An Analysis of the Markov Switching Multifractal with a Market-based Application

Matthew A. Fagan

Abstract:

I examine the Multifractal Markov Switching (MSM) volatility model of Laurent Calvet and Adlai Fisher (2004, 2008). I apply the MSM to daily log-returns of the S&P 500, 100, and their volatility counterparts, the VIX and VXO, and construct an intuitive volatility model with more than a thousand states, parameterized by only four variables. The multifractal outperforms a Gaussian-distributed GARCH in- and out-of-sample, yet does not fare as well against a Student-t distributed GARCH. However, I find the MSM to be significantly superior in forecasting accuracy at horizons of 20 to 50 days over both the Normal- and Student-distributed GARCH. In the conclusion I suggest further applications of the MSM and its multifractal Markov structure.

In statistics, the Gaussian distribution is the standard, so-coined Normal distribution that financial scientists assume their data to follow. It is the starting point for all statistical analysis, with extensions and adjustments calling for a deep understanding of the mathematics behind the theory, and graduate level study. It is both the enabler and the bane of statistical success, providing means for computation, yet misleading with its clean theory that all too often fails to hold in empirical reality.

However, in our endeavors to better match the behavior of the empirical world, we are left with few alternatives. The mathematical extensions and conditions we layer upon our main statistical hallmark grow far and far more complicated as they reach for the boundary between theory and reality, requiring ever increasingly complex mathematics and proofs to maintain practicality. We must stay true to our normal distribution, no matter how false it may seem, because it is our main means of prediction. We forever increase our sample size and pray the central limit theorem holds.

When statistics is applied to money, the disparities between theory and reality become tangible in disastrous ways. With the same brush, our statistics paint a false image of our security, misrepresenting the true potential for catastrophic failure, while creating discrepancies for those closest to the markets to take advantage of. In finance, the purity of capitalism courses through its agents, with most endlessly seeking arbitrage, and disregarding the inevitable crashes they set in motion with the ripples of their wake.

In this gluttony we are reckless. In the realm of finance, we play with our math as though it had no bearing on reality, as though the ticks on our exchange boards were mere floating points that exist only for themselves. When it comes to our financial markets; when it comes to the security of our retirement and pension funds; when it comes to our reliance on money for our continued existence, we must be more responsible in our decision making. We must understand the connection between our financial theory and its significance to the real world. We must endlessly seek ever more effective means of scientific analysis, to better guide financial activity and better serve ourselves.

In this endeavor, I present a new statistical model for volatility modeling, called the Markov Switching Multifractal, or MSM for short. As does all financial theory, the intuition it is based on comes from nature; yet, the analogy it makes is incredibly novel, and presents an entirely new perspective on human interaction. In the following sections, I will introduce this natural analogy, track its extension to finance, and its most recent culmination in the MSM. With the MSM, I will turn to statistics and compare it to the hallmark of stochastic volatility modeling, the GARCH model. My analysis is in the form of Laurent Calvet and Adlai Fisher (2004,2008), the architects of the MSM, and applies their model to both the S&P 500 and the S&P 100, as well as the volatility indices of each, the VIX and VXO, respectively.

Fractal Foundations

Developed by Mandelbrot in 1963, fractal geometry is the study of "roughness," or the dimension between flat 2-dimensionsal geometry and clean 3-dimensional objects. The concept is very simple: "A fractal is a geometric shape that can be separated into parts, each of which is a reduced-scale version of the whole," [Mandelbrot (2008)]. The driving intuition here is one of *self-affinity* or self-similarity; break up a fractal shape, and one finds the same pattern pervading through magnitudes of order. Think of veins in the body, neurons in the mind, limbs on a tree, nodes on the internet; all of these natural phenomena emerge in fractal structure, internally replicating a simple pattern to achieve incredible complexity [Andriani P, McKelvey B (2009), Schwarz & Jersey (2011)].

To be more concrete, look at a tree. See the trunk, and how many limbs it initially splits into. Follow one of these branches, and one will observe the same splitting pattern into its smaller limbs. If one traces through all the way into the leaves, she will find the same branching pattern she observed in the beginning, when the trunk first split [Schwarz & Jersey (2011)]. This is the concept of self-similarity. (Another example of fractal intuition: an old fractal joke. What does the B in "Benoit B. Mandelbrot" stand for? Benoit B. Mandelbrot.)

With extension to finance, Mandelbrot demonstrated this self-similarity in asset pricing in his widely cited 1963 paper on cotton prices [Mandelbrot (1963)]. As for modern developments, Calvet, Fisher, and Mandelbrot first constructed a financial conception of fractal nature in 1997, with their paper "A Multifractal Model of Asset Returns", or MMAR for short. Without delving too deep in the complex mathematics behind the MMAR, the model essentially uses fractional Brownian motion, in conjunction with a fractal-based account for time. This fractal timeline is constructed by taking a line segment from 0 to 1 and breaking it into parts with a self-similar generator. This generator is then applied on each new segment, and each new segment produced from that, and so on ad infinitum. To illustrate, an example is presented in graph 1 [Mandelbrot (2008)].

1 THREE-PIECE FRACTAL GENERATOR (*top*) can be interpolated repeatedly into each piece of subsequent charts (*bottom three diagrams*). The pattern that emerges increasingly resembles market price oscillations. (The interpolated generator is inverted for each descending piece.)

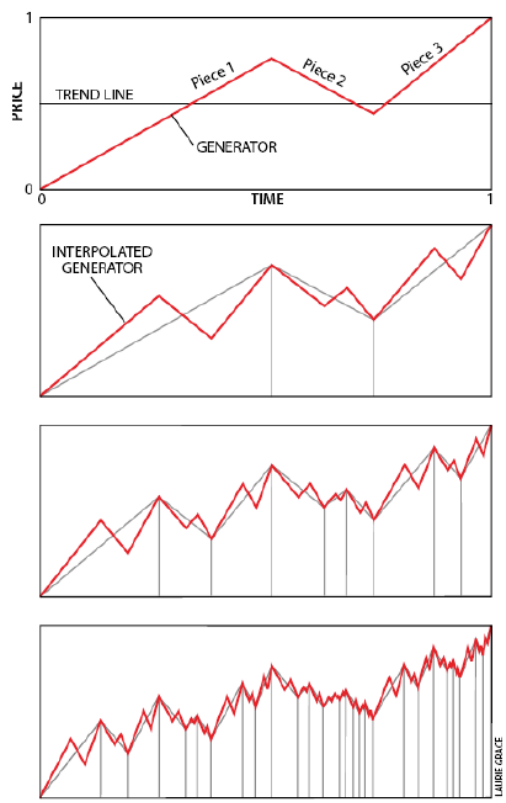

From the abstract to the MMAR: '[T]he MMAR contains long-tails, [and] long-dependence, the characteristic feature of fractional Brownian Motion (FBM) In contrast to FBM, the multifractal model displays long dependence in the absolute value of price increments, while price increments themselves can be uncorrelated." [Calvet, L., A. Fisher, and B. Mandelbrot. (1997)]. This is the purpose of applying multifractal constructs to model financial movements: to develop a model that accounts for long-tail probabilities, and long-memory dependence.

However, the MMAR was subject to one critical flaw. Due to its multiple dimensional construction and grid bound nature, it lacked stationarity, and was not practical in application. This is where the Mandelbrot, Calvet, and Fisher (1997) leave the reader in the last paragraph of their conclusion: "The main disadvantage of the MMAR is the dearth of applicable statistical methods." It could not be assessed by the same metrics as GARCH models, and thus lacked common ground to compare.

In truth, this mere brushing aside of the MMAR feels unjust. The amount captured by its construction is impressive, and worthy of recognition. Yet, Mandelbrot's incredible intelligence seems to have created something that only a contemporary can understand. In my paper, I focus on the MSM, as it is constructed in the same fashion as its econometric counterparts, but the MMAR provides invaluable information on the structure of financial data, even if not directly comparable to standard econometrics.

Regardless, this was the conflict that led me to Adlai Fisher and Laurent Calvet's Markov Switching Multifractal (MSM) model (2004, 2008). Using a Markov structure, Calvet and Fisher are able to construct a statistically stationary model of returns that incorporates a multifractal structure in the various states the Markov process can take on. The rest of this paper will delve into the MSM, and apply it to two market indices and their volatility counterparts.

The Markov Switching Multifractal (MSM)

The MSM is a mathematically rich and complex model that is simultaneously simple and intuitive to understand. Take r_t to be the log return of a financial asset or exchange rate, $r_t = ln(P_t/P_{t-1})$, where P_t is the price of the asset at time t, and P_{t-1} the price at time $t-1$. The MSM then defines returns r_t to be described by

$$r_t = \sigma(M_t)\varepsilon_t$$

I will now proceed to break up and describe these components.

M_t is the driver of the economy, a first-order Markov state vector with $\bar{k}$ components,

$$M_t = (M_{1,t}, M_{2,t}, \dots, M_{\bar{k},t}),$$

where the components of the M_t vector have equivalent marginal distribution, but grow and change at different frequencies. To display this, assume that we have constructed the volatility state vector up to $t-1$. For each $k \in (1, 2, \dots \bar{k})$, the next period multiplier is drawn from a fixed distribution M, with probability γ_k, or otherwise remains at its previous value, $M_{k,t} = M_{k,t-1}$. This can be displayed as

$$M_{k,t} \textit{ drawn from distribution } M \qquad w.p. \quad \gamma_k$$

$$M_{k,t} = M_{k,t-1} \qquad w.p \quad 1-\gamma_k$$

In this, the switching events and draws from M are assumed to be independent across k and t, a necessary assumption for the Markov process to hold. Following these principles, we require the distribution of M has the conditions, $M > 0$, and $E(M) = 1$.

