

Are Stock Return Dynamics Truly Explosive or Merely Conditionally Leptokurtic? A Case Study on the Impact of Distributional Assumptions in Econometric Modeling

Peter A. Ammermann

Department of Finance, California State University-Long Beach, Long Beach, CA, USA
Email: peter.ammermann@csulb.edu

Received 10 December 2015; accepted 22 February 2016; published 25 February 2016

Copyright © 2016 by author and Scientific Research Publishing Inc.
This work is licensed under the Creative Commons Attribution International License (CC BY).
<http://creativecommons.org/licenses/by/4.0/>



Open Access

Abstract

This paper uses the estimation of the Self-Excited Multi Fractal (SEMF) model, which holds theoretical promise but has seen mixed results in practice, as a case study to explore the impact of distributional assumptions on the model fitting process. In the case of the SEMF model, this examination shows 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 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. While the original SEMF model specification is found to yield unrealistic results for most of the series of financial returns to which it is applied, the results obtained after incorporating the Student's t distribution and a mean component into the model specification suggest that the SEMF model is a reasonable model, implying realistic return behavior, for most, if not all, of the series of stock and index returns to which it is applied in this study. In addition, reflecting the sensitivity of the sample mean to the types of characteristics that the SEMF model is designed to capture, the results of this study also illustrate the value of incorporating the mean component directly into the model and fitting it in conjunction with the other model parameters rather than simply centering the returns beforehand by subtracting the sample mean from them.

Keywords

Multifractal, Leptokurtosis, Conditional Heteroskedasticity, Maximum-Likelihood Estimation, Statistical Adequacy

1. Introduction

The problem of how best to model stock return dynamics is one that researchers have been studying for well over a century. One of the more recently developed models to describe these dynamics is the Self-Excited Multi-Fractal (SEMF) model of Filimonov and Sornette (2011). But, despite the theoretical promise shown by this model, its application to real-world sets of financial returns has yielded mixed results. The present paper uses the fitting of the SEMF model to real-world financial data to provide a case study in the potentially significant impact that distributional assumptions can have on the estimation process, as well as the importance of incorporating the estimation of key parameters, such as the mean of the process, into the model-fitting process.

The results obtained from this examination suggest that the SEMF model, as re-specified to incorporate such distributional assumptions, is a viable model for financial returns and can provide a reasonable description of the volatility dynamics for a variety of stock and index return series. As a side result, the examination also illustrates the value, due to the potentially confounding influence among various characteristics of interest, of incorporating potentially relevant features, such as the mean for the process, directly into the model building process itself rather than relying on pre-processing of the returns to remove the effects of such features.

2. Complications with Modeling Financial Returns

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." In discrete time, such a process can be described as:

$$R_t = \mu + \varepsilon_t \quad (1)$$

where the sequence of innovation or error terms $\{\varepsilon_t\}$ are presumed to be iid Normal random variables with constant variance σ^2 (*i.e.*, $\varepsilon_t \sim \text{NIID}(0, \sigma^2) \forall t$). 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.

Included among the wide variety of statistical characteristics (or "stylized features") generally found in financial returns that complicate their adequate modeling are the following:

- leptokurtosis ("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 returns (*i.e.*, strong autocorrelation among squared returns or absolute values of returns);
- multifractal properties;
- time reversal asymmetry and the "leverage effect"; and
- "bubbles" and "crashes."

A key improvement over the random walk model was provided in 1982 by Engle's Auto Regressive Conditional Heteroskedasticity (ARCH) model [1], which ultimately earned a Nobel prize for its developer and, in the intervening years, quickly spurred on the development of numerous variations (see, e.g., [2]), of which probably the most commonly applied variation is the GARCH (1, 1) model of Bollerslev [3]. Under this model, the innovations ε_t are not presumed to be iid but instead exhibit a conditional variance that is a linear function of both the previous period's conditional variance and the realized value of the previous period's innovation:

$$\sigma_t^2 = \alpha_0 + \alpha_1 \varepsilon_{t-1}^2 + \beta_1 \sigma_{t-1}^2. \quad (2)$$

One of the benefits of GARCH-type models is that they can account for at least some of the leptokurtosis within financial returns, since heteroskedastic processes exhibit leptokurtosis under aggregation. Nonetheless, the degree of leptokurtosis observed within financial returns is typically much greater than what can be explained by GARCH effects alone, so leptokurtic conditional distributions for the sequence of innovations are typically incorporated into the model in practice, with the Student's t being probably the most frequently used conditional distribution.

The GARCH family of models was clearly a step in the right direction toward providing an adequate statistical description of financial return dynamics, but such models can still account for only a portion, primarily the top two, of the variety of stock return characteristics listed above.

3. The Self-Excited Multi-Fractal Model as a Potential Model of Financial Returns

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) [4], the Poisson Multifractal Model (PMM) [5], the Markov-Switching Multifractal (MSM) model [6] [7], the Multifractal Random Walk (MRW) model [8]-[10], the Quasi-Multifractal (QMF) model [11]-[13], and, most recently, the Self-Excited Multifractal (SEMF) model [14] [15]. This latter model is a discrete-time model that incorporates a conditional volatility that is specified as follows:

$$\ln \sigma_t = \ln \sigma_0 - \frac{1}{\sigma_0} \sum_{\tau=0}^{t-1} h_{t-\tau-1} R_\tau. \quad (3)$$

In contrast to other multifractal models, the conditional volatility for the SEMF model is assumed to be endogenously rather than exogenously driven (see [16] for a discussion of the significance of this distinction), and the SEMF model can thus be viewed as a hybrid between the other multifractal stochastic volatility models and the more traditional GARCH-type models. A key difference between the SEMF model and the GARCH-type models, though, is that the former uses the returns themselves and retains the information about the sign of the returns while the latter typically only use information only about the magnitudes of the returns (e.g., squared or absolute values of the returns or return residuals; the EGARCH model [18] is a notable exception that incorporates information about both the sign and the magnitudes of the returns).

With regard to the structure of this equation defining the evolution of the conditional volatility, a key component is the SEMF's memory kernel, or h_τ , which is proposed to take one of three forms—either an exponential kernel:

$$h_\tau = h_0 e^{-\varphi\tau} \quad (4)$$

or a power-law kernel:

$$h_\tau = h_0 \tau^{-\varphi - \frac{1}{2}} \quad (5)$$

or a constant kernel:

$$h_\tau \equiv h_0. \quad (6)$$

Taking both the memory kernel and the rest of the conditional volatility equation into account, the three key parameters of the SEMF model include h_0 , φ , and σ_0 :

- h_0 , which ranges between 0 and 1 and controls the level of multifractality of the process, 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.
- φ represents the strength and direction of the memory of the process—the greater is φ , the more quickly the memory of the process recedes or tapers down; the smaller is φ , the longer the memory of the process persists 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 not only persist but will actually increase over time, reaching increasingly greater magnitudes as the process moves farther away in time from the initiating event.
- σ_0 sets the amplitude of the impact of the external flow of news on the returns being modeled, in contrast to the summation term $\omega_t = \sum_{\tau=0}^{t-1} h_{t-\tau-1} R_\tau$, which represents the extent to which the amplitude of the next, incremental return is dependent upon past realizations of the return series (hence making the dynamics “self-excited”), and which is specified as a non-decreasing, linear, deterministic function of these past returns.

As Filimonov and Sornette [14] [14] show, in theory the SEMF model can account for every single one of the “stylized features” of financial return dynamics that were shown above. Thus, compared to both the GARCH-type models as well as the other multifractal stochastic volatility models, the SEMF model would appear to have the most promise for empirical applications.

4. The Self-Excited Multi-Fractal Model in Practice

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 fi-

nancial 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 would suggest that such returns are highly non-stationary and subject to explosive volatility patterns. This leads Zeeuw van der Laan to reject the SEMF as a suitable model for financial returns and to propose an alternative “Generalized-SEMF” model, incorporating a modified conditional volatility equation that can be viewed as a restricted form of that of the EGARCH model (see [18]), as a substitute for the SEMF model. On the other hand, though, another team of researchers, Zhong and Zhao [19], fit the SEMF model to daily index returns for the Shanghai and Hong Kong stock exchanges and obtain reasonable estimates for the parameters of this model for both sets of index returns, 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 , suggesting stable processes incorporating a degree of memory that is very long-term but that does slowly dissipate over time, as would be expected of a reasonable model for financial returns.

