastic volatility models and the GARCH-type models. The key difference from those models comes from the SEMF's memory kernel, h_τ, which has a form such as:

$$h_\tau = h_0 e^{-\varphi \tau}, \tag{3}$$

for an exponential kernel or:

$$h_\tau = h_0 \tau^{-\varphi - \frac{1}{2}}, \tag{4}$$

for a power-law kernel.

Within the process, the parameter h_0, where $0 \leq h_0 < 1$, controls the level of multifractality of the process and can be considered a non-stationarity parameter – the closer it is to zero, the more stable and constant the process is; the closer it is to one, the greater the degree of non-stationarity of the process. This parameter also serves the double-duty of controlling the "leverage" or "skewness" effect within the process – if h_0 is positive, then negative returns will have a larger impact on subsequent levels of volatility than will positive returns, and vice versa. The parameter φ represents the strength and direction of the memory of the process – the greater is φ, the more quickly the memory of the process tampers down; the smaller is φ, the longer the memory of the process lasts and the longer it takes for the effects of a volatility burst to dissipate; and, if φ is negative, then the effects of a volatility burst will actually increase over time, becoming even greater in magnitude as the process moves further away in time from the instigating event.

Compared to GARCH-type models and the other multifractal stochastic volatility models, the SEMF model would appear to have the most promise for empirical applications. Unfortunately, applications of this model to real-world financial data have seen mixed results. On the one hand, one researcher, Zeeuw van der Laan [17], estimates the SEMF model for the daily returns from a variety of financial time series from around the world and obtains extremely unrealistic parameter estimates, including very large estimated values for h_0 and small but negative estimated values for φ that suggest that such returns would be highly non-stationary and subject to explosive volatility patterns. This leads Zeeuw van der Laan to reject the original specification of the SEMF model as a potential model for financial returns and to propose an alternative specification with a conditional volatility equation that could be viewed as a variation of an EGARCH model [18] to replace the original SEMF model. On the other, another team of researchers, Zhong and Zhao [19], estimate the SEMF model for daily index returns for the Shanghai and Hong Kong stock exchanges and obtain reasonable estimates for the parameters of this model, with small values of around 0.05 for the non-stationarity parameter h_0 and small but positive values for the memory dissipation parameter φ.

What is going on? How can these two conflicting sets of results for this otherwise seemingly promising model be reconciled? Is the SEMF model a reasonable and feasible model for stock returns, or is it not? Ultimately, taking both of these two sets of results into account, it appears that the resolution to this question is more likely to lie not in the specification of the conditional volatility equation, which is where Zeeuw van der Laan looks in developing his alternative model, but rather in a more basic area that Zeeuw van der Laan appears to overlook but to which Zhong and Zhao do devote significant attention – the form of the conditional probability distribution underlying the model.

LEPTOKURTOSIS VS. HETEROSKEDASTICITY AND SEMF MODEL PARAMETERS

It is a well-known result that a process whose increments are each Normally distributed but that reflect different degrees of variance will, under aggregation, exhibit excess kurtosis. In other words, leptokurtosis can be driven by heteroskedasticity across the increments of the process. Thus, in spite of the unconditional leptokurtosis observed in the data for which their models were developed, Engle's original ARCH specification as well as both Filimonov and Sornette's original specification of the SEMF model and Zeeuw van der Laan's application of it all make the typical assumption that the increments of their processes are nonetheless conditionally Normally distributed. I.e., they assume that all of the observed leptokurtosis within their data sets is driven entirely by the conditional heteroskedasticity that their respective models are capturing. As a consequence, though, this means that their model specifications require the incorporated conditional volatility equations to be able to simultaneously account for not only all of the observed memory but also all of the observed leptokurtosis within the processes to which their models are fit. In contrast, Zhong and Zhao's approach appears to follow a two-stage process in which they start their analysis by focusing on the original return series (not the conditional innovations from the model) and then fitting a leptokurtic distribution, the generalized tempered stable distribution, to the original (or raw) returns themselves. Zhong and Zhao then use these fitted distributional parameter values in the likelihood function used to estimate the other SEMF model parameters.

But, because Zhong and Zhao fit the extant leptokurtosis first, before the SEMF model parameters are estimated, this could lead to the opposite problem from what Zeeuw van der Laan faces. Under Zeeuw van der Laan's applications in which conditional Normality is assumed, the SEMF model parameters are required to perform double-duty – accounting for not just the volatility memory within the returns but also all of the leptokurtosis among them. Under Zhong and Zhao's application, however, the SEMF model parameters are required to perform only a single function and are actually precluded from accounting for any of the observed leptokurtosis, even that which should occur naturally as a consequence of the heteroskedasticity for which the SEMF model is supposed to account.

