Extreme Risk and Fractal Regularity in Finance

Laurent E. Calvet and Adlai J. Fisher

Abstract. As the Great Financial Crisis reminds us, extreme movements in the level and volatility of asset prices are key features of financial markets. These phenomena are difficult to quantify using traditional models that specify extreme risk as a rare event. Multifractal analysis, whose use in finance has considerably expanded over the past fifteen years, reveals that price series observed at different time horizons exhibit several forms of scale-invariance. Building on these regularities, researchers have developed a new class of multifractal processes that permit the extrapolation from high-frequency to low-frequency events and generate accurate forecasts of asset volatility. The new models provide a structured framework for studying the likely size and price impact of events that are more extreme than the ones historically observed.

1. Introduction

Fractals use invariance principles to parsimoniously model complex objects at multiple scales. They have proven to be of major importance in mathematics and the natural sciences, as this issue illustrates. Fractals also offer enormous potential benefits for the field of finance, in particular for modeling the price of traded securities, for computing the risk of financial portfolios, and for managing the exposure of institutions. These benefits should become more apparent as the adoption of fractal methods by the financial industry continues to gain ground. The fields of finance and economics also play a singular role in the intellectual history of fractals. Benoît Mandelbrot first discovered evidence of fractal behavior in financial returns, household income, and household wealth in the late 1950's and early 1960's and subsequently found

Electronic copy available at: http://ssrn.com/abstract=2126466

similar patterns in the coastline of Brittany and other natural phenomena. These observations prompted the development of the fractal and multifractal geometry of nature ([**M82**]).[1]

In financial markets, the distribution of price changes is of key importance because it determines the risk, as well as the potential gains, of a position or a portfolio of assets. Different investors may measure price changes at different horizons. For instance, a high-frequency trader may look at price changes over microseconds, while a pension fund or a university endowment may have horizons of months, years, or decades. For this reason, researchers have tried to uncover invariant properties in the distribution of price changes observed over different time increments. The French economist Jules Regnault (1863) may have been the first to observe that the standard deviation of a price change over a time interval of length Δt scales as the square root of Δt ([**R**]). This observation motivated Louis Bachelier ([**Ba**]) to formalize the definition of the Brownian motion as a possible model of a stock price. That is, he postulated that price changes are Gaussian, identically distributed and independent. While some of these assumptions are controversial, Bachelier opened up the field of financial statistics, which has remained vibrant ever since.

In the early 1960's, Benoît Mandelbrot discovered that price changes have much thicker tails than the Gaussian distribution permits ([**M63**]). He proposed to replace the Brownian motion with another family of scale-invariant processes with independent increments – the stable processes of Paul Lévy ([**L24**]). Let $p(t)$ denote the logarithm of a stock price or an exchange rate. If $p(t)$ follows a Brownian or a Lévy process, the distributions of returns $p(t+\Delta t) - p(t)$ observed over various horizons Δt can be obtained from each other by linear rescaling. Self-similarity turns out to be a rather crude approximation of price changes. In addition, the Brownian motion and Lévy processes assume that price changes are independent, while there is substantial evidence that the size of price changes,$|p(t+\Delta t) - p(t)|$, is persistent ([**E82**]). In fact for many series, the size of price changes is a long-memory process characterized by a hyperbolically declining autocorrelogram ([**DGE**],[**D**]).

Since the mid–1990's, researchers have uncovered different forms of scale-invariance in financial returns, based on multifractal moment-scaling. [**G**], [**CFM**], and [**CF02**] found evidence that the moments

[1] Benoît Mandelbrot eventually returned to finance in the mid-1990's, when he taught a course on *Fractals in Finance* at Yale University. We attended this course and went on to develop with Benoît Mandelbrot the first applications of multifractals to financial series.

Electronic copy available at: http://ssrn.com/abstract=2126466

of the absolute value of price changes, $\mathbb{E}(|p(t+\Delta t)-p(t)|^q)$, scale as power functions of the horizon Δt. This observation motivated the construction of the first family of multifractal diffusions ([**CFM**], [**CF01**], [**BDM**]). The Markov-Switching Multifractal (MSM) captures well the dynamics of asset prices, including fat tails, long memory and moment scaling ([**CF01**], [**CF04**], [**CF**]). MSM assumes that the size of the price change is driven by a product of components that have invariant distributions but different degrees of persistence. Each component follows a Markov chain in its own right. MSM thus constructs a multifractal measure stochastically over time, which improves over earlier algorithms with predetermined switching points. MSM permits proper statistical estimation and generates accurate forecasts of the conditional distribution of returns, and therefore of the downside risk of a position.