So, 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? One possible answer lies in the specific sets of returns that the two sets of researchers examined. Zeeuw van der Laan examines a variety of financial time series, including daily returns for IBM, Shell, General Electric (GE), Coke, and the DJIA and S&P 500 indices. He also examines half-hourly returns for IBM, GE, the DJIA, and the US Dollar/Euro, Yen/US Dollar, and Swiss Franc/Euro exchange rates. Zhong and Zhao examine only two time series, including daily returns for both the Hang Seng Index (HSI) and the Shanghai Stock Exchange (SSE) Composite Index. In terms of the types of returns the two studies examine, the former focuses on returns for older, larger companies and more developed markets, while the latter studies the returns for newer, more emerging markets, especially in the case of the SSE. Thus, intuitively, if any of the series were likely to exhibit explosive dynamics, it would be the latter series, for the SSE, while the sets of returns for the older and more established companies and markets studied by Zeeuw van der Laan would seem more likely to exhibit more stable dynamics, rather than the other way around.

5. Leptokurtosis versus Heteroskedasticity and the Fitting of SEMF Models

Instead, a much more likely explanation for the differences between the two sets of results is suggested by the differences in the distributional assumptions made in the estimation or fitting of the SEMF model to these sets of returns. Probably the most prominent of the stylized features of financial returns is the first feature listed, leptokurtosis, or fat tails in the distribution of the returns, and one of the noteworthy aspects of the various GARCH-type models and many of the multifractal models is that the heteroskedasticity that they engender can help to account for the leptokurtosis in financial returns, since processes whose increments are drawn from a distribution whose volatility varies across the increments will exhibit excess kurtosis, relative to processes whose increments are drawn from a distribution exhibiting constant volatility.

Nonetheless, the degree of leptokurtosis observed in financial returns is typically much greater than what can be accounted for simply through the observed heteroskedasticity, so, as was noted above with regard to GARCH-type models, most heteroskedastic models of financial returns also incorporate a leptokurtic conditional distribution, most commonly the Student’s t distribution, to account for the portion of the observed leptokurtosis that cannot be accounted for by the conditional heteroskedasticity. Moreover, the degree of leptokurtosis that remains and must be explained by these leptokurtic conditional distributions is typically extremely high, as reflected in estimated degrees of freedom for fitted conditional Student’s t distributions that are typically below eight, and often as low as four or less, in which case the tails would be so fat that the kurtosis would be undefined. And, notably, the original specification of the EGARCH model [18], whose conditional variance equation can be viewed as a generalized form of the volatility equation developed by Zeeuw van der Laan for his Generalized SEMF model, also incorporated a conditional distribution, the generalized error distribution (GED), that allows for leptokurtosis within the fitted returns, and both of the data series to which Nelson originally fitted the EGARCH model had estimated tail parameters suggesting high degrees of conditional leptokurtosis among the returns, even after accounting for the effects of the conditional heteroskedasticity.

Despite these typical empirical findings, Filimonov and Sornette’s original specification of the SEMF model incorporates a simple Normal or Gaussian conditional (or residual) distribution (as did the original ARCH and GARCH model specifications), and Zeeuw van der Laan fits the SEMF model as it was originally specified,

with the assumption that the returns are conditionally Normally distributed. Zhong and Zhao, by contrast, take explicit account of the high degree of leptokurtosis among the financial returns they are attempting to model and, rather than assume conditional Normality in fitting the SEMF model to their data, they instead appear to follow a two-stage process in which they first fit a leptokurtic distribution, the generalized tempered stable distribution, to the actual returns themselves and then incorporate the estimated values for the distributional parameters into the likelihood function for the SEMF model, which is then maximized with respect to the parameters of the SEMF's conditional volatility equation (*i.e.*, h_0 , φ , and σ_0) in order to estimate the values of these SEMF model parameters. It is likely to be this difference between the two studies, the assumptions they make regarding the underlying distributions of the returns they are fitting, that plays the key role in driving the dramatically different results they obtain.

Thus, Zeeuw van der Laan implicitly assumes, in effect, that the self-excited multifractal behavior being fitted is the sole factor driving the extant leptokurtosis within the financial return series he examines. Consequently, the parameters for the fitted SEMF models must perform double duty, simultaneously explaining or accounting for both the excess kurtosis and the presumed underlying SEMF behavior, and, to the extent that at least a part of the excess kurtosis is being driven by separate factors other than the SEMF behavior, this can render it impossible for the fitted model to simultaneously account for both sets of characteristics adequately and can in turn lead to unstable parameter estimates, parameters that imply explosive volatility dynamics, or simply a poorly fitting model. On the other hand, though, 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. While Zeeuw van der Laan's estimated SEMF parameters are required to perform double duty, Zhong and Zhao's fitted parameters, by contrast, would be precluded from even fully performing the single duty for which they were designed—capturing the full effect of the self-excited multifractal behavior, together with the degree of leptokurtosis that would naturally be driven by that behavior, within the index returns Zhong and Zhao study.

In other words, we are 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 they obtain could clearly be driven by the role that leptokurtosis plays in their respective model specification and estimation processes, and the correct response to resolve the question of whether the SEMF model is reasonable or not is to incorporate a potentially leptokurtic conditional distribution, especially one such as the Student's *t* distribution, in which the Normal distribution is nested and can be obtained as a special case, into the SEMF model specification and to jointly and simultaneously estimate both the parameters of this leptokurtic conditional distribution along with the parameters of the SEMF volatility equation 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 clearer insight into the extent to which the leptokurtosis in financial returns is a consequence of the conditional heteroskedasticity within these returns (*i.e.*, what the SEMF conditional volatility equation accounts for) versus being an inherent characteristic of financial returns that is driven by something deeper.

6. Exploratory Results for SPY Returns

As a preliminary effort to explore this issue, two variations of an exponential kernel SEMF model, *i.e.*, two versions of the model:

$$R_t = \sigma_0 \xi_t e^{\frac{\omega_t}{\sigma_0}} \quad (7)$$

with:

$$\omega_t = \sum_{\tau=0}^{t-1} \left[h_0 e^{-\varphi(t-\tau-1)} \right] R_\tau, \quad (8)$$

one the original version incorporating a Normal conditional distribution, *i.e.*, $\xi_t \sim \text{NIID}(0, 1)$, and the other version incorporating a Student's *t* conditional distribution with an unknown degrees of freedom parameter ν —*i.e.*, $\xi_t \sim \text{St}(\nu) \forall t$ —were fit to a representative financial time series, the daily returns for the SPDR ETF (SPY) from 1 February 1993 through 8 June 2015 (a total of 5629 observations) that were obtained through Bloomberg. The Student's *t* distribution was chosen both because it is widely used in financial applications and because it incor-

porates the Normal distribution as a limiting case as the degrees of freedom for the distribution tend toward infinity.

Under the conditionally Normal SEMF specification, the log-likelihood function is given as follows:

$$\mathcal{L}(h_0, \varphi, \sigma_0) = \sum_{t=1}^T \ln f(R_t | \underline{R}_{t-1}; h_0, \varphi, \sigma_0) = -\frac{T}{2} \ln 2\pi - \frac{1}{2} \sum_{t=1}^T \ln \sigma_t^2 - \frac{1}{2} \sum_{t=1}^T \left(\frac{R_t}{\sigma_t} \right)^2, \quad (9)$$

where the conditional volatility σ_t is given by:

$$\sigma_t = \text{std.dev.}(R_t | \underline{R}_{t-1}) = \sigma_0 e^{\frac{\omega_t}{\sigma_0}} \quad (10)$$

and for which, for the exponential kernel case, ω_t is specified as shown above, as a linear function of the prior returns up through R_{t-1} :

$$\omega_t = \sum_{\tau=0}^{t-1} \left[h_0 e^{-\varphi(t-\tau-1)} \right] R_\tau = h_0 \sum_{\tau=0}^{t-1} \left[e^{-\varphi(t-\tau-1)} \right] R_\tau. \quad (11)$$

This model was fit using Matlab, under the initial condition for the estimation that R_0 for the conditional volatility recursion is presumed to equal zero. Thus, as a consequence:

$$\omega_1 = 0 \quad (12)$$

so that the conditional volatility for the first return observation, R_1 , is presumed to be equal to the unconditional volatility, σ_0 , while, for the following observations:

$$\omega_2 = \sum_{\tau=0}^1 h_{t-\tau-1} R_\tau = h_0 R_1 + h_1 R_0 = h_0 R_1, \quad (13)$$

and so forth; in general:

$$\omega_t = \sum_{\tau=0}^{t-1} h_{t-\tau-1} R_\tau = \sum_{\tau=1}^{t-1} h_{t-\tau-1} R_\tau. \quad (14)$$

(Note that this approach is somewhat different from the approach followed by Zeeuw van der Laan; his approach excludes the first 200 observations from being fitted in the likelihood itself and uses them solely for training the variance equation for the first return to which the model is fit, which is the 201st observation.)