γ, the driver of transition probabilities, is a vector, composed of $\bar{k}$ components,

$$\gamma = (\gamma_1, \gamma_2, \dots, \gamma_{\bar{k}}),$$

where each component is specified by

$$\gamma_k = 1 - (1 - \gamma_1)^{(b^{k-1})},$$

and $\gamma_1 \in (0,1)$ and $b \in (1, \infty)$. This is to ensure 1. that the probability of switching increases as the state, k, increases, 2. that this rate of increase slows as k increases, and 3. that the parameter γ_k remains < 1 for all k. The intuition is that the more stable, lower frequency multipliers switch less often than the higher frequency multipliers, therefore the γ_k vector increases in switching probability as k increases.

To construct $M_{k,t}$, the last piece missing is the M distribution. In my paper I apply the simplest, binomial MSM, in which the random variable M takes one of two values:

$$\begin{array}{ll} m_0 & w.p.\ 1/2 \\ 2 - m_0 & w.p.\ 1/2. \end{array}$$

This form of MSM, which sets $m_1 = 2 - m_0$, allows a single parameter, m_0, to establish the distribution M. In the code, this M distribution is inherent in the g(M) construction by dividing the γ_k values by 2, and separating them into 2 columns, 1 – $\gamma_k/2$, and $\gamma_k/2$.

Compiling the $M_{k,t}$ generators together, the MSM creates a state space A, with 2^k states that are each composed of $\bar{k}$ M multipliers. To illustrate, the value $M_{k,t}$ is constructed as follows:

$$\sigma(M_t) = \bar{\sigma}(\prod_{i=1}^{\bar{k}} M_{k,t})^{1/2}$$

Where, due to the independence of the $M_{k,t}$ multipliers, $\bar{\sigma}$ is approximately the long run standard deviation of returns.

In terms of conceptualizing this, it is useful to think of the M_t states as binary numbers with $\bar{k}$ placeholders. Take a $\bar{k}$ of 4. In binary representation, with 4 placeholders, the possible numbers are {0000, 0001, 0011, …. 1111}, which constitute a state space of 2^4, or 16 possible states. If we take 0 as a symbol for m_0, and 1 as a symbol for m_1, then the $(\prod_{i=1}^{\bar{k}} M_{k,t})$ value is the product of each of these values: 0000 = ${m_0}^4$; 0001 = ${m_0}^3 * m_1$; etc. Continuing this, the right side of the $\sigma(M_t)$ equation is completed by taking the square root of these products and scaling them by $\bar{\sigma}$. This is how the model is represented in the code and is how the state space is constructed. Each state is given an $M_{k,}$ value, and its probability of moving to a different M_k value at time $t + 1$ is equivalent to the combined probability of each possible switch from $k = (1, 2, \dots, \bar{k})$. As an example, taking this with the construction of $g(M)$ the probability of remaining in a given state equal to $\prod_{i=1}^{\bar{k}}(1 - .5 * \gamma_i)$.

All combined, the MSM is thus defined by 4 parameters:

$$\psi \equiv (m_0, \bar{\sigma}, b, \gamma_{\bar{k}}) \in \mathbb{R}_+^4 ,$$

with $\bar{k}$ left out as a means of determining model selection. This small parameter space is one of the primary strengths of the MSM: in traditional Markov regime models, a free parameter is generally required for each state; in this construction, the MSM is allowed a vast state space that is specified by only 4 parameters. For instance, when $\bar{k}$ = 10, the state space takes on 2^{10} states, or 1024.

For an in-depth explanation of how the MSM is estimated in closed-form, I direct the reader to Calvet and Fisher's book, "Multifractal Volatility: Theory, Forecasting, and Pricing," or their 2004 paper in the *Journal of Financial Econometrics*, "How to Forecast Long-Run Volatility: Regime Switching and the Estimation of Multifractal Processes," both cited below (Calvet, Fisher 2004, 2008), as I could only repeat, if not butcher, their explanations here.

GARCH Model

To assess the MSM's efficacy, I compare it to the industry standard in stochastic volatility analysis, the generalized auto-regressive conditional heteroscedasticity model, or GARCH as it is well known. In my testing of numerous GARCH models, I found the Normal- and Student-distributed GARCH modes to be the best models for each data set. In line with Calvet and Fisher (2004, 2008), I considered a fractionally integrated GARCH, as well as a generalized-error GARCH in the spirit of Mandelbrot, however in the former the model failed to converge and optimize, and in the latter the likelihoods obtained were far below those of the Normal- and Student-GARCH.

These models are defined as

$$x_t = h_t^{1/2} e_{t,}$$

Where e_t are standard Normal Gaussian innovations $(0,1)$ in the Normal-GARCH, and i.i.d. Student innovations with v degrees of freedom in the Student-GARCH, which requires an additional free parameter to estimate.

In the pursuit of parsimony, and corroboration from the literature [Calvet and Fisher (2004, 2008), Bollerslev (1987), Chuang et al. (2013)], I chose GARCH $(1,1)$ to be the best descriptor of the data sets, which constructs h_t as follows:

$$h_{t+1} = \omega + \alpha\varepsilon_t^2 + \beta h_t$$

Monte Carlo Analysis

Log-Likelihood Evaluation

To analyze the efficacy of the MSM code itself, I ran Monte Carlo simulations by generating data sets with the MSM code, and then evaluated the code's ability to select the true $\bar{k}$ value. My data is presented in table 1.

In my trials I found the code to effectively select the true $\bar{k}$, although with some caveats. In the table, I present the percentage of success in selecting the correct $\bar{k}$ value, and an interesting outcome came out of this analysis: as expected, $\bar{k}$ = 5 is selected far more often than all other $\bar{k}$ values at 77.75 %; however, the second most selected is $\bar{k}$ = 1. The caveat I mention: in assessing the LL selection, I included a line that selected $\bar{k}$ = 5 if it was *among* the highest log-likelihoods. In essence, if there were multiple maximum log-likelihoods from a given simulation and 5 was among the models, I counted 5 as the model chosen. In removing this caveat, the results skew down to the lower $\bar{k}$ models, with $\bar{k} = 1$ providing the maximum log-likelihood 33.75% of all 400 trials, and $\bar{k} = 5$ exclusively only accounting for 6.75% of all trials. In my simulated series and trials, multiple models provided the same likelihoods, suggesting some indifference in the model selection code.

This indicates that the function may underestimate the true $\bar{k}$ for a given data set; however, in looking at the average log-likelihood across all simulations, $\bar{k} = 5$ provides the highest value. This suggests that, although when considered individually the lower-$\bar{k}$ models produce comparable log-likelihoods to the true k=5 model, overall, the true model provides the largest log-likelihood. This reinforces the legitimacy of the code and provides confidence in its model-selection capabilities.

Running another set of Monte Carlo simulations displayed some interesting behavior with model selection tendencies based on parameter set-up. The second set of ln L analysis in Table A displays behavior that is more expected of the MSM – the percentage of time the true $\bar{k}$ of 5 is chosen is still the highest among all possible values, but we see more spread in the neighboring $\bar{k}$ values. The outlier here is the abnormally large ln L value for $\bar{k}$ = 8, suggesting there exists some trade-off between parameter selection and possible frequencies, and the possibility that a more complex MSM model can effectively represent data generated by less complex parameters.

I believe this second table better represents the model selection behavior of the MSM MLE function, as the values obtained in this analysis seemed more along the lines of my empirical data, and I worked to remove any failures or complications in the output.

Parameter Evaluation

To evaluate the parameters, I performed the same Monte Carlo simulation as in the log-likelihood evaluation, using the same parameters, yet with a $\bar{k} = 8$, as this resembled empirical results.

In comparing the mean estimate for each parameter to its respective true value, the estimated values are close to the true values, but fail to capture them in a 95% confidence interval. The most pronounced differences are in the estimates for b, but these are reasonable when considering the range of values that b can take. In the mathematical model, b can be any positive integer; in the code, it is considered initially to be nearest 1.5, 3, 6, or 20, and is more precisely tuned in the maximizing function after the fact. Due to its wide range, the larger root-mean square error and standard error are to be expected.

In analyzing the parameter outputs, I see the model trend towards estimates of 6, 1.5, .999, and 5 for b, $m_0, \gamma_{\bar{k}}$, and $\bar{\sigma}$, respectively. In the code, these are possible initial values the startingvals function proposes to enter into the likelihood function, suggesting the function either failed to alter these values, or found the initial values to be acceptable.

Noticing this possible failure, I reconstructed another Monte Carlo simulation with starting values of {3, 1.5, .9, .25} and a $\bar{k} = 5$. These results are listed in table B directly below the original trial. In this run, the sample mean values and error terms were much more reasonable, and far closer to the true values; in line with this, a 95% confidence interval does capture the true values for both $\bar{\sigma}$ and m_0, yet is still barely off with b and $\gamma_{\bar{k}}$. Comparing the two data sets to one another, I observe similar tendencies, such as the high RMSE for b, and the standard errors declining across the parameters. Given this new trial, and the efficacy of the MLE function in achieving precise parameter estimates on average, I am confident in the simulation means, and the MSM code overall.