Fractals provide a natural structure for modeling large risks. A common approach in finance is to represent an asset price as the sum of an Itô diffusion and a jump process. The diffusion describes "ordinary" fluctuations, while jumps are meant to capture "rare events." Difficulties in empirical implementation of this type of approach are readily apparent. Because rare events are modeled as intrinsically different from regular variations, inference on rare events must be conducted on a small set of observations and is therefore imprecise. Of increasing importance, researchers would like to understand the implications for asset prices of events that have never been previously observed ("peso effects," [**R88**], [**B06**], [**G12**], [**W**]). However, statistical inference on an empty set is a notoriously challenging exercise! Fractal modelling offers a promising tool in light of growing awareness of the importance of rare events. Scale invariance properties permit researchers to model all price variations using a single data-generating mechanism. As a consequence, models constructed using fractal principles are extremely parsimonious. A small number of well-identified parameters, combined with testable assumptions on scale-invariance, specify price dynamics at all timescales. Implications for rare events, even those more extreme than have been observed in existing data, are a natural outcome of a fractal approach to modeling financial prices.

The organization of the paper is as follows. Section 2 discusses early fractal models and presents fractal regularities in financial markets. Section 3 discusses the Markov-Switching multifractal model and its empirical applications. In Section 4, we show the pricing implications of fractal models. Section 5 concludes.

2. Fractal Regularities in Financial Markets

2.1. Self-Similar Proposals. Let $P(t)$ denote the price at date t of a financial asset, such as a stock or a currency, and let $p(t)$ denote its logarithm. The asset's return between dates t and $t + \Delta t$ is given by:

$$p(t + \Delta t) - p(t).$$

For over century, one of leading themes in finance has been to understand the dynamics of returns.

In his 1900 doctoral dissertation, the French mathematician Bachelier introduced an early definition of the Brownian motion as a model of stock prices ([**Ba**]). The Brownian motion assumes that the return $p(t+\Delta t)-p(t)$ has a Gaussian distribution with mean $\mu\Delta t$ and variance $\sigma^2\Delta t$, or more concisely

$$p(t + \Delta t) - p(t) \overset{d}{=} \mathcal{N}(\mu\Delta t; \sigma^2\Delta t).$$

The Brownian motion now pervades modern financial theory and notably the Black–Merton–Scholes approach to continuous time valuation ([**BS**], [**M**]). Its lasting success arises from its tractability and consistency with the financial concepts of no-arbitrage and market efficiency.

In the late 1950's and early 1960's, advances in computing technology made it possible to conduct more precise tests of Bachelier's hypotheses. In a series of pathbreaking papers, Benoît Mandelbrot ([**M63**], [**M67**]) uncovered major departures from the Brownian motion in commodity, stock and currency series. His main observation was that the tails of return distributions are thicker than the Brownian motion permits. Benoît Mandelbrot understood that this phenomenon was not a mere statistical curiosity, as some researchers suggested at the time, but major failures of the Brownian paradigm. In layman's terms, extreme price changes are key features of financial markets that the Brownian motion cannot capture. Since the purpose of risk management is to weather financial institutions against storms, underestimating the size of these storms, as the thin-tailed Brownian model does, is a recipe for financial disaster, panic and bankruptcy. The Great Financial Crisis reminds us of the severity of the shocks that can be unleashed on financial institutions, especially those who took on excessive risk as a result of poor risk management models and practices.

Benoît Mandelbrot advocated that financial prices should be modelled by a broader class of stochastic processes.

Definition 1. **(Self-similar process)** The real-valued process $\{p(t); t \in \mathbb{R}_+\}$ is said to be self-similar with index H if for all $c > 0$,

$n > 0$, and $t_1, ..., t_n \in \mathbb{R}_+$, the vector $(p(ct_1), ..., p(ct_n))$ has the same distribution as $(c^H p(t_1), ..., c^H p(t_n))$, or more concisely:

$$(p(ct_1), ..., p(ct_n)) \stackrel{d}{=} (c^H p(t_1), ..., c^H p(t_n)).$$

The Brownian motion is self-similar with $H = 1/2$. The stable processes of Paul Lévy ([**L24**]) are self-similar processes with independent increments and Paretian tails:

$$\mathbb{P}\{|p(t + \Delta t) - p(t)| > x\} \sim K_\alpha \, \Delta t \, x^{-\alpha}$$

as $x \to +\infty$, where $\alpha = 1/H \in (0; 2)$ and K_α is a positive constant. Lévy processes have thicker tails than the Brownian motion and are more likely to accommodate large price changes ([**M63**], [**M67**]). One difficulty with Lévy-stable processes, however, is that they have infinite variances, which is at odds with the empirical evidence available for a large number of financial series ([**BG**], [**FR**], [**AB**]). Furthermore, infinite variances pose theoretical difficulties, since a large body of the portfolio selection and asset pricing literature is based on the trade-off between mean and variance (see, e.g, [**CV**], [**MK**], [**S**], [**T**], [**M**]).