The parameter estimates generated under this approach for the conditionally Normal SEMF model specification are as shown in **Table 1**. As can be seen in this set of results, the parameter estimates obtained for this model specification for the daily SPY returns are similar to those found by Zeeuw van der Laan for the data sets he examined. Specifically, the estimate of h_0 , the non-stationarity parameter, is very large, approaching a value of one, suggesting a highly non-stationary process, while the estimate for φ , the memory dissipation parameter, is actually negative, suggesting an explosive volatility process, with the impact of memory actually increasing over time rather than slowly fading. Notably, however, each of these parameters is extremely unstable, with each exhibiting standard errors that dwarf the magnitudes of the parameters themselves, leading to statistically insignificant parameter values with p-values close to 1.00.

A reflection of the poor fit of this model is provided by a comparison of the summary statistics for the original series of returns versus those for the standardized residuals from the fitted model, as shown in **Table 2**. As is typical of financial time series, the degree of kurtosis among the SPY returns is extremely high, 13.94 (note: for a Normal series, the kurtosis equals three), which suggests a large number of extreme magnitude returns, of either

Table 1. Normal SEMF model parameter estimates for SPY returns.

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

Table 2. Summary statistics for SPY returns versus normal SEMF model residuals.

Summary statistic	Original series	Standardized residuals
Mean	0.0004	0.0002
Standard deviation	0.0119	0.0069
Skewness	0.1144	0.0220
Kurtosis	13.9358	12.4217

sign, among the original series. These extreme returns or “outliers” would need to be accounted for by the fitted model. Unfortunately, the SEMF specification does not appear to be able to succeed very well at this task, and the degree of kurtosis for the standardized residuals is still extremely high, 12.42, reduced only slightly compared to the original returns.

In contrast, though, the standard deviation for the standardized residuals is extremely low, only 0.0069, compared to the magnitude of one that the standard deviation of standardized residuals should have, indicating that the estimated conditional volatility is being dramatically inflated or over-stated. The extreme degree of this estimated volatility inflation can be estimated either by the inverse of the standard deviation of the standardized residuals (145.26x), which would reflect the impact of the estimated conditional volatility, given past returns, for each observation, or by the ratio of the estimated value of the unconditional volatility parameter, σ_0 , to the standard deviation of the original returns (146.16x); either method suggests that the scale or volatility estimated by the model is inflated by an extremely high factor of nearly 150-to-one.

This suggests that, since the model cannot directly account for the extreme degree of kurtosis within the original returns, the attempt nonetheless to fit the large number of extreme returns within the data leads 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 cannot do an adequate job of fitting the tails of the distribution (the key driver of what kurtosis, the 4th central moment, measures), it overcompensates by instead adjusting the estimated scale of the distribution (which is what standard deviation, or the 2nd central moment, measures). Moreover, given the extremely high magnitude of the unconditional standard deviation, σ_0 , that is estimated for the model, the magnitude of any conditional heteroskedasticity effects that are being 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 would, in effect, in comparison to the estimated value of σ_0 , be merely a “tempest in a teapot”.

7. Fitting the Conditionally Student’s t Alternative to SPY Returns

Given the poor fit of the Normal SEMF model specification and the fact that, given the high degree of kurtosis they exhibit, the residuals from this model are clearly not Normally distributed (as the model specification assumes them to be), a natural alternative is to re-specify the SEMF model to incorporate a Student’s t distribution as the conditional or residual distribution. This would bring an additional parameter into the model, the degrees of freedom parameter ν , to describe the tail behavior of the returns and allow for leptokurtosis among the residuals.

Note that the Student’s t distribution can be derived as the marginal distribution that results from the combination of a Normal distribution conditional on the value for the variance, $N(0, \sigma^2)$, with the value of the variance σ^2 itself also being a random variable whose value follows an inverse gamma distribution. Conceptually, then, a Student’s t SEMF model could be interpreted as incorporating two sources of variability—the returns vary Normally about their mean of zero on a scale determined by the variance, but this variance itself also varies, in accordance with an inverse gamma distribution, around its own respective mean, which is determined by the SEMF conditional volatility equation, with the degree of this variability being a function of ν —the smaller is ν , the greater is the variability of the variance, while, as ν increases toward infinity, the variance tends toward a constant and the marginal distribution tends toward the Normal. These two sources of variability, (1) the variability of the variance itself around its mean of $\sigma_0^2 e^{-\frac{2}{\sigma_0} \sum_{r=0}^{t-1} [h_0 e^{-\varphi(t-r-1)}] R_r}$, and then, given the specific value for the variance σ_t^2 , (2) the variability of the return itself around its own mean of zero, interact with each other to al-

low for the extreme-magnitude returns and fat tails observed in practice.

An alternative interpretation for the value of incorporating a leptokurtic conditional distribution into the model in place of the Normal distribution is that the conditionally Normal SEMF specification allows for multifractal dynamics at time scales of daily (assuming daily returns are being fit) and wider but implies that, on an intra-daily scale, they follow much simpler dynamics that can be approximated by a Normal random walk. Incorporating a Student's t or other leptokurtic distribution as the model's conditional distribution allows the model to reflect the aggregate effects of more complicated, and more realistic, dynamics at the intra-daily level.

The log-likelihood function for the Student's t SEMF model incorporates four parameters—three of which, h_0 , φ , and σ_0 , determine the evolution of the conditional volatility σ_t , and ν , which describes the behavior of the tails of the distribution—and is given as follows:

$$\begin{aligned} \mathcal{L}(\theta) &= L(h_0, \varphi, \sigma_0, \nu) = \sum_{t=1}^T \ln f(R_t | \mathbf{R}_{t-1}; h_0, \varphi, \sigma_0, \nu) \\ &= T \ln \left[\frac{\Gamma\left(\frac{\nu+1}{2}\right)}{\Gamma\left(\frac{\nu}{2}\right) \sqrt{\pi(\nu-2)}} \right] - \frac{1}{2} \sum_{t=1}^T \ln \sigma_t^2 - \left(\frac{\nu+1}{2}\right) \sum_{t=1}^T \ln \left[1 + \frac{1}{(\nu-2)} \left(\frac{R_t}{\sigma_t}\right)^2 \right]. \end{aligned} \quad (15)$$

where σ_t and ω_t are specified the same as given above for the conditionally Normal SEMF model.

The results of fitting the Student's t SEMF model are provided in **Table 3**, and they paint a picture that is very different from that of the previous specification. Most notably, in contrast to the Normal SEMF specification case, all of the parameter estimates for the Student's t specification exhibit extremely low standard errors and, consequently, extremely high significance levels (*i.e.*, extremely low p-values), suggesting a model that is much better fit to the data.

Among the estimated parameter values, the relatively low value for the degrees of freedom parameter ν implies that a substantial portion of the leptokurtosis with the SPY returns is properly accounted for by the conditional or residual distribution rather than being a consequence of the heteroskedasticity across the returns; consequently, rather than fitting the tails of the data, the conditional volatility parameters are freed to better fit the scale of the returns, at which they appear to succeed much better, with regard to both the unconditional and the conditional volatility. In terms of the former, the estimated value of σ_0 now has a much more reasonable value, close to the estimated standard deviation of the original series; consequently, in contrast to its highly inflated value from the Normal case, the value of σ_0 under the Student's specification is low enough so that the movements of the conditional volatility around this estimate of the unconditional volatility could be of sufficient magnitude to be meaningful. And, in terms of the estimated values of the parameters driving the variations in the conditional volatility equation, the results found here are more similar to the results of Zhong and Zhao—namely, h_0 is relatively small, indicating a relatively low degree of non-stationary, and φ is both small and positive, indicating a volatility process with a memory that is long but that does dissipate slowly over time, rather than increasing over time and becoming explosive as would be suggested by the Normal case parameters.

The much better fit this model specification provides is also reflected in the summary statistics for its standardized residuals, shown in **Table 4**. In marked contrast to the Normal specification case, the standardized residuals from the Student's t specification exhibit a sample standard deviation that is close to one, suggesting it is doing a much better job of accounting for and predicting the volatility within the returns. Moreover, the variation in the conditional volatility incorporated into this model now appears to be able to account for an appreciable degree of the leptokurtosis within the original returns, so that the degree of kurtosis among the standardized

Table 3. Student's t SEMF model parameter estimates for SPY returns.

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

Table 4. Summary statistics for SPY returns versus alternative residual series.