In other words, in comparing the two we are ultimately faced with a "Goldilocks" type of situation in which Zeeuw van der Lan's SEMF model parameters must account for too much leptokurtosis while Zhong and Zhao's SEMF model parameters must account for too little. Thus, the differences in the degree of "reasonableness" of the respective parameter estimates that are obtained could clearly be driven by the role that leptokurtosis plays in the model specification and estimation process, and the correct response to resolving the question of whether the model is reasonable or not is to incorporate a potentially leptokurtic conditional distribution, such as the Student's t distribution, into the SEMF model specification and to jointly and simultaneously estimate the Student's t distributional parameters in conjunction with the SEMF volatility equation parameters so that the latter set of parameters can be allowed to account for just the right amount of leptokurtosis within the set of financial returns. This would further allow us to gain a clearer picture of the extent to which the leptokurtosis in financial returns is driven by the conditional heteroskedasticity within these returns (what the volatility equation accounts for) versus being an inherent characteristic of these returns that is driven by something deeper.

RESULTS FOR SPY RETURNS

In order to explore this issue further, two variations of an exponential kernel SEMF model, one the original version incorporating a Normal conditional distribution and the other version incorporating a Student's t conditional distribution, are fit to a representative financial time series, daily returns for the SPDR ETF (SPY) from 1 February 1993 through 8 June 2015. The Student's t distribution was chosen both because it is widely used in financial applications and because it incorporates the Normal distribution as a limiting case, as the degrees of freedom for the distribution tends toward infinity.

The parameter estimates obtained for the SPY returns under the conditionally Normal specification are as follows:

Parameter	Estimated Value	Standard Error	t-statistic	p-value
h_0	0.7405	287.2350	0.0026	0.9979
φ	-0.0005	0.4711	-0.0010	0.9992
σ_0	1.7432	2395.6400	0.0007	0.9994

As can be seen in this table, the parameter estimates obtained are similar to those found by Zeeuw van der Laan. Specifically, the estimate of h_0 is very large, suggesting a highly non-stationary process, while the estimate of φ is actually negative (albeit insignificantly so), suggesting an explosive volatility process.

The estimates obtained under the conditionally Student's t SEMF specification, on the other hand, paint a very different picture:

Parameter	Estimated Value	Standard Error	t-statistic	p-value
h_0	0.0716	0.000014	4961.33	0.0000
φ	0.0191	0.000005	3827.87	0.0000
σ_0	0.0120	0.000000	283912.00	0.0000
ν	4.5279	0.000379	11945.80	0.0000

These latter results are more consistent with those found by Zhong and Zhao. Specifically, the parameter estimates under the conditionally Student's t specification are all found to have reasonable values, and are all of these parameter values are also found to be statistically significant. Among these, h_0 is found to be relatively small, indicating a relatively low degree of non-stationary, and the estimate of φ is of somewhat greater magnitude (as compared to the value found under the Normal SEMF specification) and, much more importantly, is positive, which indicates that the memory within the volatility of the process dissipates over time, although at a relatively slow rate, in contrast to being explosive with an impact that expands in magnitude over time, as would be suggested by a negative value.

In addition, the unconditional standard deviation of the process, σ_0, is estimated at 0.0120, which is relatively close to the standard deviation for the original returns, 0.0119. Under the Normal SEMF specification, this parameter was estimated to be nearly 150 times larger, suggesting that the estimated conditional volatility equation under that specification was having a difficult time tracking the volatility of the returns. This implies that, since the conditionally Normal specification could not directly account for the extreme degree of kurtosis within the original returns, the attempt by the model to nonetheless fit the large number of extreme returns within the data leads instead to a dramatic inflation or explosion of the estimated volatility, and especially the unconditional volatility, to try to account for these extreme returns. In other words, since the model could not do an adequate job of fitting the tails of the distribution (which is the key driver of what kurtosis, the 4^{th} central moment, measures), it overcompensates by instead adjusting the estimated scale of the distribution (which is what standard deviation, or the 2^{nd} central moment, measures). Moreover, as a consequence of the extremely high magnitude of the unconditional standard deviation, σ_0, that is estimated for the conditionally Normal specification, the magnitude of any conditional heteroskedasticity effects that are driven by past returns around this unconditional volatility would be proportionately miniscule, so the apparent explosively nonstationary dynamics that are suggested by the estimated values of h_0 and φ under that specification would, in effect, in comparison to the estimated value of σ_0, be merely a "tempest in a teapot." The extreme extent of the conditional leptokurtosis that actually exists within the SPY returns is indicated by the very low estimated value for ν, the degrees of freedom parameter, under the Student's t SEMF specification. While the Student's t distribution tends toward the Normal for high degrees of freedom, as the degrees of freedom approach four or below, the process becomes so extremely leptokurtic that the moments of the distribution cease to be defined.