Table 1 Monte Carlo ln L and Parameter Analysis

Table A - Ln L Analysis

$\bar{k}$	1	2	3	4	5	6	7	8
Mean ln L	5911.12	5504.26	5792.64	6526.94	6749.77	4101.57	5452.93	4519.97
ln L SEs	123.86	84.96	88.77	122.39	133.72	89.98	95.69	100.74
Percent of times k chosen	0.1575	0.0050	0.0025	0.0275	0.7775	0.0000	0.0275	0.0025
$Mean\ [\ln L\ (\bar{\boldsymbol{k}}') - \ln L(\bar{\boldsymbol{k}} = 5)]$	-838.65	-1245.51	-957.13	-222.83	0	-2648.20	-1296.83	-2229.80
SE of ln L Differences	110.02	110.81	105.90	88.66	0	95.69	109.82	93.62

$\bar{k}$	1	2	3	4	5	6	7	8
Mean ln L	651.22	839.99	902.12	907.38	912.30	908.04	910.63	922.72
ln L SEs	23.07	23.78	23.59	23.67	23.28	23.25	23.00	21.31
Percent of times k chosen	0.0000	0.0000	0.0148	0.2136	0.5015	0.1157	0.1157	0.0386
$Mean\ [\ln L\ (\bar{\boldsymbol{k}}') - \ln L(\bar{\boldsymbol{k}} = 5)]$	-261.09	-72.31	-10.19	-4.92	0	-4.26	-1.67	10.42
SE of ln L Differences	16.07	11.31	6.19	5.14	0	3.38	3.66	8.67

Table B - Parameter Analysis

$\bar{k}$ = 8	b	$\gamma_{\bar{k}}$	m_0	$\bar{\sigma}$
True Values	8	1.5	0.75	0.5
Mean Simulated Value	8.8472	1.3878	0.6853	0.3097
SE of simmed values	0.2399	0.0131	0.0134	0.0010
RMSe of simmed values	4.8671	0.2854	0.2748	0.1913

$\bar{k}$ = 5	b	$\gamma_{\bar{k}}$	m_0	$\bar{\sigma}$
True Values	3	1.5	0.95	0.25
Mean Simulated Value	3.1814	1.4565	0.9473	0.2498
SE of simmed values	0.0524	0.0073	0.0024	0.0007
RMSe of simmed values	1.0630	0.1512	0.0470	0.0131

These tables present the Monte Carlo estimation data for the MSM MLE function. The first set of Table A was generated by the True Values in first set of Table B; likewise for the second set. We see some discrepancy in model selection based on the generating parameters and k value. However, the true model is chosen most often on average, which provides confidence in the code's estimation procedure.

Empirical Application

My empirical analysis applies the MSM to two market exchanges, the S&P 500 and 100, and their volatility counterparts, the VIX and VXO respectively. I chose these due to their access, as Cal Poly lacks formal finance data providers to pull from, as well as their implications in the finance world. The S&P 500 is the most liquid and active stock exchange, with many pension funds, retirement funds, and long-term investments that depend on it. Its influence spreads throughout the finance world, and it serves as the standard benchmark for market comparison. If a model were found to better describe its volatility than current standard models, the model may help avoid disastrous pitfalls and failures such as the housing collapse, and better maintain all those persons' livelihoods invested the market.

Beyond the implications of the S&P indices, the volatility counterparts provide a kind of "second-moment" analysis, as the movements in these indices are entirely dependent on the S&P's movement. In addition, the volatility indices exhibit more general volatility, and widen the range of volatility to test the application of the MSM. The assumption behind the multifractal structure and its application is that it follows the structure of spontaneous human interaction and reaction, and with the volatility indices I get a direct example of human reaction.

Investigating the series and the relations of the VIX and VXO, I observe some interesting phenomena. Initially, I assumed the VIX would provide the largest returns around the GFC; however, in analyzing the actual returns, the maximum daily return for the VIX comes on February 27th, 2007, where the S&P 500 experienced a 3.5% decline. This drop, although a small blip on the screen in comparison to the volatility of the late 90s and late 2000s, was large relative to the most recent market behavior, during which the market was characterized by very mild volatility and moderate stable growth. The positive return on the VIX was due to the calm nature the S&P 500 exhibited prior to the 3.5% drop, displaying that the VIX's movement is more due to expectation. As the concept goes, its agents trade off of "fear," and sudden, unexpected changes spike cortisol levels in traders.

Finally, these indices are much more "mild" in terms of volatility than the exchange rates first analyzed by Calvet and Fisher. In this respect I aim to provide more breadth of application for the MSM, and determine if its volatility components, which are well attuned to capture the wild nature of exchange rates, can be tamed to the milder, standard volatility of market indices.

For sample size, my data closely mirrors the length of Calvet and Fisher's data sets for analysis: S&P 500 is matched to the VIX, contains 6552 data, from 1990 until December 31st 2015; the VIX is the total of its lifespan up to January 8th 2016, containing 6557 observations; the S&P 100 runs from 1983 until January 11th 2016, providing 8327 observations; and the VXO runs the whole of its lifespan, from 1986 until December 31st 2015, providing 7563 observations.

Figure 2 displays the time series graphs for all 4 series time, whereas Figure 3 displays the two S&P indices superimposed over the volatility indices, to illustrate the relativity of volatility. Figure 2 displays the. Apparent spikes are prevalent, as well as volatility clustering, and wild jumps throughout each series. The S&P 500, with its stable growth prior to the housing collapse and the hectic volatility following it display the persistence of volatility trends.