Summary statistic	Original series	Standardized residuals	
		Normal SEMF	Student's t SEMF
Mean	0.0004	0.0002	0.0344
Standard deviation	0.0119	0.0069	1.1270
Skewness	0.1144	0.0220	-0.3931
Kurtosis	13.9358	12.4217	6.4696

residuals is less than half that of the original returns. Finally, and importantly for the sake of statistical adequacy but in direct contrast to what was found in the Normal case, the degree of kurtosis that still remains among the standardized Student's t model residuals is roughly consistent with the degree of kurtosis that would be expected under the assumed residual distribution, a Student's t distribution with relatively low degrees of freedom (For a discussion of statistical adequacy and the issues and problems associated with misspecified statistical models, see, e.g., [20] and [21]).

These results clearly suggest that the SEMF model can be a reasonable model for financial returns and that the findings of non-stationary and explosive volatility dynamics obtained by Zeeuw van der Laan likely reflect, not a poorly specified equation for the conditional volatility, but simply the need to incorporate a conditional distribution that allows for a degree of leptokurtosis beyond what the evolution of the conditional volatility, by itself, can account for. To explore the generalizability of this conclusion, this paper will next examine the fit of various specifications of the SEMF model to a variety of daily stock and index returns, including, in addition to the SPY returns, all of the daily return series examined by Zeeuw van der Laan and Zhong and Zhao in their respective studies.

8. Analysis of the Initial Estimation Results for the Alternative Sets of Daily Returns

Among the series of returns that Zeeuw van der Laan examines are daily stock returns for IBM (including both a longer series, starting in 1962, and a shorter series, starting in 2002), Royal Dutch Shell (which has two classes of stock outstanding, RDS.A and RDS.B, which differ in terms of whether their dividends are subject to Dutch withholding tax; since Zeeuw van der Laan does not specify which class of shares he examines, both classes are examined here), General Electric (GE), and Coca-Cola (KO), and daily index returns for the Dow-Jones Industrial Average (DJIA) and the S&P 500 (SP500). Zhong and Zhao examine daily index returns for the Hang-Seng Index (HSI) for the Hong Kong stock market and the Shanghai Stock Exchange Composite Index (SSE). (Zeeuw van der Laan also examines half-hourly return series for IBM, GE, the DJIA, and three foreign exchange series, the U.S. Dollar/Euro, the Japanese Yen/U.S. Dollar, and the Swiss Franc/Euro exchange rates. Due to a lack of data availability, along with the possibility that market microstructure effects could introduce additional complications into the dynamics of such higher frequency returns, they are not included in the present study.)

For the current study, daily returns for the above-listed stocks and for the S&P 500 were obtained from the University of Chicago's CRSP database (accessed through Wharton's WRDS), and daily returns for the other indices were obtained through Bloomberg, for approximately the same time frames as the two previous studies. Summary information and summary statistics for the daily return series examined are provided in **Table 5**. As can be seen from this table, the lengths of the return series range from slightly fewer than 1500 observations for the shortest series, for the two classes of Royal Dutch Shell stock, up to more than 15,000 observations for the two U.S. index series. With regard to the summary statistics, the sample means are all positive but close to zero (with magnitudes that, on average, are about one-thirtieth of a standard deviation away from zero). In terms of skewness, the various series appear to show a variety of left- and right-skewness, though with most exhibiting relatively small degrees of skewness, either way. Two notable exceptions are the Hang Seng and Shanghai Stock Exchange index returns, which each exhibit much greater degrees of skewness, but in opposite directions—positive for the SSE returns and negative for the HSI returns.

The most notable summary statistic, though, is the fourth one—the sample kurtosis. Similar to what was found and extensively discussed with regard to the SPY returns, every single one of the series of returns is extremely

Table 5. Summary information and statistics of return series examined.

Series	Beg. date	End. date	# Obs.	Mean	Std. Dev.	Skew.	Kurt.
SPY	2/1/93	6/8/15	5629	0.00042	0.0119	0.114	13.94
GE	1/2/62	5/20/11	12,433	0.00052	0.0164	0.186	11.52
IBM (long)	1/2/62	5/20/11	12,431	0.00047	0.0162	0.075	12.93
IBM (short)	2/6/02	6/13/11	2355	0.00036	0.0159	0.335	9.78
KO	1/2/62	5/23/11	12,434	0.00061	0.0152	-0.005	15.36
RDS.A	7/26/05	6/10/11	1481	0.00047	0.0201	0.531	13.83
RDS.B	7/21/05	6/10/11	1484	0.00046	0.0205	0.474	13.57
DJIA	3/1/50	6/13/11	15,498	0.00031	0.0096	-0.782	31.11
SP500	1/3/50	5/25/11	15,536	0.00046	0.0096	-0.603	23.33
HSI	1/2/87	4/1/11	6025	0.00053	0.0176	-1.277	36.54
SSE	1/4/93	4/1/11	4449	0.00056	0.0230	1.742	28.63

leptokurtic, with degrees of kurtosis for each series well above the value of three that a Normal series would exhibit. The series that is the least leptokurtic is the shorter IBM series, which is the only series for which the degree of kurtosis is (only slightly) less than ten, while, interestingly, the daily returns for DJIA, SP500, HSI, and SSE, the four index series (which are not actually traded securities themselves but rather are mathematical constructs that represent weighted averages of the returns of traded securities and hence should reflect the returns of well-diversified portfolios) exhibit the most extreme degrees of leptokurtosis, with sample kurtosis measures of greater than twenty for each, substantially higher than the degrees of kurtosis exhibited by the individual securities' returns.

The results of fitting the conditionally Normal specification of the SEMF model to these series of returns are provided in [Table 6](#) (for comparison, the sample kurtosis from the original sets of returns is also provided, in the final column). As could be expected, given the previous findings for the SPY returns and the extreme degrees of kurtosis similarly exhibited by rest of the sets of returns, this specification of the SEMF model does not seem to provide a good or reasonable description of the dynamics for these sets of returns.

With regard to the model parameters, for four of the series—SPY, the two RDS series, and SSE—the nonstationarity parameter h_0 and the unconditional volatility parameter σ_0 are both estimated to be very large, implying that these series are highly volatile and nonstationary, and the memory strength parameter φ is estimated to be negative, implying memory that strengthens rather than diminishes over time, leading to potentially explosive volatility dynamics. And, as was discussed for the SPY series, the conditional volatility is dramatically over-estimated for each of these four series, so that, for each of these series, the standardized residuals from this model have sample standard deviations of approximately 0.03 or less, rather than their theoretical values of close to one. But, consistent with these highly inflated estimated volatility levels, the estimated model parameters also all have extremely unstable values, which are reflected in indicated p-values that are approximately equal to one for all three of the model parameters for each of these four series.

For all but one of the remaining series, by contrast—specifically, for GE, the longer IBM series, KO, DJIA, SP500, and HSI—the memory strength parameter φ is estimated to be so large (with values greater than 20) that, rather than building up and becoming explosive over time, any memory within these series of returns would not only fade over time but would actually dissipate so quickly that there would effectively be no memory past the first lag (a result that would be inconsistent with the extremely long memory that is typically found within the volatility of financial return series). For the three index series among these six, however, there is a much more dramatic issue of potential concern—the volatility parameter σ_0 is estimated to actually have a negative value, which, especially since it occurs for multiple sets of returns (all of which are index series), suggests there is something clearly going on within these returns to which the optimization algorithm is having difficulty adequately adjusting the model parameters.

Table 6. Results of fitting normal SEMF model specification.