CONCLUSIONS

These initial results clearly suggest both the promise of the SEMF model as a model for financial time series as well as the significant importance that the form of the conditional probability distribution that is incorporated into the analysis can play in conducting econometric analysis and interpreting the results obtained from the modeling of financial time series, for which the distributional and volatility memory assumptions have a complicated and interactive relationship. Additional research documented in [20] finds that incorporating the estimation of the mean into the model can also dramatically impact both the value that is estimated for the mean itself as well as the estimated values that are obtained for the other model parameters.

REFERENCES

[1] Engle, R. Autoregressive conditional heteroskedasticity with estimates of the variance of U.K. inflation. *Econometrica*, 1982, 50, 987-1008.
[2] Bollerslev, T., Chou, R. and Kroner, K. ARCH modeling in finance: a review of the theory and empirical evidence, *Journal of Econometrics*, 1992, 52, 5–59.
[3] Bollerslev, T. Generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics*, 1986, 31, 307-327.

[4] Mandelbrot, B., Fisher, A. and Calvet, L. A multifractal model of asset returns. Cowles foundation discussion paper #1164, 1997.
[5] Calvet, L. and Fisher, A. Forecasting multifractal volatility. *Journal of Econometrics*, 2001, 105 (1), 27-58.
[6] Calvet, L. and Fisher, A. How to forecast long-run volatility: Regime-switching and the estimation of multifractal processes. *Journal of Financial Econometrics*, 2004, 2 (1), 49-83.
[7] Calvet, L. and Fisher, A. Multifractal volatility: Theory, forecasting, and pricing, Academic Press, 2008.
[8] Bacry, E., Delour, J. and Muzy, J.-F. Multifractal random walk. *Physical Review E*, 2001, 64 (2), 026103+.
[9] Muzy, J.-F. and Bacry, E. Multifractal stationary random measures and multifractal random walks with log infinitely divisible scaling laws. *Physical Review E*, 2002, 66, 056121.
[10] Muzy, J.-F., Delour, J. and Bacry, E. Modeling fluctuations of financial time series: from cascade process to stochastic volatility model. *European Physical Journal B*, 2000, 17, 537-548.
[11] Saichev, A. and Sornette, D. Generic multifractality in exponentials of long memory processes. *Physical Review E*, 2006, 74 (1), 011111+.
[12] Saichev, A. and Filimonov, V. On the spectrum of multifractal diffusion process. *Journal of Experimental and Theoretical Physics*, 2007, 105 (5), 1085-1093.
[13] Saichev, A. and Filimonov, V. Numerical simulation of the realizations and spectra of a quasi-multifractal diffusion process. *JETP Letters*, 2008, 87 (9), 506–510.
[14] Filimonov, V. and Sornette, D. Self-Excited Multifractal Dynamics. *Europhysics Letters*, 2011, 94 (4), 46003.
[15] Filimonov, V. Modeling financial time series: Effective multifractality and self-excited multifractal process. Paper presented at Fractals and Related Fields II, Porquerolles Island, France, 2011. http://congres-math.univ-mlv.fr/sites/congres-math.univ-mlv.fr/files/FIlimonov.pdf
[16] Sornette, D., Malevergne, Y. and Muzy, J.-F. Volatility fingerprints of large shocks: Endogenous versus exogenous, in Application of Econophysics, Proceedings of the second Nikkei symposium on econophysics, ed. Takayasu, H., Springer Verlag, 2004.
[17] Zeeuw van der Laan, P. Application of the SEMF process to financial time series. Masters thesis, 2011. http://www.er.ethz.ch/publications/MAS_Zeeuw-van-der-Laan_Final_version_nov_2011.pdf
[18] Nelson, D. Conditional hetereoskedasticity in asset returns: a new approach. Econometrica, 1991, 59 (2), 347-370.
[19] Zhong, J. and Zhao, X. Modeling complicated behavior of stock prices using discrete self-excited multifractal process. *Systems Engineering Procedia*, 2012, 3, 110-118.
[20] Ammermann, P. Are Stock Return Dynamics Truly Explosive or Merely Conditionally Leptokurtic? A Case Study on the Impact of Distributional As-sumptions in Econometric Modeling. *Journal of Data Analysis and Information Processing*, Forthcoming.