Figure 2

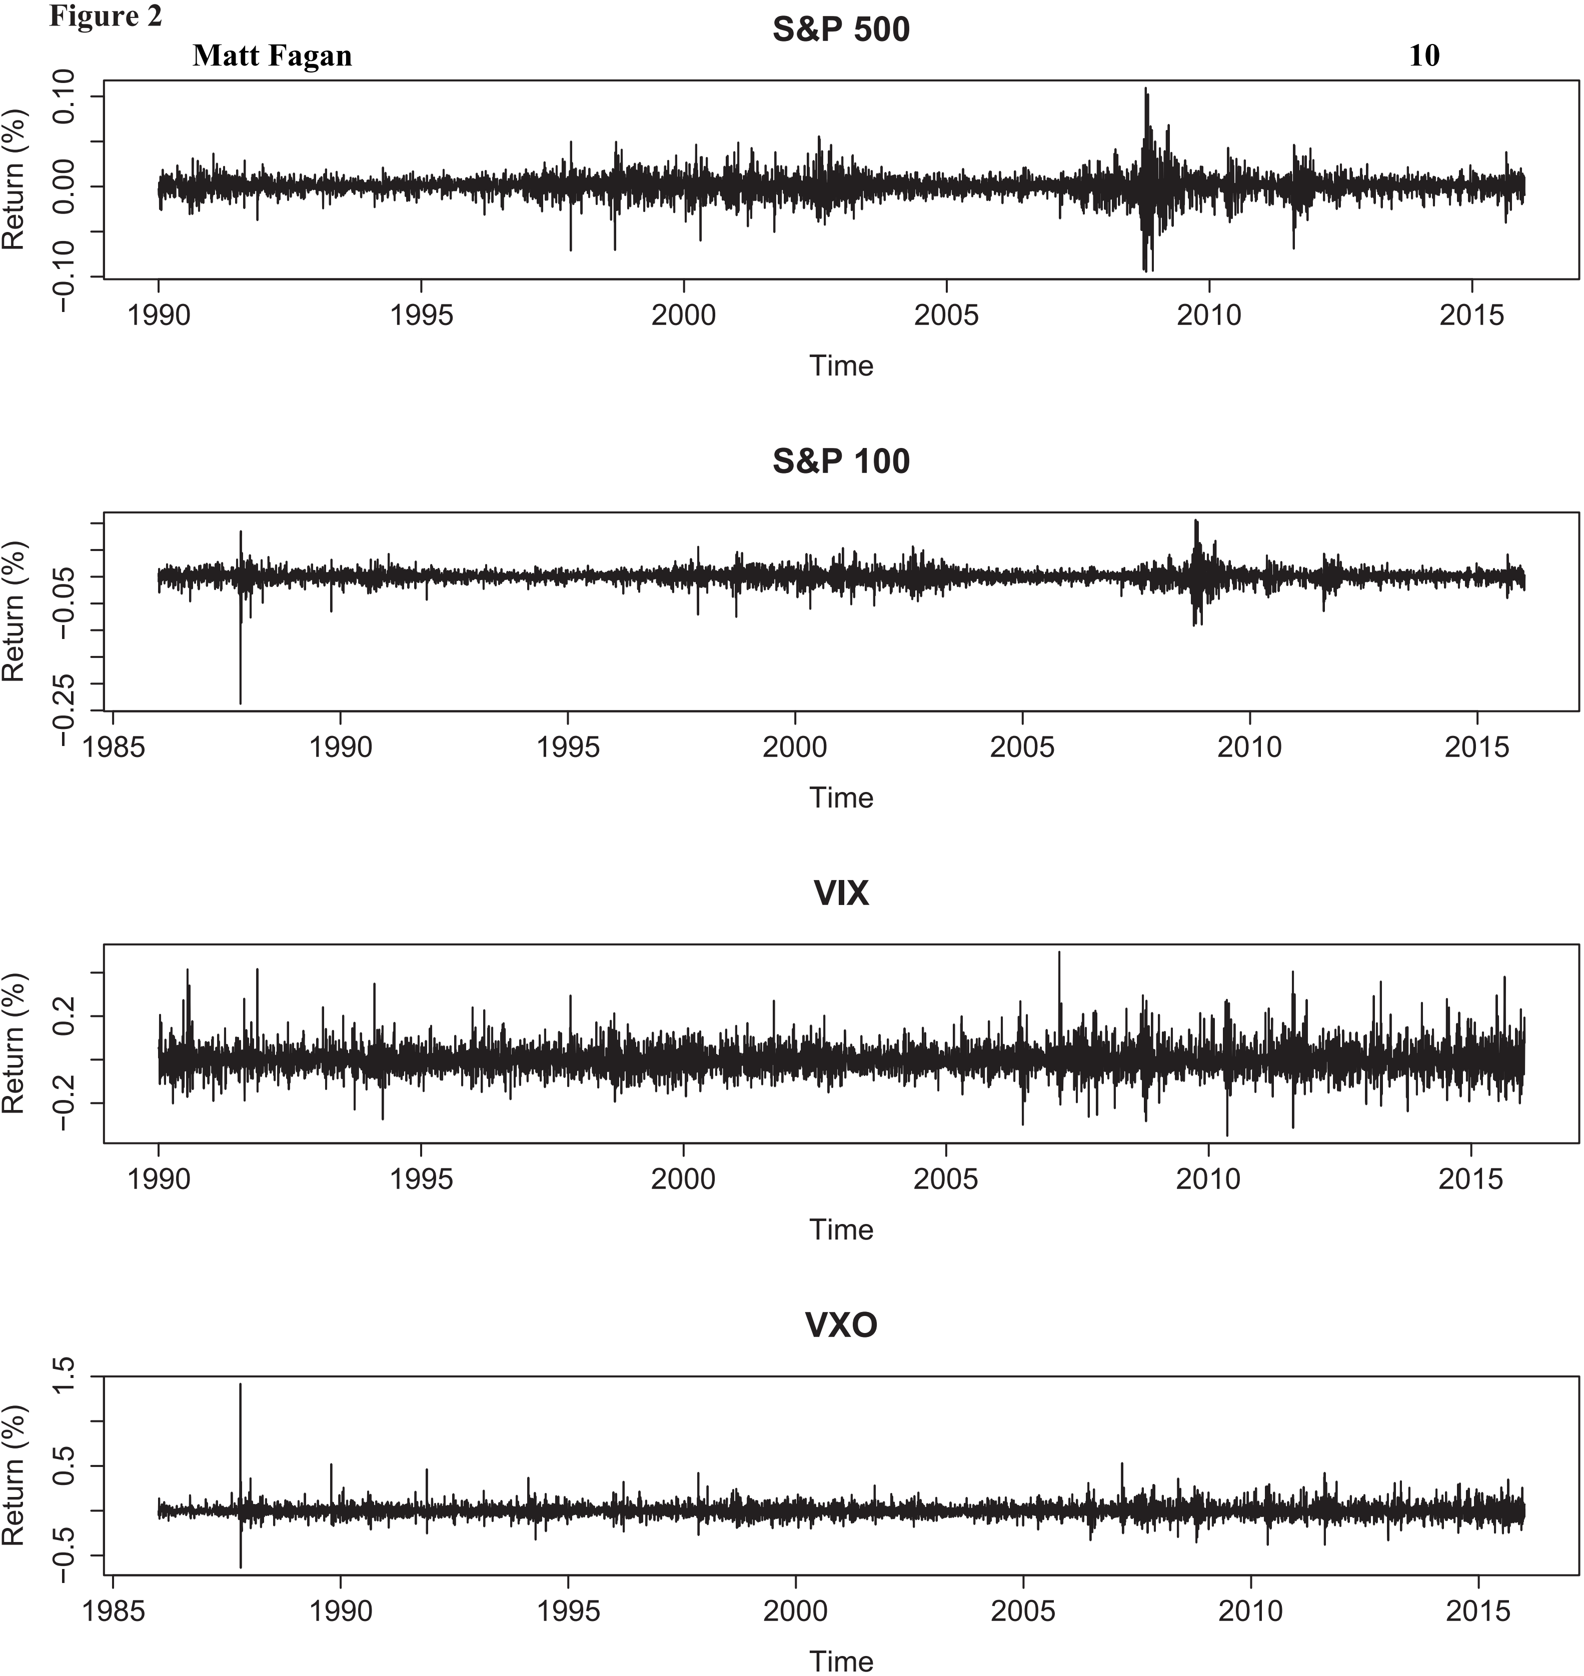

This figure presents the daily log-returns, in percent, of the two market indices and two volatility indices. The y-axis is left scaled to the specific series due to wild nature of the VXO with its outlier 141% returns, and to observe the inherent nature of volatility within each specific index.