Series	Estimated parameter values (<i>p</i> -values in parentheses)			Standardized residuals			Orig. returns
	h_0	Φ	σ_0	St. Dev.	Skew.	Kurt.	Kurt.
SPY	0.7405 (0.998)	-0.0005 (0.999)	1.7432 (0.999)	0.0069	0.022	12.42	13.94
GE	0.0266 (0.000)	25.2173 (0.000)	0.0164 (0.000)	0.9995	0.235	11.60	11.52
IBM (long)	0.0312 (0.000)	26.6106 (0.000)	0.0162 (0.000)	0.9996	0.109	12.24	12.93
IBM (short)	0.0416 (0.000)	0.0076 (0.000)	0.0157 (0.000)	0.9902	0.211	7.43	9.78
KO	0.0257 (0.000)	20.1040 (0.000)	0.0152 (0.000)	0.9992	0.066	12.65	15.36
RDS.A	0.7632 (1.000)	-0.0030 (1.000)	1.7495 (1.000)	0.0114	0.375	12.26	13.83
RDS.B	0.8350 (0.998)	-0.0020 (0.995)	0.6501 (0.928)	0.0310	0.151	10.77	13.57
DJIA	-0.0503 (0.000)	63.8986 (0.000)	-0.0095 (0.000)	0.9994	0.315	20.34	31.11
SP500	-0.0546 (0.000)	44.1210 (0.000)	-0.0095 (0.000)	0.9987	0.310	16.10	23.33
HSI	-0.0935 (0.000)	35.6430 (0.000)	-0.0169 (0.000)	0.9995	0.414	15.65	36.54
SSE	0.7002 (0.978)	-0.0006 (0.958)	1.7537 (0.013)	0.0134	1.518	25.68	28.63

A key goal of this paper was to let the data provide information to us about how well the specified model suits the data in which the modeler is interested. As such, the model parameters were purposely left unconstrained in the optimization algorithm, so that their values could adjust to where they need to be to best reflect the true characteristics and dynamics of the underlying data rather than where the model suggests the parameters “should” be. That being said, the interpretation of a negative value for σ_0 is complicated by the fact that, wherever σ_0 appears in the likelihood function, it is either squared (so a negative value for σ_0 would not shrink the density function down to a singularity or a complex-valued function) or it is being divided into ω_t , whose sign is determined by the sign of h_0 , which is also negative in these cases, so the negative signs for these two variables would cancel out. Nonetheless, while the parameter values were left unconstrained within the optimization algorithm, the starting values input for these parameters were all positive and set at “reasonable” values that would imply non-explosive multifractal behavior. Specifically, the starting values used were $h_{0,\text{start}} = 0.8$, $\varphi_{\text{start}} = 0.08$, and $\sigma_{0,\text{start}}$ was set equal to the sample standard deviation for the original series. Thus, the fact that optimization process still pushes h_0 and σ_0 into negative territory to provide the best fit for the model suggests that, in its current specification, it does not fit these sets of returns very well. Or, put another way, the conditionally Normal SEMF model does not provide a good simultaneous description for both the dynamics and statistical distributional characteristics of these sets of returns.

But, even without examining the specific values of the estimated parameters more closely, this lack of a good statistical fit is also reflected in the extreme degrees of kurtosis, ranging from a low of just under 7.5 all the way up to a high of greater than 25, that are found among the standardized residuals for every single one of the series. Such extreme kurtosis is inconsistent with a Normal conditional distribution, and, as was discussed for the SPY returns in the earlier part of this paper, its presence suggests that the estimated SEMF model parameters are thus required to do double-duty of trying to capture both the volatility dynamics within the returns and the excess kurtosis among the innovations driving the returns, and the estimated models thus end up performing poorly on at least one, and in most cases both, of these tasks. This suggests that, as had previously been found for the SPY returns, an alternative model specification, such as the conditionally Student's t SEMF model specification, that explicitly allows for leptokurtosis within the innovations to the returns, may provide both a better fit to the returns and more reasonable estimates for the parameters describing the volatility dynamics among these returns.

9. Impact on the Estimation Results of Incorporating a Leptokurtic Conditional Distribution

The results obtained from fitting the Student's t SEMF specification to the daily return series are shown in **Table 7**. Notably, the degrees of freedom parameter ν is relatively low for all eleven cases, ranging from 8.63 at its highest, for the RDS.A returns, all the way down to a low of only 2.44 for the SSE returns, and this estimated

Table 7. Results of fitting student's t SEMF model specification.

Series	Estimated parameter values (p-values in parentheses)				Standardized residuals			Orig. returns
	h_0	Φ	σ_0	ν	St. Dev.	Skew.	Kurt.	Kurt.
SPY	0.0716 (0.000)	0.0191 (0.000)	0.0120 (0.000)	4.5279 (0.000)	1.1270	-0.393	6.47	13.94
GE	0.0247 (0.000)	0.0103 (0.000)	0.0169 (0.000)	4.3504 (0.000)	1.0635	0.072	8.78	11.52
IBM (long)	-0.0227 (0.000)	0.0237 (0.000)	-0.0164 (0.000)	3.9840 (0.000)	1.0644	-0.093	12.03	12.93
IBM (short)	0.0422 (0.000)	0.0079 (0.000)	0.0159 (0.000)	4.4697 (0.000)	0.9741	0.211	7.44	9.78
KO	0.0231 (0.000)	0.0194 (0.000)	0.0158 (0.000)	4.0320 (0.000)	1.0768	-0.060	10.69	15.36
RDS.A	-0.0600 (0.000)	0.0075 (0.000)	-0.0195 (0.000)	8.6364 (0.000)	1.0009	0.176	7.81	13.83
RDS.B	0.0544 (0.000)	0.0069 (0.000)	0.0199 (0.000)	6.8707 (0.000)	1.0111	-0.069	9.71	13.57
DJIA	0.0410 (0.000)	0.0241 (0.000)	0.0095 (0.000)	4.3174 (0.000)	1.1804	-0.665	15.14	31.11
SP500	-0.0002 (0.000)	-0.0001 (0.000)	0.0072 (0.000)	3.7805 (0.000)	1.3367	-0.601	23.46	23.33
HSI	-0.0573 (0.000)	0.0445 (0.000)	-0.0172 (0.000)	3.4392 (0.000)	1.0906	0.670	14.40	36.54
SSE	-0.0213 (0.000)	12.5153 (0.000)	0.0283 (0.000)	2.4436 (0.000)	0.8126	1.765	29.14	28.63

degrees of freedom parameter ν is less than five for all series save for the two sets of RDS returns. All of these estimated values imply a high degree of conditional leptokurtosis; the lower is ν , the more extreme the degree of leptokurtosis, and if $\nu < 3$, as is estimated to be the case for the SSE returns, the tails of the distribution taper off so slowly that the theoretical kurtosis would be infinite. (Note: for the Student's t distribution, the theoretical value of the kurtosis is $3 + 6/(\nu - 4)$, for $\nu > 4$, and is ∞ for $2 < \nu \leq 4$; for $\nu < 6$, as is the case for all but the two RDS series, the kurtosis itself is finite, but the variance of the estimate for the kurtosis would be undefined, while, for ν greater than but close to six, this variance would be finite but extremely large; thus, in finite samples, the value of the sample kurtosis in such cases would be large but variable, which is consistent with the results found here for the various daily return series.) Thus, the relatively low estimated values for ν found for all of these sets of returns mean that, in each case, a significant amount of the residual leptokurtosis can be accounted for by the leptokurtosis of the conditional distribution, leaving fewer of the statistical characteristics and dynamics within the data that the SEMF model parameters have to account for. The fact that the sample standard deviations of the standardized residuals are now relatively close to one for all eleven of the series is consistent with this conclusion, suggesting that the estimated SEMF model parameters are now able to do a better job of capturing the volatility dynamics within the series of returns.

Nonetheless, despite these indications of improvement, and notwithstanding the positive parameter results that were found for the SPY returns under the Student's t SEMF model specification, the same is not necessarily true for all of the other sets of returns under examination here. For approximately half of the sets of returns, including SPY, GE, the shorter IBM series, RDS.B, and DJIA, the Student's t SEMF model specification yields fairly reasonable parameter estimates. The remaining five series, however, all entail parameter estimates that are problematic in some way, including, in all five cases, negative values for the nonstationarity parameter h_0 . Three of these series, the longer IBM series, RDS.A, and HSI, also exhibit negative estimates for σ_0 , which introduces the complications in interpretation noted above but which also cancels out the effect of a negative estimated value for h_0 .

For SP500 and SSE, on the other hand, σ_0 is estimated to be positive, but the estimated values for ϕ are problematic— ϕ is estimated to be negative for SP500, suggesting potentially explosive volatility dynamics (although the estimated value is at least very small and close to zero, reducing the degree of potential explosiveness), while ϕ is estimated to be extremely large and positive for SSE, implying rapid fading of memory beyond the first lag. For both of these series, moreover, since they are not offset by negative values for σ_0 , the negative estimated values for h_0 are also problematic. Not only does h_0 control the level of the multifractality and nonstationarity within the process, but it also drives the so-called “leverage effect” within the returns. This effect is arguably a reflection of the degree of risk aversion among investors and refers to the fact that the volatility process within financial returns typically appears to be skewed, so that a given magnitude of below average return will lead to a greater increase in the subsequent period's volatility than will the same magnitude of above-average return. But, the fact that h_0 is estimated to be negative for SP500 and SSE suggests that this skewness works in the “wrong” direction, with above-average returns leading to greater subsequent volatility and below-average returns leading to a decrease in subsequent volatility for these series.