Fractional Brownian motions represent another important class of self-similar processes ([**K**], [**M65**], [**MV68**]). Their increments are stationary, normally distributed, and strongly dependent. Their autocorrelation declines at the hyperbolic rate:

$$Cov(r_{t+n}; r_t) \sim H(2H - 1)n^{2H-2} \text{ as } n \to \infty,$$

where $r_t = p(t) - p(t - 1)$ denotes the return on a time interval of unit length. The fractional Brownian motion rarely represents a practical model of asset prices, because their unconditional distributions are Gaussian. Long memory in returns is inconsistent with arbitrage-pricing in continuous time ([**MS**]).

The applicability of self-similar processes to finance is limited by another, common shortcoming. By definition, returns observed at various frequencies have identical distributions up to a scalar renormalization:

$$p(t + \Delta t) - p(t) \stackrel{d}{=} (\Delta t)^H p(1).$$

Most financial series, however, are *not* exactly self-similar, but have thicker tails and are more peaked in the bell at shorter horizons. This observation is consistent with the economic intuition that higher frequency returns are either large if new information has arrived, or close to zero otherwise. For this reason, self-similar processes cannot be fully satisfactory models of the asset returns.

2.2. Evidence on Fat Tails and Long Memory. Following [**M63**] and [**M67**], a number of researcher have documented Paretian tails in financial series, albeit with a tail index larger than 2 and therefore a finite variance. The early evidence was parametric ([**BG**], [**FR**], [**AB**]). In the 1970's, statisticians developed precise techniques for the nonparametric estimation of the tail index of a distribution ([**Hi**]). This lead to a series of papers showing that financial series have finite tail indexes that are most often strictly larger than 2 (e.g., [**KK**], [**KSV**], [**PMM**], [**JD**], [**LP**], [**G09**]).

In the 1990's, researchers uncovered evidence of strong persistence in volatility ([**D**], [**DGE**]). Long memory is often defined by a hyperbolic decline in the autocovariance function as the lag goes to infinity. For every moment $q \geq 0$ and every integer n, let

$$\rho_q(n) = Corr(|r_t|^q, |r_{t+n}|^q) \tag{2.1}$$

denote the autocorrelation in levels. We say the asset exhibits long memory in the size of returns if $\rho_q(n)$ is hyperbolic in n.

These key features of financial data can be seen by casual observation of standard asset returns. Figure 1, Panel A, shows the Japanese Yen / U.S. dollar exchange rate series from 1973, following the demise of the Bretton-Woods system of fixed exchange rates, to the present day. The series shows both fat tails and volatilility clustering at different time scales, including over periods as long as several years, as occurs in the presence of long memory. Panel B shows the same features in a long time series of daily U.S. stock index returns, obtained from the Center for Research in Securities Prices (CRSP).

[Insert Figure 1 here]

Figure 2 shows long memory in the volatility of these two series. We display on a double logarithmic scale the lag length n on the horizontal axis versus the autocorrelation of squared returns $\rho_2(n)$ on the vertical axis. For both series the plots are approximately linear, indicating a hyperbolic decay in volatility.

[Insert Figure 2 here]

2.3. Multifractal scaling. In the mid-1990's, the combination of fat tails and long memory in volatility led researchers to consider that asset prices may exhibit multifractal properties. This suggested that financial series should exhibit multifractal moment-scaling, a generalization of self-similarity defined by:

$$\mathbb{E}\left(|p(t+\Delta t)-p(t)|^q\right)=c_q(\Delta t)^{\tau(q)+1} \tag{2.2}$$

which holds for every (finite) moment q and time interval Δt. A self-similar process satisfies this relation, with $\tau(q) = Hq - 1$. The process p is said to be strictly multifractal if (2.2) holds for a strictly concave function $\tau(q)$. Strict multifractality has been observed in fields as diverse as fluid mechanics, geology, or astronomy. We now also have strong evidence of strict multifractal moment-scaling in a variety of financial series, including currencies and equities ([**CF02**], [**CFM**], [**G**], [**VA**]).