Figure 3

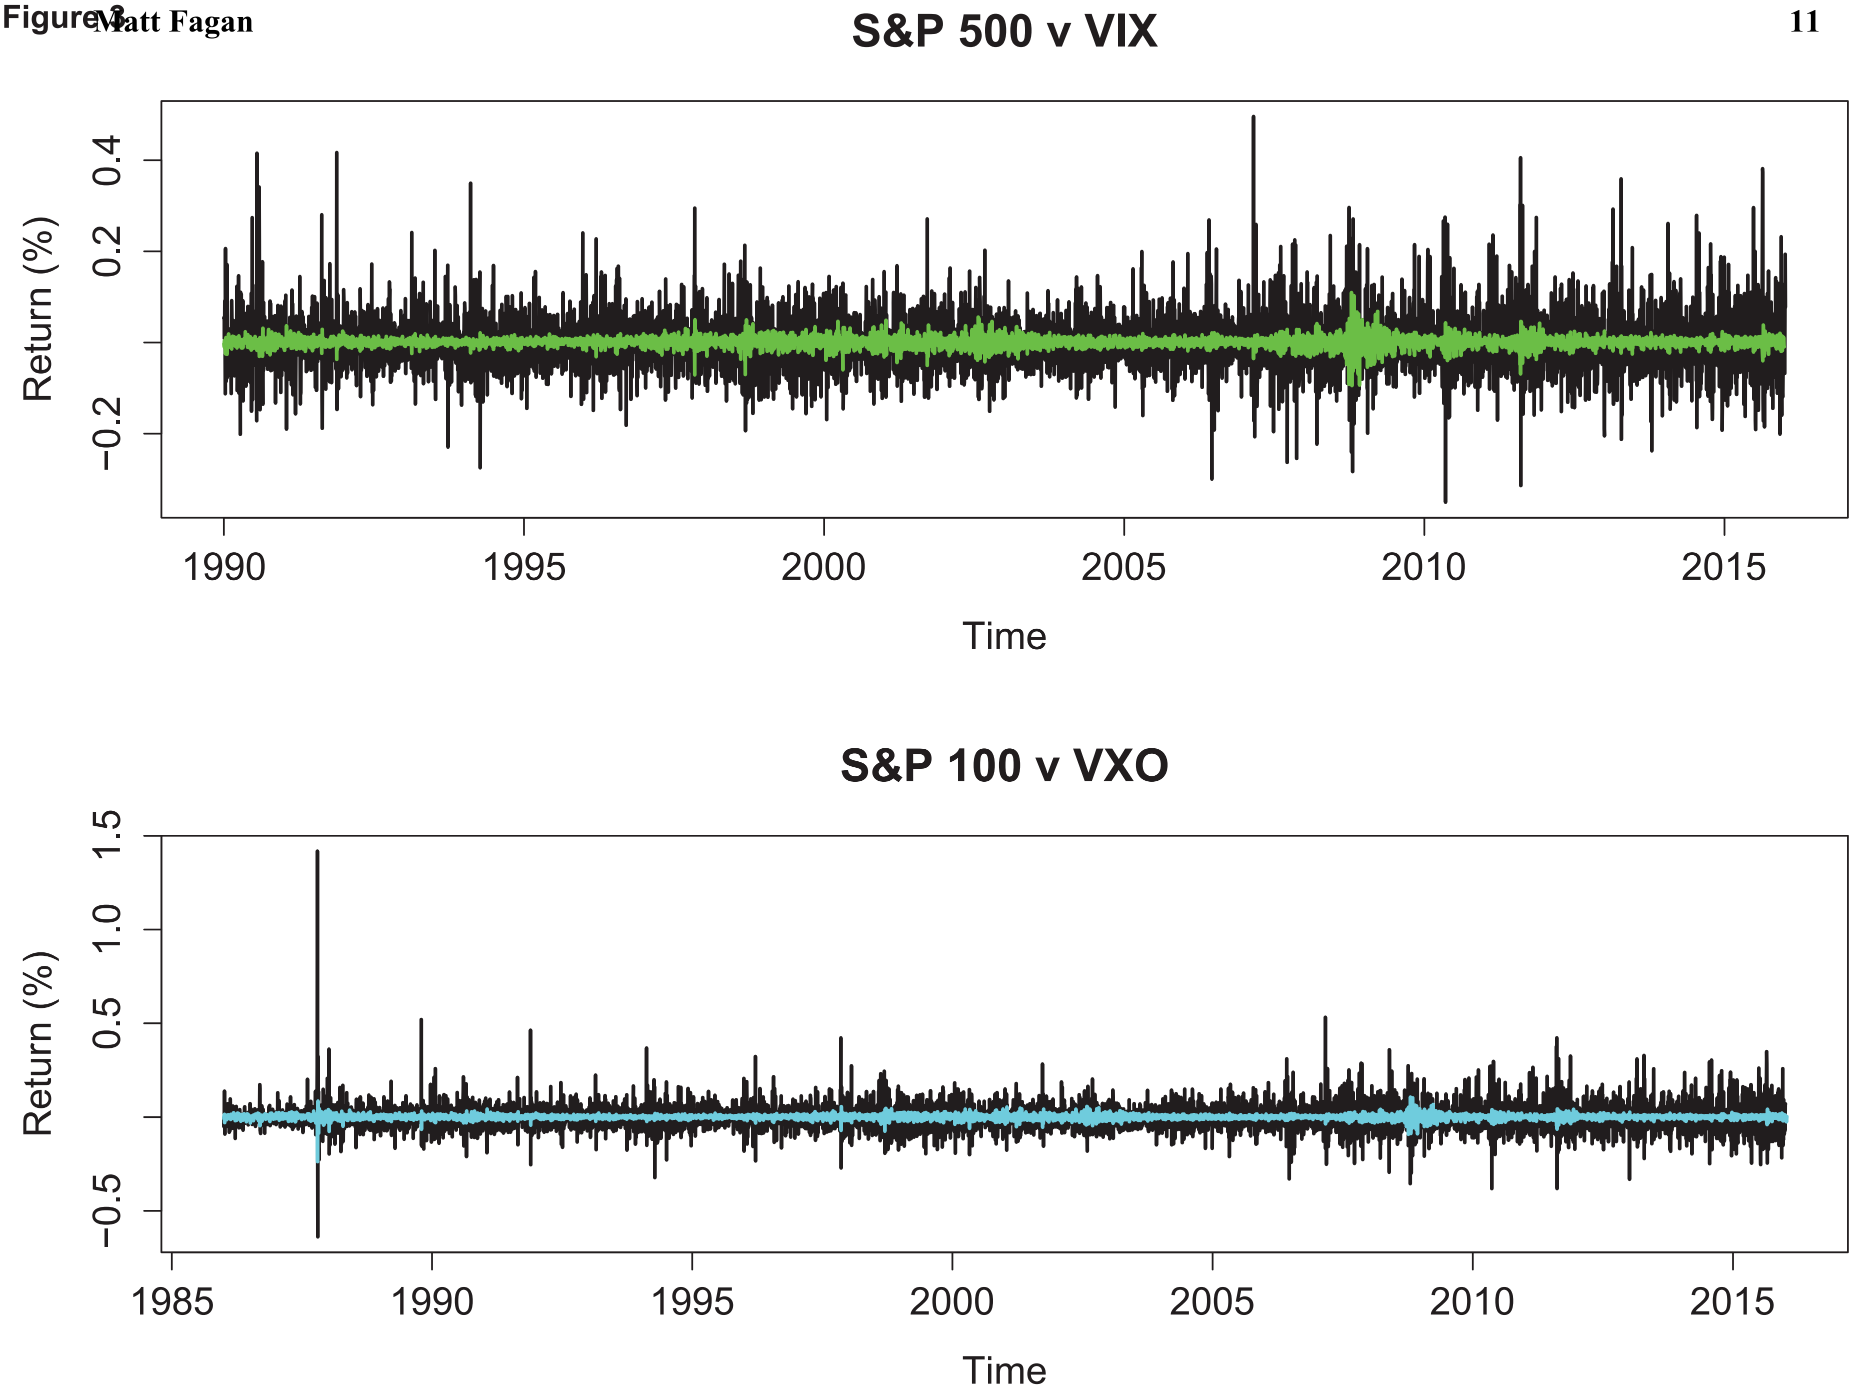

These graphs display the VIX and VXO with their S&P counterparts graphed on top in color to illustrate the counter movements with which the volatility indices react to the market.

ML Estimation Results

Table 2 presents all 4 datasets with their parameter estimates and log-likelihoods for each $\bar{k}$ level. The volatility components increase at a rate of $2^{\bar{k}}$, meaning that with $\bar{k} = 1$ only 2 states are present, whereas a $\bar{k}$ of 10 yields 1024. As was observed in Calvet and Fisher, m_0 tends to decline as $\bar{k}$ increases: the presence of more volatility states allows for more granular levels of volatility with less variable values of the $M_{k,t}$ multipliers.

The most persistent low level multiplier comes in the S&P 100 $(1/\gamma_1)$, with a duration of approximately 825 years, or 25 times the sample size; in contrast, the most persistent component for the VIX lasts roughly 79 days. In essence, this means we would expect the S&P 100's most stable frequency to change only once every 825 years, whereas the most stable for the VIX is expected to change every 79 days.

Using the same two indices as examples, by contrast, the most frequent component for the VIX switches about once every 1.4 days, and the same component for the S&P 100 switches about once every day. This indicates that the more stable volatility levels remain constant for the S&P 100, with only the highest frequency multipliers switching often, whereas the VIX multipliers change consistently and thoroughly across all frequencies.

Model Selection

Table 3 presents an analysis of the significance of log-likelihood differences between models to select the most accurate $\bar{k}$ value. To assess significance, I employ two tests: Vuong's (1989), and Clarke's (2007) significance test. Vuong's test is well known in model comparison literature and is the standard for comparing the difference in log-likelihoods. It evaluates the sum of the pointwise log-likelihood differences with respect to the variance of each model's log-likelihoods, and uses a BIC approach to penalize models with more parameters.

However, Vuong's test imposes a normal assumption on the distribution of log-likelihoods, and this assumption does not hold well in the MSM log-likelihood data; this is why I employ the lesser-known Clarke's test in conjunction with Vuong. Clarke's test looks at the pointwise log-likelihood differences as a binomial trial, with the number of differences greater than 0 as to the successes, and the total number of log-likelihoods as the number from which the successes are drawn. It then tests whether the probability of the number of successes is equal to .5, the expected probability of success if the two models are equivalent. The intuition is that, if the models were equivalent descriptors for the data, then half of the log-likelihoods would be greater than 0, and half would be less than 0.

The power in Clarke's parameter-free test is that we avoid assuming normality for the distribution of log-likelihoods; in fact, Clarke found his test to outperform Vuong's when the distribution of log-likelihoods displays high kurtosis, and that it was most effective when the data were distributed double-exponentially. In my analysis, the log-likelihoods displayed very high levels of kurtosis, so I found it apt to apply this test. However, a drawback to the Clarke's test is that it is somewhat binary in its decision process: with respect to capturing a 50% success rate in a binomial test, the confidence intervals are very tight, providing p-values with little grey area in between to assess significance level. This is why I use both tests: Clarke's to provide a more definitive answer, and Vuong's to provide more middle ground to the results.

In my table I present these two tests' p-values, where the null hypothesis is that the 2 models are equivalent descriptors for the data. The results are about as expected. As $\bar{k}$ increases, the log-likelihood increases, to a point. Across the board, the differences in log-likelihoods decrease as $\bar{k}$ increases, which follows the intuition, as a higher $\bar{k}$ can accommodate more granular changes in the data. The differences in log-likelihoods are least pronounced on those models nearest to the k with the highest log-likelihood, with the exception of the S&P 500, where $\bar{k}$s of 5 and 9 were found to be worse descriptors of the data at only the 25% significance level. The Clarke's test provides even more resounding evidence to dismiss the null hypothesis at the 1% significance level for all $\bar{k}$.

Table 2 All MSM Parameters (Model with highest ln *L* in bold)

$\bar{k}$	1	2	3	4	5	6	7	8	9	10
					S&P 500					
b	3.8110	9.2250	7.9217	4.7030	3.1782	3.0239	**2.1366**	2.2593	1.8247	3.3925
m_0	1.7297	1.6141	1.5188	1.4662	1.4222	1.4208	**1.3590**	1.3605	1.3102	1.4229
$\gamma_{\bar{k}}$	0.0248	0.0425	0.0617	0.0663	0.0617	0.0596	**0.0621**	0.0654	0.0785	0.0647
$\bar{\sigma}$	0.2174	0.2315	0.2469	0.2614	0.2480	0.2084	**0.2177**	0.2723	0.2075	0.9801
ln L	21137.40	21377.65	21431.25	21449.22	21464.26	21464.18	**21466.71**	21465.77	21465.69	21460.22
					S&P 100					
b	5.4520	11.6025	7.4998	4.7522	11.6479	7.1713	7.1499	7.1520	**5.2387**	2.7102
m_0	1.7124	1.5977	1.5349	1.4785	1.4501	1.4046	1.4043	1.4044	**1.3682**	1.3892
$\gamma_{\bar{k}}$	0.0208	0.0524	0.0689	0.0592	0.7268	0.8199	0.8183	0.8181	**0.9346**	0.0659
$\bar{\sigma}$	0.2261	0.2248	0.2544	0.2954	0.2709	0.2366	0.3062	0.2586	**0.1673**	0.7564
ln L	26695.58	26943.71	27028.28	27063.36	27066.59	27075.64	27075.06	27075.10	**27078.93**	27070.96
					VIX					
b	3.1410	7.8858	5.3058	4.0003	2.9393	2.3865	2.0802	1.8924	1.7607	**1.6616**
m_0	1.6395	1.5378	1.4616	1.4061	1.3633	1.3312	1.3065	1.2870	1.2709	**1.2571**
$\gamma_{\bar{k}}$	0.0893	0.2314	0.3593	0.5407	0.6058	0.6368	0.6583	0.6792	0.6954	**0.7076**
$\bar{\sigma}$	1.1362	1.1338	1.1168	1.0841	1.0539	1.0322	1.0214	1.0150	1.0103	**1.0064**
ln L	9339.32	9404.81	9422.23	9428.22	9430.32	9431.35	9431.90	9432.25	9432.51	**9432.69**
					VXO					
b	4.3370	7.7262	6.7933	7.5685	4.7721	4.3037	3.5349	2.7958	**2.3164**	3.6765
m_0	1.6834	1.5884	1.5160	1.5144	1.4494	1.4324	1.3952	1.3591	**1.3293**	1.3971
$\gamma_{\bar{k}}$	0.0862	0.1817	0.3754	0.4640	0.6846	0.6220	0.7617	0.8062	**0.8239**	0.7804
$\bar{\sigma}$	1.1979	1.1893	1.1063	1.4456	1.2707	1.5213	1.3743	1.2895	**1.1892**	1.2726
ln L	10467.32	10600.27	10629.20	10649.66	10654.26	10661.54	10665.18	10667.30	**10669.03**	10664.09

This table shows ML estimation results for the binomial multifractal model for all four market indices. Columns correspond to the number of frequencies *k* in the estimated model. The likelihood function increases with the number of volatility frequencies overall, yet reaches a maximum log-likelihood at various levels of complexity. For example, the S&P 500 requires only a *k* of 7, whereas the VIX multifractal model attains its highest log-likelihood at the maximum *k* of 10