This finding points to the possibility of another possible source of distributional misspecification that has not yet been addressed but is much more basic—the mean or “level” of the distribution of returns. The mean is the central value around which the variable described by the distribution varies, and if it is misspecified then this will bias the estimates for the other parameters and central moments of the distribution. E.g., volatility reflects the typical distance of possible values for the return away from the mean value for the return, so, if this mean value for the return is misspecified, then the volatility parameters involved with measuring the distance away from it will be distorted; similarly, skewness reflects the relative distances away from the mean of greater-than mean returns versus less-than-mean returns, so, again, skewness-related parameters will be distorted by a mis-identification of the location of the mean.

The SEMF model as originally developed and as applied thus far presumes that the series of returns has a mean of zero. Concomitantly, under this original specification a positive return would, by presumption, also be a greater-than-mean return and, conversely, a negative return would be a less-than-mean one. However, while average daily stock and index returns tend to be very small and close to zero (especially in comparison to the typical degree of day-to-day volatility seen in financial returns), they nonetheless do tend to be positive, which is what allows them to gradually accumulate over time to generate the long-run returns that investors expect of them (And, consequently, a return can be positive and nonetheless still be a below-mean return).

These characteristics are reflected in the eleven sets of stock and index returns examined in this paper—their sample means, as noted above, are very small and close to zero in comparison to their respective sample standard deviations; nonetheless, adjusting for (the square root of) the large number of observations in each of these series entails estimate standard errors of the mean compared to which the sample means appear to be significantly different (in a positive direction) from zero for seven of the series—SPY, GE, the longer IBM series, KO, DJIA, SP500, and HSI; they are positive but not large enough to appear significant for the shorter IBM series and the two RDS series; finally, the returns for SSE appear to be at the edge of being significant, with an implied p-value of 0.105. Thus, the fact that the majority of these sets of returns have means that appear to be positive and significantly different from zero necessitates that efforts to obtain model parameters which reflect the true dynamics and statistical characteristics of these returns must allow for the effects of such non-zero means.

10. Impact on the Estimation Results of Pre-Centering the Returns

When the main focus of interest is on characteristics such as multifractality, volatility dynamics, or other nonlinear behavior within a given financial (or other) series, researchers will sometimes pre-process their data before applying the model of interest to it. Such pre-processing can range from, at one end, extensive efforts to pre-whiten the data, including centering it by subtracting off the sample mean, standardizing the result by dividing through by the sample standard deviation, and then filtering the data to remove any linear dependencies or autocorrelation within the data; to simply centering and standardizing the data; to, at the most basic, simply centering the data. This latter would typically be the choice if one were interested in using the model to be fitted to explore the behavior of the volatility of the series.

Since modeling the volatility dynamics is likewise a key goal of the SEMF model, the latter alternative of simply centering the returns by subtracting off their respective sample means is the level of pre-processing that is initially adopted here. The results from fitting the Student's t SEMF model specification to the respective sets of pre-centered returns are shown in [Table 8](#).

Table 8. Results of fitting student's t SEMF model to pre-centered returns.

Series	Estimated parameter values (p-values in parentheses)				Standardized residuals			Orig. returns
	h_0	Φ	σ_0	ν	St. Dev.	Skew.	Kurt.	Kurt.
SPY	0.0610 (0.000)	0.0178 (0.000)	0.0105 (0.000)	4.4261 (0.000)	1.026	-0.360	6.46	13.94
GE	-0.0227 (0.000)	0.0101 (0.000)	-0.0157 (0.000)	4.3943 (0.000)	1.021	-0.073	8.29	11.52
IBM (long)	-0.0220 (0.000)	0.0233 (0.000)	-0.0160 (0.000)	4.0101 (0.000)	1.002	-0.086	11.64	12.93
IBM (short)	-0.0387 (0.000)	0.0082 (0.000)	-0.0142 (0.000)	4.4118 (0.000)	1.031	-0.215	7.62	9.78
KO	-0.0217 (0.000)	0.0188 (0.000)	-0.0150 (0.000)	4.0661 (0.000)	1.001	0.058	10.49	15.36
RDS.A	0.0470 (0.000)	0.0054 (0.000)	0.0160 (0.000)	9.3721 (0.000)	1.133	-0.151	7.83	13.83
RDS.B	0.0419 (0.000)	0.0043 (0.000)	0.0164 (0.000)	7.3821 (0.000)	1.155	-0.045	9.66	13.57
DJIA	0.0382 (0.000)	0.0231 (0.000)	0.0090 (0.000)	4.2857 (0.000)	1.011	-0.656	14.57	31.11
SP500	-0.0386 (0.000)	0.0188 (0.000)	-0.0089 (0.000)	3.9905 (0.000)	1.014	0.600	11.97	23.33
HSI	0.0553 (0.000)	0.0448 (0.000)	0.0165 (0.000)	3.4307 (0.000)	0.991	-0.876	17.59	36.54
SSE	-0.0236 (0.000)	9.3089 (0.000)	0.0285 (0.000)	2.4344 (0.000)	0.807	1.768	29.20	28.63

Unfortunately, the results here show little improvement over the results for the un-centered returns—reasonable parameter estimates are obtained for only five of the series (SPY, RDS.A, RDS.B, DJIA, and HSI), while negative values are obtained for both h_0 and σ_0 for another five of the series (GE, both the IBM series, KO, and SP 500), and a negative value for h_0 by itself for the remaining series (SSE). But, rather than focus on these parameter estimates, it is important to discuss a more fundamental problem involving the approach used to pre-process the returns in the first place.

Specifically, the process of “centering” the data is typically conducted using the sample mean, simply subtracting the value of the sample mean for the series from each observation of the series. But, the sample mean is an ordinary least squares (OLS) estimator, and, in the presence of heterogeneous variance such as is found in financial returns, generalized least squares (GLS) or weighted least squares (WLS) estimators are generally more efficient, assuming one has reasonable estimates of the volatility for each observation, and such estimators typically lead to different estimates from OLS estimators. Similarly, the sample mean is also well known for its lack of robustness as an estimator for the level of a series and for its extreme sensitivity to the many outliers that are typically found in highly leptokurtic series. Consequently, for the heavy-tailed Laplace distribution, for example, the median of the data rather than the sample mean is the MLE for the level of the process, since the sample mean is too sensitive to the impact of the many outliers observed in series that follow such a distribution. Financial returns, unfortunately, present the modeler with both of these problems, heterogeneous variance and a leptokurtic conditional distribution, so that the sample mean is likely to be a very poor estimator for the true mean of such series.

11. Impact on the Estimation Results of Centering the Returns within the Model

A more appropriate alternative to the simple use of the sample mean to center the returns is to instead incorporate the mean μ as a separate parameter that is incorporated into and subtracted from the returns within the SEMF model specification itself, with its value then being estimated along with the rest of the model parameters via MLE. This would lead to an estimator that implicitly controls for the effects of both heterogeneous variance across the data points (through the SEMF’s volatility evolution equation) and the leptokurtosis among them (assuming the conditionally Student’s t specification) and is thus likely to be a much more efficient estimator for the true “level” of the returns than the sample mean.

Incorporating a mean component into the Student’s t SEMF model specification leads to a model with five parameters— h_0 , φ , σ_0 , ν , and μ —and with the following likelihood function:

$$\begin{aligned} \mathcal{L}(\theta) &= L(h_0, \varphi, \sigma_0, \nu, \mu) = \sum_{t=1}^T \ln f(R_t | x_t, \Phi_{t-1}; h_0, \varphi, \sigma_0, \nu, \mu) \\ &= T \ln \left[\frac{\Gamma\left(\frac{\nu+1}{2}\right)}{\Gamma\left(\frac{\nu}{2}\right) \sqrt{\pi(\nu-2)}} \right] - \frac{1}{2} \sum_{t=1}^T \ln \sigma_t^2 - \left(\frac{\nu+1}{2}\right) \sum_{t=1}^T \ln \left[1 + \frac{1}{(\nu-2)} \left(\frac{R_t - \mu}{\sigma_t}\right)^2 \right]. \end{aligned} \quad (16)$$

where σ_t , in turn, can be written as:

$$\sigma_t = \sigma_0 e^{\frac{\omega_t}{\sigma_0}} = \sigma_0 \exp\left(-\frac{h_0}{\sigma_0} \sum_{\tau=0}^{t-1} \left[e^{-\varphi(t-\tau-1)} \right] (R_\tau - \mu)\right). \quad (17)$$

The results of fitting this model specification are shown in **Table 9**. As can be seen in this table, once the SEMF model specification has been generalized to incorporate such widely recognized features of stock and index returns as (1) a non-zero mean and (2) a high degree of conditional leptokurtosis, in addition to (3) the effects of volatility dynamics on the returns, this more comprehensive SEMF model appears to be much better “behaved” than what was found for the less comprehensive model specifications, and, notably, for nearly all of the series (with only one notable exception, to be described shortly), all of the expanded set of parameters for this model are estimated to have reasonable values that would describe realistic behavior that is consistent with the “stylized features” of financial returns.

Specifically, the estimates for the unconditional volatility parameter σ_0 are positive for all of the return series, the estimated values for the memory strength parameter φ for all the return series are small but positive, indicating

Table 9. Results of fitting centered student's t SEMF model specification.

Series	Estimated Parameter Values (p-values in parentheses)					Standardized residuals			Orig. returns
	h_0	Φ	σ_0	ν	μ	St. dev.	Skew.	Kurt.	Kurt.
SPY	0.0542 (0.000)	0.0169 (0.000)	0.0096 (0.000)	4.3797 (0.000)	0.0007 (0.000)	0.977	-0.341	6.52	13.94
GE	0.0235 (0.000)	0.0102 (0.000)	0.0162 (0.000)	4.3741 (0.000)	0.0003 (0.000)	1.038	0.072	8.49	11.52
IBM (long)	0.0224 (0.000)	0.0235 (0.000)	0.0162 (0.000)	3.9940 (0.000)	0.0002 (0.000)	1.036	0.090	11.86	12.93
IBM (short)	0.0438 (0.000)	0.0078 (0.000)	0.0167 (0.000)	4.4916 (0.000)	-0.0002 (0.000)	0.950	0.210	7.38	9.78
KO	0.0223 (0.000)	0.0191 (0.000)	0.0154 (0.000)	4.0468 (0.000)	0.0003 (0.000)	1.035	-0.059	10.58	15.36
RDS.A	0.0340 (0.000)	0.0047 (0.000)	0.0121 (0.000)	10.1373 (0.000)	0.0010 (0.000)	1.391	-0.126	7.84	13.83
RDS.B	0.0293 (0.000)	0.0039 (0.000)	0.0119 (0.000)	7.9638 (0.000)	0.0010 (0.000)	1.476	-0.027	9.60	13.57
DJIA	0.0372 (0.000)	0.0228 (0.000)	0.0088 (0.000)	4.2793 (0.000)	0.0004 (0.000)	0.958	-0.653	14.36	31.11
SP500	0.0367 (0.000)	0.0183 (0.000)	0.0086 (0.000)	3.9756 (0.000)	0.0006 (0.000)	0.935	-0.597	11.74	23.33
HSI	0.0545 (0.000)	0.0449 (0.000)	0.0163 (0.000)	3.4328 (0.000)	0.0007 (0.000)	0.953	-0.976	19.26	36.54
SSE	-0.0171 (0.000)	0.0267 (0.000)	0.0276 (0.000)	2.4724 (0.000)	0.0004 (0.000)	0.827	2.116	35.37	28.63

a long but gradually fading memory within the volatility of each series, and, for all but the SSE returns, the estimated values for the nonstationarity parameter h_0 are, similarly, positive but relatively small, indicating both a slight degree of multifractality within these series of returns as well as a negative volatility skewness effect, consistent with the leverage effect found widely in financial returns.

With regard to the standardized residuals from these fitted Student's t SEMF models, each set of residuals exhibits a high degree of leptokurtosis that is consistent with the relatively low estimated degrees of freedom parameter ν that is found in each case (varying in a very low range from 2.47 to 4.49 for most of the series, and a bit higher but still leptokurtic 10.14 and 7.96 for RDS.A and RDS.B, respectively), and, for all but the SSE residuals, the degree of kurtosis is substantively less than that for the original returns, indicating that the volatility dynamics described by the fitted model is successfully able to account for a sizable portion of the leptokurtosis within the original returns. Similarly, and in contrast to the results from the fitted Normal SEMF model specification, the sample standard deviations for the standardized residuals are all relatively close to one, although, among these, the estimated volatility equation appears to slightly underestimate the volatility for the two RDS series, since the sample standard deviations for their respective standardized residuals are somewhat greater than one (just under 1.4 and 1.5, respectively).

The lone series whose results appear to be anomalous is the SSE series, for which the nonstationarity parameter h_0 is negative rather than positive (and not offset by a negative estimate for σ_0), suggesting a positive rather than a negative volatility skewness effect, which is the opposite of the direction of the traditional leverage effect

and could be taken to imply risk-seeking rather than risk-averse behavior among investors in this market. Ironically, this is one of the series for which Zhong and Zhao had obtained more reasonable parameter values. Nonetheless, as was seen in the summary statistics shown in **Table 5**, SSE was the only series to exhibit a substantial degree of positive skewness within the original return series, so the implied positive volatility skewness may be a reflection of this underlying return skewness. Moreover, the apparent degree of risk-seeking behavior among investors that this would imply may be a consequence not of risk-seeking, per se, but of investment activities related to specific structural aspects of this market, its price limits on daily fluctuations. As a story on the Financial Times Online [22] notes, such a consequence could reflect the activities of “zhangting gansidui or ‘limit-up kamikaze’ investors. These managers would get a small group of friends to buy into a stock, then talk up its prospects to a wider group of friends. ... The second group, who were in fact being duped, would rush to buy the stock and the manager would reap his profits by offloading his holdings when the stock hit its ‘limit up’-the maximum it could rise in a single day.” Thus, the anomalous parameter values for this series may be indicative not necessarily of a model that is misspecified and thus fits the data poorly but rather of a model that does in fact fit the data well but reflects features that are unique to this set of returns.

Thus, after the model has been reasonably re-specified to incorporate a non-zero mean component and a leptokurtic conditional distribution, the fitted SEMF model now appears to be a reasonable model, describing realistic return dynamics, for most, if not all, of the stock and index return series examined. But, after having made such adjustments to the model specification, an important observation from the results of fitting this model can be made—not only does incorporating a mean component into the model affect both how well the model fits the returns and the conclusions that the model yields about the dynamics of those returns, but estimating the mean in conjunction with the other features incorporated into the model specification also has a substantial impact on what the value of the mean is estimated to be.

Table 10 shows a comparison of the original sample means versus the means estimated from within the model itself (*i.e.*, the “model-fitted means”) for each of the series, along with the percentage change of the latter from the former. Interestingly, all of the sample means were positive, although those for the shorter IBM series, the two RDS series and the SSE series were not of sufficient magnitude in comparison to their standard errors to be significantly different from zero. The means μ estimated in conjunction with the model, however, are significantly different from zero ($p \approx 0.000$) in every single case, although the sign for the mean of the shorter IBM series’ returns has switched from positive to negative. But, while this is the only mean that actually changes signs, the means for all of the rest of the series also exhibit dramatic differences in value, with declines in value for about half of the series ranging in magnitude from 28.4% to 143.0% and increases in value for the other six with magnitudes of change of at least 36.5% all the way up to 127.7%. Across all the series, the average percentage change between the two estimates for the mean is only 9.9%, but the average magnitude of difference between the two estimates is a substantial 67.3%. Thus, taking both the volatility dynamics and leptokurtosis into account in the estimation process for these means can clearly have a substantial impact on the estimates for the mean that are obtained.

Table 10. Sample Means versus model-fitted means.

Series	Sample mean: original series	Mean (μ) from centered student’s t SEMF model	%-age difference
SPY	0.00042	0.00071	68.6%
GE	0.00052	0.00030	-41.3%
IBM (Long)	0.00047	0.00021	-56.0%
IBM (Short)	0.00036	-0.00015	-143.0%
KO	0.00061	0.00033	-46.7%
RDS.A	0.00047	0.00100	112.4%
RDS.B	0.00046	0.00104	127.7%
DJIA	0.00031	0.00042	36.5%
SP500	0.00046	0.00064	38.2%
HSI	0.00053	0.00074	41.2%
SSE	0.00056	0.00040	-28.4%

One final observation before leaving this section, though, is that there is a clear difference in the estimates for the mean (as well as for the other model parameters) for the shorter IBM series in comparison to the longer IBM series, even though the latter series is partially comprised of the former. Thus, even though a reasonable SEMF model can be fit to the longer series, it appears that such a model may be fitting an average of a variety of different types of return characteristics as the underlying company, IBM, has changed and evolved over time to meet the challenges of its competitive environment and as beliefs, preferences, and attitudes toward risk have changed and evolved over the generations of investors who have traded the stock. And, as a result of the combined effect of IBM's weaker competitive situation over the most recent decade especially compared to the situation it faced in the 1960s and 1970s, along with the concomitant changes in investors' attitudes toward IBM stock, the stock's return dynamics appear to have become much more nonstationary and at least slightly more volatile (in terms of σ_0), but subject to the effects of a long-term memory that has become stronger and fades away much more slowly. In general, as is illustrated by these IBM results, regardless of how well a given model fits, explains, and accounts for the dynamics within a long-horizon series of financial returns, it must always be remembered that, due to the dynamic and evolutionary nature of the stock market, the behavior such a model is describing will always be an average of a moving target as the behavior and characteristics of that stock or other financial entity evolves over time.