Table 3 Multifractal Model Selection

$\bar{k}$	1	2	3	4	5	6	7	8	9	10
					S&P 500					
ln *L*	21137.4	21377.65	21431.25	21449.22	21464.26	21464.18	21466.71	21465.77	21465.69	21460.22
V Score	4.06827	1.100276	0.43808	0.216115	0.030269	0.031236	0	0.011615	0.012599	0.080112
V Test	0.0000	0.0000	0.0002	0.0013	0.2143	0.1932	0	0.1014	0.2432	0.0403
Clarke Score	3667	3442	3628	3763	3673	3189	0	4015	3267	3850
Clarke Test	0.0000	0.0000	0.0000	0.0000	0.0000	0.0010	0	0.0000	0.0096	0.0000
					S&P 100					
ln *L*	26695.58	26943.71	27028.28	27063.36	27066.59	27075.64	27075.06	27075.1	27078.93	27070.96
V Score	4.200972	1.481811	0.555023	0.170652	0.135236	0.036019	0.042341	0.041963	0	0.087296
V Test	0.0000	0.0000	0.0004	0.0970	0.0273	0.0878	0.1799	0.1438	0	0.2517
Clarke Score	4389	4256	4206	4240	4420	4211	4299	4223	0	4374
Clarke Test	0.0000	0.0011	0.0057	0.0021	0.0000	0.0051	0.0001	0.0037	0	0.0000
					VIX					
ln *L*	9339.316	9404.815	9422.226	9428.222	9430.325	9431.353	9431.897	9432.254	9432.508	9432.686
V Score	1.153061	0.344193	0.129179	0.055131	0.029159	0.016458	0.009739	0.00533	0.002192	0
V Test	0.0000	0.0000	0.0050	0.0439	0.0783	0.0949	0.0993	0.1039	0.1208	0
Clarke Score	3643	3901	3651	3887	3981	4017	4011	4056	4056	0
Clarke Test	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0
					VXO					
ln *L*	10467.32	10600.27	10629.2	10649.66	10654.26	10661.54	10665.18	10667.3	10669.03	10664.09
V Score	2.319406	0.790627	0.457983	0.222678	0.169754	0.086052	0.044196	0.019823	0	0.05678
V Test	0.0000	0.0003	0.0113	0.0021	0.0135	0.0051	0.0108	0.0287	0	0.0462
Clarke Score	4044	3865	3576	4194	4651	3973	4480	4506	0	4267
Clarke Test	0.0000	0.0015	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0	0.0000

This table reports the *t* -ratios and one-sided *p* -values for the log-likelihood difference of the model in a given column the model with the highest log-likelihood. The Vuong Test uses the Vuong (1989) method to asses log-likelihood difference, the Clarke Test uses the non-parametric method developed by Clarke (2003) to assess the same difference. A low *p* -value for the each test means the given model would be rejected in favor of the model yielding the highest log-likelihood

In Sample Comparison

In comparing the MSM to normal-GARCH models, the MSM significantly outperforms at all significance levels and for all indices. It provides higher log likelihoods and lower BIC values, and the Vuong (1989) and Clarke (2007) tests confirm these differences as statistically significant. The GARCH parameters are available in table 4, and the in-sample comparison data is available in table 5.

However, the MSM does not fare as well against the student-GARCH. Examining the log-likelihoods for the MSM and student-GARCH, the student-GARCH provides higher likelihoods for all indices except the VXO, where the MSM provides a slight bump over the student-GARCH model. The BICs follow a similar trend, with the MSM only coming in lower than the student-GARCH with the VXO, albeit by a mere thousandths-place difference.

To make claims as to the significance of these log-likelihood and BIC differences, I turn to Vuong's and Clarke's tests. With the student-GARCH, we fail to reject the null hypothesis that the MSM outperforms the student-GARCH at all valuable significance levels for all indices; Clarke's test affirms this position with its p-value of 1, indicating that the null hypothesis of a binomial distribution of .5 is rejected in favor of the student-GARCH. In context of the test, this means the number of positive differences in pointwise log-likelihoods was significantly below half in comparing the MSM to the student-GARCH.

Yet, there is some saving information here for the MSM. Looking across the indices, we see the p-values in favor of the student-GARCH decline with each index – or, rather, as general volatility increases. As the volatility of the series increases, it seems the evidence increases in favor of rejecting the null-hypothesis that the student-GARCH is a better descriptor for the data. In essence, we see the MSM's performance directly correlated with the volatility of the index, and its efficacy more pronounced when the index is characterized by extremely large outliers.

Looking into the VXO, we observe a single daily return of 1.42, or 142%, followed by a -0.64, or -64%, return two days later; in the original MSM analysis by Calvet and Fisher (2004, 2008), the currency exchanges they used were characterized by numerous instances of these extreme outliers, some even reaching as far as 500%. In addition, the standard deviations of my indices are significantly lower than those of Calvet and Fisher: as a highlight, the standard deviation for the VXO data is .069, whereas the lowest standard deviation of Calvet and Fisher's indices is .274 (CAD v USD Exchange Rate). These findings, together with those of Calvet and Fisher, suggest the MSM describes better those indices in which extreme values are more the norm and where changes are more "wild," whereas the GARCH remains superior for less volatile time series.

Table 4 GARCH Parameters

	ω	α	β	v
		S&P 500		
Student-GARCH	7.61E-07	0.0700	0.9251	7.1364
Normal-GARCH	1.19E-06	0.0797	0.9106	-
		S&P 100		
Student-GARCH	8.67E-07	0.0621	0.9317	6.3694
Normal-GARCH	1.60E-06	0.0838	0.9045	-
		VIX		
Student-GARCH	0.000306	0.1228	0.8019	5.7918
Normal-GARCH	0.000346	0.1088	0.8041	-
		VXO		
Student-GARCH	0.000206	0.1308	0.8310	4.9951
Normal-GARCH	0.000432	0.1589	0.7547	-

This table shows all ML parameter estimates for the two comparative GARCH models for all indices. Not included are the standard errors or significance tests, as these are more accessible in the included R code. All terms were deemed significant, except the Normal-GARCH coefficients.

Table 5 In-sample model comparison

	No. of Params	ln L	MSM BICs	BIC p-value vs. multifractal	
				Vuong	Clarke
			S&P 500		
Binomial MSM	4	21466.71	-6.5474		
Norm-GARCH	3	21366.07	-6.5180	0.0000	0.0091
Student-GARCH	4	21481.76	-6.5519	0.8952	1.0000
			S&P 100		
Binomial MSM	4	27078.93	-6.4995		
Norm-GARCH	3	26875.56	-6.4518	0.0003	0.0000
Student-GARCH	4	27116.98	-6.5087	0.9869	1.0000
			VIX		
Binomial MSM	4	9432.69	-2.8718		
Norm-GARCH	3	9179.47	-2.7959	0.0000	0.0000
Student-GARCH	4	9436.11	-2.8728	0.4510	1.0000
			VXO		
Binomial MSM	4	10669.03	-2.8167		
Norm-GARCH	3	10199.17	-2.6936	0.0000	0.0000
Student-GARCH	4	10665.79	-2.8158	0.3278	1.0000

This table presents the in-sample comparison statistics. The BICs are calculated by $(-2 \ln L + NP \ln T)/T$, where T is the length of a given series, and NP the number of parameters. The last two columns display p-values of a given test of the MSM vs a corresponding GARCH model. A low p-value indicates the given GARCH model would be rejected in favor of the MSM. The first p-value uses Vuong's (1989) method, and the second uses Clarke's (2003) method.

Out of Sample Forecasting

I now test each model's forecasting ability at horizons of 1, 5, 10, 20, and 50 days. To evaluate out of sample forecasting I estimate each model's parameters on the first half of each data set and use the last half for fitting forecast data. To compare each model's forecasting efficacy, I used the Mincer-Zarnowitz ordinary least squares regression model, as it was the means employed by Calvet and Fisher (2004) in their own paper. The equation is as follows:

$$x_{t+n}^2 = \alpha + \beta \mathrm{E}_t[x_{t+n}^2] + \mu_t$$

Where n is the forecast horizon, and $\mathrm{E}_t[x_{t+n}^2]$ is the expected value of the squared return at time t from n days prior. The traditional Mincer-Zarnowitz uses γ_0 and γ_1 as opposed to α and β respectively; I choose these as they are easier to represent in the code. A perfect predictor would have an α of 0 and a β of 1, therefore I use these as my null hypotheses in evaluating the coefficients. The results are presented in table 6.

Results

Put simply, all models are comparatively equal predictors, with idiosyncratic differences arising across indices and horizons.

For 1 step ahead predictions, the MSM is only slightly more effective in forecasting for the VXO, as we fail to reject the null hypothesis that the Beta coefficient is equal to 1 at the 5% significance level, whereas this same hypothesis is rejected for the normal- and student-GARCH at the 5% level; however, the 3 models are equivalent at the 10% and 20% significance levels. Beyond this, in comparing the mean squared errors (MSEs) for each predictor, the MSM is rejected in favor of the normal-GARCH and the student-GARCH for the VIX at the 5% level. This is the only instance across all horizons and across all indices that the difference in MSEs is significant, with the difference in MSEs being approximately zero for all other indices at each prediction horizon.

As the prediction horizon increases, some interesting trends emerge. The following discussion focuses on the 50-step horizon data. With respect to the less volatile indices (S&P 500, S&P 100), the Beta coefficient stays closer to 1, and for the S&P 500 I fail to reject the null hypothesis that the Beta is equal to 1, whereas I reject this for the normal- and student-GARCH at all possible significance levels. With the more volatile indices (VIX, VXO), the Beta coefficient for the MSM, although not equal to 1, is significantly lower than the Betas for both the normal and standard-GARCH, with each value equal to about 6.3, 72.5, and 11.9, respectively. Looking at the p-values for the Beta coefficient, I see the null hypothesis for the MSM is rejected at the 1% significance level, whereas I fail to reject this for the Normal- and Student-GARCH; however, this is due to their massive standard errors, which show that the 95% confidence interval for the Beta coefficients contain 0 as well, meaning that the Beta term for the Normal- and Student-GARCH models are altogether insignificant.

This anomaly is even more pronounced in estimating the VXO, with the Beta coefficient blowing up to about 912 in the Normal-GARCH fit, and the standard error reaching 2188. This suggests the Normal-GARCH is a wholly ineffective predictor for the VIX and VXO. The student-GARCH proves to be far more reasonable than its normally distributed counterpart, with its Beta coefficient remaining at a mild 1.42, and its error term 1.08; however, this too contains 0 in its 95% confidence interval, leading me to conclude that the term is insignificant. The MSM, with its beta of .65 and its standard error of .13, remains significant, although the null hypothesis that the Beta coefficient is equal to 1 is rejected.