12. Conclusions

In theory, the SEMF model of Filimonov and Sornette (2011) appears to be a promising model to describe the evolution and volatility dynamics for financial returns, but when the original, conditionally Normal, zero-mean SEMF model specification was fit to a representative sample of stock and index return series, the parameter estimates obtained implied unrealistic behavior for nearly all the series under study. Such implied behavior included potentially explosive volatility dynamics ($\varphi < 0$) for four of the series (SPY, RDS.A and RDS.B, and SSE), and, conversely, a volatility memory that essentially erased itself after a single lag ($\varphi > 20$) for six of the other series (GE, the longer IBM series, KO, DJIA, SP500, and HSI). Reasonable parameter values implying realistic financial return behavior were found for only one of these sets of returns, the shorter IBM series.

The statistical context in which such results were obtained, however, failed to adequately account for and incorporate the effects of other well-known and widely recognized statistical characteristics of stock and index returns, including both a non-zero mean for the returns and a high degree of leptokurtosis, even after accounting for the effects of underlying volatility dynamics, among these returns. The analysis in this paper showed that properly taking such statistical distributional assumptions into account can dramatically change the conclusions yielded from a given model as the parameters for that model are allowed to better reflect the features of the data they were designed to describe.

Specifically, once the SEMF model specification is adjusted to incorporate both a mean component and a Student's t rather than a Normal conditional distribution, the fitted SEMF model is found to yield reasonable parameters for the volatility dynamics that would describe realistic behavior for most, if not all, of these series. Thus, these results will suggest that the stock return dynamics are subject to neither explosive volatility nor rapid-onset amnesia within the volatility memory process, but are instead merely conditionally leptokurtic, and with non-zero means, and the SEMF model, as re-specified, once again appears to be a promising model for stock returns.

Acknowledgements

The author wishes to thank his statistics dissertation advisor, Dr. George Terrell of the Department of Statistics at Virginia Tech, for ongoing advice, support, and encouragement during the course of the dissertation research that led to the work presented in this paper; Dean Michael Solt and the College of Business Administration at CSULB, for encouragement and financial support during the course of this research, and Mr. Nupur Shah, Head of Instructional Technology at the College of Business Administration at CSULB, for software and other support for this research.

References

- [1] Engle, R. (1982) Autoregressive Conditional Heteroskedasticity with Estimates of the Variance of U.K. Inflation. *Econometrica*, **50**, 987-1008. <http://dx.doi.org/10.2307/1912773>

- [2] Bollerslev, T., Chou, R. and Kroner, K. (1992) ARCH Modeling in Finance: A Review of the Theory and Empirical Evidence. *Journal of Econometrics*, **52**, 5-59. [http://dx.doi.org/10.1016/0304-4076\(92\)90064-X](http://dx.doi.org/10.1016/0304-4076(92)90064-X)
- [3] Bollerslev, T. (1986) Generalized Autoregressive Conditional Heteroskedasticity. *Journal of Econometrics*, **31**, 307-327. [http://dx.doi.org/10.1016/0304-4076\(86\)90063-1](http://dx.doi.org/10.1016/0304-4076(86)90063-1)
- [4] Mandelbrot, B., Fisher, A. and Calvet, L. (1997) A Multifractal Model of Asset Returns. Cowles Foundation Discussion Paper #1164.
- [5] Calvet, L. and Fisher, A. (2001) Forecasting Multifractal Volatility. *Journal of Econometrics*, **105**, 27-58. [http://dx.doi.org/10.1016/s0304-4076\(01\)00069-0](http://dx.doi.org/10.1016/s0304-4076(01)00069-0)
- [6] Calvet, L. and Fisher, A. (2004) How to Forecast Long-Run Volatility: Regime-Switching and the Estimation of Multifractal Processes. *Journal of Financial Econometrics*, **2**, 49-83. <http://dx.doi.org/10.1093/jfinec/nbh003>
- [7] Calvet, L. and Fisher, A. (2008) Multifractal Volatility: Theory, Forecasting, and Pricing. Academic Press, Burlington, MA.
- [8] Bacry, E., Delour, J. and Muzy, J.-F. (2001) Multifractal Random Walk. *Physical Review E*, **64**, Article ID: 026103. <http://dx.doi.org/10.1103/PhysRevE.64.026103>
- [9] Muzy, J.-F. and Bacry, E. (2002) Multifractal Stationary Random Measures and Multifractal Random Walks with Log Infinitely Divisible Scaling Laws. *Physical Review E*, **66**, Article ID: 056121. <http://dx.doi.org/10.1103/PhysRevE.66.056121>
- [10] Muzy, J.-F., Delour, J. and Bacry, E. (2000) Modeling Fluctuations of Financial Time Series: From Cascade Process to Stochastic Volatility Model. *European Physical Journal B*, **17**, 537-548. <http://dx.doi.org/10.1007/s100510070131>
- [11] Saichev, A. and Sornette, D. (2006) Generic Multifractality in Exponentials of Long Memory Processes. *Physical Review E*, **74**, 011111. <http://dx.doi.org/10.1103/PhysRevE.74.011111>
- [12] Saichev, A. and Filimonov, V. (2007) On the Spectrum of Multifractal Diffusion Process. *Journal of Experimental and Theoretical Physics*, **105**, 1085-1093. <http://dx.doi.org/10.1134/S1063776107110209>
- [13] Saichev, A. and Filimonov, V. (2008) Numerical Simulation of the Realizations and Spectra of a Quasi-Multifractal Diffusion Process. *JETP Letters*, **87**, 506-510. <http://dx.doi.org/10.1134/S0021364008090129>
- [14] Filimonov, V. and Sornette, D. (2011) Self-Excited Multifractal Dynamics. *Europhysics Letters*, **94**, Article ID: 46003. <http://dx.doi.org/10.1209/0295-5075/94/46003>
- [15] Filimonov, V. (2011) Modeling Financial Time Series: Effective Multifractality and Self-Excited Multifractal Process Paper Presented at Fractals and Related Fields II, Porquerolles Island, France. <http://congres-math.univ-mlv.fr/sites/congres-math.univ-mlv.fr/files/FIlimonov.pdf>
- [16] Sornette, D., Malevergne, Y. and Muzy, J.-F. (2004) Volatility Fingerprints of Large Shocks: Endogenous versus Exogenous, in Application of Econophysics. *Proceedings of the Second Nikkei Symposium on Econophysics*, Tokyo, 12-14 November 2002, 91-102.
- [17] Zeeuw van der Laan, P. (2011) Application of the SEMF Process to Financial Time Series. Master Thesis. [http://www.er.ethz.ch/content/dam/ethz/special-interest/mtec/chair-of-entrepreneurial-risks-dam/documents/dissertation/master thesis/MAS Zeeuw-van-der-Laan Final version nov 2011.pdf](http://www.er.ethz.ch/content/dam/ethz/special-interest/mtec/chair-of-entrepreneurial-risks-dam/documents/dissertation/master%20thesis/MAS_Zeeuw-van-der-Laan_Final_version_nov_2011.pdf)
- [18] Nelson, D. (1991) Conditional Heteroskedasticity in Asset Returns: A New Approach. *Econometrica*, **59**, 347-370. <http://dx.doi.org/10.2307/2938260>
- [19] Zhong, J. and Zhao, X. (2012) Modeling Complicated Behavior of Stock Prices Using Discrete Self-Excited Multifractal Process. *Systems Engineering Procedia*, **3**, 110-118. <http://dx.doi.org/10.1016/j.sepro.2011.11.015>
- [20] Spanos, A. (1995) On Theory Testing in Econometrics: Modeling with Nonexperimental Data. *Journal of Econometrics*, **67**, 189-226. [http://dx.doi.org/10.1016/0304-4076\(94\)01633-B](http://dx.doi.org/10.1016/0304-4076(94)01633-B)
- [21] Spanos, A. (1999) Probability Theory and Statistical Inference: Econometric Modeling with Observational Data. Cambridge University Press, Cambridge. <http://dx.doi.org/10.1017/CBO9780511754081>
- [22] Davies, P. and Rabinovitch, S. (2014) China: Funds on the Edge. Financial Times Online.