Looked at holistically, these results suggest that the MSM underperforms at shorter time horizons, but performs significantly better as time horizon increases. In addition, this disparity in performance with respect to the GARCH is more pronounced in the more volatile indices (VIX, VXO), suggesting that the MSM better captures outliers, more extreme values, and more erratic time series behavior in large prediction horizons. Of the two GARCH models, the normal GARCH's loss of efficacy as time horizon increases is most pronounced, with the Beta growing to capture more extreme observations at the 20 step horizon, whereas the student-GARCH remains on par with the MSM up until the 50 step horizon.

This is due to the nature of the models. In examining the projected values, the normal GARCH trends towards the mean of the data, projecting the estimated mean of squared returns for each return. The student-GARCH accounts for more variability with its t-distribution, and therefore projects near the mean with some miniscule variability. In addition, the student-GARCH more accurately captures the observed mean of the data: for example, with the VXO, the normal-GARCH projects .00398 constant at the 50-step horizon, whereas the student-GARCH projects .00408, .00405, .00403, .00415, etc., and the observed mean of the squared returns is .00473.

Herein lies the power of the MSM. Although it too does eventually trend down towards the mean, due to its Markov driver and conditional independence, the projections persist longer, and the model takes far longer to reach the mean. In testing persistence, the MSM still maintains some variability at a 1000-step horizon, and requires a prediction horizon greater than 5000 to project a constant mean. However, the mean estimated by the MSM tends to be above that of the observed mean – for example, with the VXO data, the MSM trends towards .00553 as an estimate for squared returns, whereas the observed mean is .00473.

Table 6 Out-of-Sample Forecast Data

		1 Step Ahead				
		Estimate	Std. Error	Pr(>|t|)	R^2	MSE
S&P 500	MSM-Alpha	-8.9E-05	0.000012	0.0000	0.17016	0.00000023
	MSM-Beta	1.89502	0.073073	0.0000		---
	Norm-Alpha	0.000004	0.000009	0.6550	0.233051	0.00000022
	Norm-Beta	0.962555	0.030497	0.2200		(0.4118)
	Std-Alpha	0.000005	0.000009	0.5900	0.224656	0.00000022
	Std-Beta	0.9422	0.03057	0.0590		(0.4234)
S&P 100	MSM-Alpha	-6.7E-05	0.00001	0.0000	0.170178	0.00000019
	MSM-Beta	1.538285	0.052616	0.0000		---
	Norm-Alpha	0.000001	0.000008	0.8530	0.215857	0.00000018
	Norm-Beta	1.015833	0.029995	0.5980		(0.4261)
	Std-Alpha	-3E-06	0.000008	0.6960	0.194012	0.00000019
	Std-Beta	1.08748	0.034336	0.0110		(0.4613)
VIX	MSM-Alpha	-0.00059	0.000387	0.1260	0.067647	0.00011979
	MSM-Beta	1.257776	0.081387	0.0020		---
	Norm-Alpha	-0.00852	0.000294	0.0000	0.447945	0.00007093
	Norm-Beta	3.438573	0.066661	0.0000		(0.0281)
	Std-Alpha	-0.00643	0.000275	0.0000	0.41429	0.00007525
	Std-Beta	2.808291	0.058309	0.0000		(0.0408)
VXO	MSM-Alpha	0.000057	0.000388	0.8830	0.064339	0.00015360
	MSM-Beta	0.962567	0.059582	0.5300		---
	Norm-Alpha	0.001543	0.000304	0.0000	0.070267	0.00015263
	Norm-Beta	0.81726	0.048262	0.0000		(0.4892)
	Std-Alpha	0.00132	0.000308	0.0000	0.074381	0.00015195
	Std-Beta	0.817997	0.046851	0.0000		(0.4818)

These tables report the results from the Mincer-Zarnowitz OLS regression of a given n-ahead forecast. For an unbiased forecast we expect an Alpha of 0 and a Beta of 1; the *p* -values provided in the coefficent estimates compare the observed estimates to these null hypotheses. In the last two columns I report the R^2 and MSEs for each model and index, as well as the *p* -values of MSE difference in parentheses underneath the MSEs. A low *p* -value indicates the given model produces a significantly lower MSE than the MSM

Table 6 cont.

		5 Ahead				
		Estimate	Std. Error	Pr(>|t|)	R^2	MSE
S&P 500	MSM-Alpha	-9.8E-05	0.000013	0.0000	0.14994	0.00000028
	MSM-Beta	1.976499	0.082214	0.0000		---
	Norm-Alpha	0.000006	0.000009	0.5060	0.218941	0.00000022
	Norm-Beta	0.944824	0.031185	0.0770		(0.4062)
	Std-Alpha	0.000009	0.00001	0.3420	0.208145	0.00000022
	Std-Beta	0.910214	0.031023	0.0040		(0.4207)
S&P 100	MSM-Alpha	-6.7E-05	0.000011	0.0000	0.142419	0.00000023
	MSM-Beta	1.549049	0.058899	0.0000		---
	Norm-Alpha	-1E-06	0.000008	0.9260	0.205059	0.00000019
	Norm-Beta	1.038196	0.031682	0.2280		(0.4029)
	Std-Alpha	0.000002	0.000009	0.8530	0.176063	0.00000019
	Std-Beta	1.064736	0.035696	0.0700		(0.4475)
VIX	MSM-Alpha	-0.00126	0.00064	0.0490	0.027194	0.00012976
	MSM-Beta	1.524526	0.158494	0.0010		---
	Norm-Alpha	-0.00076	0.00061	0.2110	0.025393	0.00012533
	Norm-Beta	1.495116	0.160938	0.0020		(0.5028)
	Std-Alpha	0.000468	0.00048	0.3290	0.026204	0.00012522
	Std-Beta	1.115844	0.118213	0.3270		(0.5015)
VXO	MSM-Alpha	0.000979	0.000463	0.0350	0.029023	0.00016476
	MSM-Beta	0.792293	0.074239	0.0050		---
	Norm-Alpha	0.000456	0.000548	0.4050	0.024361	0.00016024
	Norm-Beta	1.153869	0.118193	0.1930		(0.5082)
	Std-Alpha	0.002146	0.00036	0.0000	0.031035	0.00015914
	Std-Beta	0.681611	0.061717	0.0000		(0.4965)

Table 6 cont.

		10 Ahead				
		Estimate	Std. Error	Pr(>|t|)	R^2	MSE
S&P 500	MSM-Alpha	-8.3E-05	0.000014	0.0000	0.115973	0.00000025
	MSM-Beta	1.847448	0.089148	0.0000		---
	Norm-Alpha	0.000018	0.00001	0.0580	0.175716	0.00000023
	Norm-Beta	0.860429	0.032586	0.0000		(0.4220)
	Std-Alpha	0.000022	0.00001	0.0260	0.16735	0.00000024
	Std-Beta	0.820359	0.031995	0.0000		(0.4328)
S&P 100	MSM-Alpha	-5.1E-05	0.000012	0.0000	0.109042	0.00000021
	MSM-Beta	1.435399	0.063598	0.0000		---
	Norm-Alpha	0.000009	0.000009	0.3150	0.163999	0.00000020
	Norm-Beta	0.985274	0.034494	0.6690		(0.4182)
	Std-Alpha	0.000015	0.000009	0.0950	0.141147	0.00000020
	Std-Beta	0.986637	0.037734	0.7230		(0.4520)
VIX	MSM-Alpha	-0.00204	0.000943	0.0310	0.015345	0.00012679
	MSM-Beta	1.8096	0.251084	0.0010		---
	Norm-Alpha	-0.00458	0.001189	0.0000	0.018147	0.00012643
	Norm-Beta	2.669026	0.340569	0.0000		(0.4957)
	Std-Alpha	-0.00111	0.000765	0.1480	0.017667	0.00012649
	Std-Beta	1.615547	0.208927	0.0030		(0.4964)
VXO	MSM-Alpha	0.001964	0.00052	0.0000	0.01342	0.00016223
	MSM-Beta	0.615157	0.085056	0.0000		---
	Norm-Alpha	-0.00629	0.001456	0.0000	0.016943	0.00016165
	Norm-Beta	2.867196	0.352904	0.0000		(0.4938)
	Std-Alpha	0.002462	0.000443	0.0000	0.014548	0.00016204
	Std-Beta	0.64359	0.085481	0.0000		(0.4980)

Table 6 cont.

		20 Ahead				
		Estimate	Std. Error	Pr(>|t|)	R^2	MSE
S&P 500	MSM-Alpha	-5.2E-05	0.000015	0.0010	0.073611	0.00000026
	MSM-Beta	1.585544	0.098384	0.0000		---
	Norm-Alpha	0.000037	0.00001	0.0000	0.119379	0.00000025
	Norm-Beta	0.732845	0.034843	0.0000		(0.4445)
	Std-Alpha	0.000042	0.00001	0.0000	0.113568	0.00000025
	Std-Beta	0.682773	0.033389	0.0000		(0.4515)
S&P 100	MSM-Alpha	-2.2E-05	0.000013	0.0780	0.070283	0.00000022
	MSM-Beta	1.2375	0.069806	0.0010		---
	Norm-Alpha	0.000025	0.000009	0.0070	0.106367	0.00000021
	Norm-Beta	0.893561	0.040193	0.0080		(0.4497)
	Std-Alpha	0.000037	0.000009	0.0000	0.092595	0.00000021
	Std-Beta	0.855931	0.041574	0.0010		(0.4688)
VIX	MSM-Alpha	-0.00155	0.001709	0.3660	0.003712	0.00012859
	MSM-Beta	1.759235	0.485247	0.1180		---
	Norm-Alpha	-0.00895	0.005126	0.0810	0.001841	0.00012884
	Norm-Beta	4.057425	1.532503	0.0460		(0.5029)
	Std-Alpha	-0.0012	0.002227	0.5910	0.001793	0.00012884
	Std-Beta	1.714089	0.654801	0.2760		(0.5029)
VXO	MSM-Alpha	0.002876	0.000603	0.0000	0.005047	0.00016381
	MSM-Beta	0.450924	0.100617	0.0000		---
	Norm-Alpha	-0.0209	0.012607	0.0970	0.000892	0.00016450
	Norm-Beta	6.597361	3.160385	0.0770		(0.5072)
	Std-Alpha	0.003836	0.000734	0.0000	0.001061	0.00016447
	Std-Beta	0.36342	0.162597	0.0000		(0.5070)

Table 6 cont.

		50 Ahead				
		Estimate	Std. Error	Pr(>|t|)	R^2	MSE
S&P 500	MSM-Alpha	0.000023	0.000017	0.1790	0.021234	0.00000028
	MSM-Beta	0.962298	0.114214	0.7410		---
	Norm-Alpha	0.000085	0.000011	0.0000	0.029468	0.00000028
	Norm-Beta	0.403087	0.040522	0.0000		(0.4906)
	Std-Alpha	0.000089	0.000011	0.0000	0.030284	0.00000028
	Std-Beta	0.364708	0.036157	0.0000		(0.4896)
S&P 100	MSM-Alpha	0.00005	0.000014	0.0000	0.01984	0.00000023
	MSM-Beta	0.742262	0.080856	0.0010		---
	Norm-Alpha	0.000068	0.000011	0.0000	0.025776	0.00000023
	Norm-Beta	0.63005	0.060114	0.0000		(0.4921)
	Std-Alpha	0.000083	0.00001	0.0000	0.026984	0.00000023
	Std-Beta	0.569056	0.053044	0.0000		(0.4905)
VIX	MSM-Alpha	-0.01638	0.006618	0.0130	0.002811	0.00012976
	MSM-Beta	6.320644	1.988766	0.0080		---
	Norm-Alpha	-0.23562	0.447141	0.5980	-0.00022	0.00013016
	Norm-Beta	72.50552	134.9373	0.5960		(0.5047)
	Std-Alpha	-0.03493	0.064682	0.5890	-0.00019	0.00013015
	Std-Beta	11.91873	19.48207	0.5750		(0.5046)
VXO	MSM-Alpha	0.001666	0.000787	0.0340	0.006322	0.00016476
	MSM-Beta	0.670383	0.134764	0.0140		---
	Norm-Alpha	-3.6211	8.701799	0.6770	-0.00022	0.00016585
	Norm-Beta	911.8896	2188.06	0.6770		(0.5114)
	Std-Alpha	-0.00038	0.004445	0.9310	0.000192	0.00016578
	Std-Beta	1.421449	1.084407	0.6980		(0.5107)

Conclusion

In looking at the entirety of my analysis, I can conclude the MSM does outperform the Normal-GARCH model in describing the S&P 500, S&P 100, VIX, and VXO; however, I cannot conclude the MSM outperforms the Student-GARCH. It seems the mixture of Gaussians that the MSM uses to describe the data tend to overfit the realm of empirical probability observed in the market indices, whereas a Student-t distribution, with its longer tails and larger accommodation and outliers than a standard Gaussian, is significantly sufficient in describing the volatility of the S&P 500 and 100, with a declining significance towards the VIX and VXO.

However, consider this a mere rake across the surface of the MSM. With its numerous extensions in continuous time form, multivariate form, and equilibrium theory, the vast realm of possibility available to the MSM is not to be discounted by these findings. Looking at the trend in increased efficacy as overall volatility increases, it leads me to suspect the MSM is a superior predictor for increasingly volatile series. In application, this could be useful for options pricing, better accounting for the range of possible values; Chuang et al. (2013) suggest this in their paper focused on individual stock options of the S&P 100. Going further, as the world of financial engineering grows more creative in its mathematical ingenuity to create more volatility and other reactionary indices, I suspect the volatility observed in these will be far more wild, and the MSM holds promise to account for their realm of possibility.

Looking at the mathematical tools which constitute the MSM, the multifractal Markov generator is also extremely promising for application beyond mere financial and daily return description. Given my romantic tendencies and evidence in nature, I truly believe the multifractal mathematics to be the best descriptor for human interaction. With the growing trend in looking at more micro-economics; in applying multifractal philosophy to model market behavior and to better describe Human agents as opposed to Econs, in the terms of Kahneman (2011); in looking at the interaction between agents and how these coalesce to produce the macro-effects we experience [McKelvey, Lichtenstein, and Andriani (2012), McKelvey, Yalamova (2011)]; I believe the multifractal markov model can be applied to this realm of interaction, and help bridge the gap between empirical dependence and statistical independence. Statistical science, the only available tool for economic scientists, requires clean distributions for finite moments; yet, in the empirical world, the arena for applying economic science, we observe disjoint and heterogeneous results, the entirety of which finite moments and clean distributions fail to capture.

So much of our investigations into nature, into its fundamental interaction and laws, are made more fruitful in applying multifractal intuition. We see results in bloodwork, in trees, in avalanches, in our neural networks, in our social networks, in the architecture of galaxies and space [Andriani, McKelvey, (2007), Andriani, McKelvey (2007), Mandelbrot (2010), Schwarz & Jersey (2011)]. Coming back to economics, no realm of work is more tangible, or reaches further in impact, than the work with money; and, the mathematics and models we apply to our data, the level of understanding we possess of our fiscal nature, could help practitioners prevent disaster, and allow for more global prosperity.

In our work as economists, we must take our efforts seriously, we must take these results from nature as legitimate, and we must work to conceive of new ways to apply what they reveal of our world to our monetary policy. In the spirit of McKelvey and Andriani, I call for a new null hypothesis [Andriani P., McKelvey, B. (2007), Andriani P., McKelvey, B. (2007)]. I call for an assumption of interaction and interdependence, as this is the true nature of animate, and inanimate, behavior. And, in the efforts of applying and assessing this new null hypothesis, I see multifractal statistics as an invaluable asset.

As for monetary application, I believe the MSM would better suit individual firm returns, and the options tied to each. I believe the MSM could provide better guides to hedging strategies and the nature of interaction – especially with more extensions, such as a multivariate MSM, or continuous-time, both of which Calvet and Fisher have outlined in their book on multifractal volatility. I strongly recommend their book to any financial academics and practitioners who seek a new perspective on finance. Looking beyond the realm of finance, I believe the Multifractal Markov driver could aid models of economic agents and their interaction, and describe other observed behavior, such as the sociological analysis of “viral” trends and memes, or the social tokens that catch on and become widespread, as well as the behavior of post-market drift, and other anomalies to the perfect market hypothesis. It is time we reassess our process of inquiry, and work ceaselessly to better capture the breadth of our behavior

References

Andriani P, McKelvey B (2007) Beyond Gaussian averages: redirecting international business and management research toward extreme events and power laws. *Journal of International Business Studies*, 38: 1212-1230

Andriani P, McKelvey B (2009) From Gaussian to Paretian thinking: Causes and implications of power laws in organizations. Organ. Sci. 20(6):1053–1071.

Bollerslev, T. (1987). ‘‘A Conditional Heteroskedastic Time Series Model for Speculative Prices and Rates of Return.’’ Review of Economics and Statistics 69, 542--547.

Calvet, L., A. Fisher, and B. Mandelbrot. (1997). Cowles Foundation Discussion Papers no. 1164--1166, Yale University. Available at http://users.math.yale.edu/~bbm3/web_pdfs/Cowles1164.pdf

Calvet L., Fisher A. (2004). How to Forecast Long-Run Volatility: Regime Switching and the Estimation of Multifractal Processes. *Journal of Financial Econometrics*, 2(1): 49-83.

Calvet L., Fisher A. (2008). *Multifractal Volatility: Theory, Forecasting, and Pricing*. Burlington, MA: Elsevier Inc.

Chuang W., Huang T., Lin B. (2013). Predicting volatility using the Markov-switching multifractal model: Evidence from S&P 100 index and equity options. *North American Journal of Economics and Finance* 25: 168– 187

Clarke, K. (2007). A Simple Distribution-Free Test for Nonnested Model Selection. *Political Analysis*, 15(3), 347-363. Retrieved from http://www.jstor.org/stable/25791898

Kahneman, D. (2011). Thinking, fast and slow. New York: Farrar, Straus and Giroux.

Mandelbrot, B. B. (1963), The Variation of Certain Speculative Prices, *Journal of Business* **36**, 394-419

Mandelbrot B. B. (2008, Sept 15). How Fractals Can Explain What’s Wrong with Wall Street. *Scientific American.* Retrieved from http://www.scientificamerican.com/article/multifractals-explain-wall-street/

Mandelbrot, B. B. (2010, Feb). Benoit Mandelbrot: Fractals and the art of roughness [Video file]. Retrieved from http://www.ted.com/talks/benoit_mandelbrot_fractals_the_art_of_roughness

McKelvey B., Yalamova R. (2011) Explaining What Leads Up to Stock Market Crashes: A Phase Transition Model and Scalability Dynamics. *Journal of Behavioral Finance*, 12: 169–182

McKelvey, B., Lichtenstein, B.B. and Andriani, P. (2012) ‘When organisations and ecosystems interact: toward a law of requisite fractality in firms’, *Int. J. Complexity in Leadership and Management*, Vol. 2, Nos. 1/2, pp.104–136.

Schwarz, M. (Producer), & Jersey, B. (Director). (2011). *Hunting the Hidden Dimension* [Motion Picture]. United States: Nova.

Vuong, Q. (1989). Likelihood Ratio Tests for Model Selection and Non-Nested Hypotheses. *Econometrica* 57, 307--333.