Figure 3 shows evidence of multifractal moment scaling in the Yen / U.S. dollar exchange rate series and in CRSP stock index returns. The panels of the figure plot the partition interval Δt on the horizontal axis versus an empirical estimate of $\mathbb{E}\left(|p(t+\Delta t)-p(t)|^q\right)$ on the vertical axis. The empirical estimate is obtained by taking the sample analogue of (2.2), as explained in the figure caption, for a variety of moments q. The dotted lines in the figure represent the scaling implied by Brownian motion, which satisfies self-similarity with $H = 1/2$. The panels both show evidence of moment-scaling that is linear in Δt, but the scaling coefficients $\tau(q)$ cannot be captured as a linear function of a single index H. These empirical facts are characteristics of multifractal scaling.

[Insert Figure 3 here]

A class of stochastic processes consistent with (2.2) is provided by the Multifractal Model of Asset Returns (MMAR), the first example of a multifractal diffusion ([**CFM**], [**CF02**]). This approach builds on multifractal measures ([**M74**]), which are constructed by the iterative random reallocation of mass within a time interval. One of the simplest examples is the binomial measure[2] on $[0, 1]$, which we derive as the limit of a *multiplicative cascade*. Consider the uniform probability measure μ_0 on the unit interval. Let M denote a binomial distribution taking the high value $m_0 \in [1/2, 1]$ or the low value $1-m_0$ with equal probability. In the first step of the cascade, we draw two independent values M_0 and M_1 from the binomial. We define a measure μ_1 by uniformly spreading the mass M_0 on the *left* subinterval $[0, 1/2]$, and the mass M_1 on the *right* subinterval $[1/2, 1]$. The density of μ_1 is a step function.

[2] The binomial is sometimes called the Bernoulli or Besicovitch measure.

In the second stage of the cascade, we draw four independent binomials $M_{0,0}$, $M_{0,1}$, $M_{1,0}$ and $M_{1,1}$. We split the interval $[0, 1/2]$ into two subintervals of equal length; the left subinterval $[0, 1/4]$ is allocated a fraction $M_{0,0}$ of $\mu_1[0, 1/2]$, while the right subinterval $[1/4, 1/2]$ receives a fraction $M_{0,1}$. Applying a similar procedure to $[1/2, 1]$, we obtain a measure μ_2 such that:

$$\begin{array}{ll} \mu_2[0,1/4] = M_0 M_{0,0}, & \mu_2[1/4,1/2] = M_0 M_{0,1}, \\ \mu_2[1/2,3/4] = M_1 M_{1,0}, & \mu_2[3/4,1] = M_1 M_{1,1}. \end{array}$$

Iteration of this procedure generates an infinite sequence of random measures (μ_k) that weakly converges to the binomial measure μ.

Consider the interval $[t, t+2^{-k}]$, where $t = \sum_{i=1}^{k} \eta_i 2^{-i}$ and $\eta_1, ..., \eta_k \in \{0, 1\}$, [3] and let φ_0 and φ_1 denote the relative frequencies of 0s and 1s in $(\eta_1, ..., \eta_k)$. The measure of the dyadic interval is then

$$\mu[t, t + 2^{-k}] = M_{\eta_1} M_{\eta_1,\eta_2} \dots M_{\eta_1,\dots,\eta_k} \Omega,$$

where Ω is a random variable determined by the change in the mass generated by stages $k+1, \dots, \infty$ of the cascade. Furthermore, we note that

$$\mathbb{E}(\mu[t, t + 2^{-k}]^q) = [\mathbb{E}(M^q)]^k = (\Delta t)^{\tau_\mu(q)+1},$$

where $\Delta t = 2^{-k}$ and $\tau(q) = -\log_2(E(M^q)) - 1$. The moments of the limiting measure of a dyadic interval of length Δt is therefore a power of Δt.

The MMAR extends multifractals from measures to diffusions. The asset price is specified by compounding a Brownian motion with an independent random time-deformation:

$$p(t) = p(0) + B[\theta(t)],$$

where θ is the cumulative distribution of the multifractal measure μ. The price process inherits the moment-scaling properties of the measure, in the sense that $\mathbb{E}(|p(t + \Delta t) - p(t)|^q) = (\Delta t)^{\tau_\mu(q/2)+1}$ on any dyadic interval $[t, t+\Delta t]$. These moment restrictions represent the basis of estimation and testing ([**CFM**], [**CF02**], [**L08**]). The MMAR provides a well-defined stochastic framework for the analysis of moment-scaling, which has generated extensive interest in econophysics (for example, [**LB**]). The multifractal model is also related to recent econometric research on power variation, which interprets return moments at various frequencies in the context of traditional jump-diffusions (for example, [**ABDL**], [**BNS**]). The MMAR also captures nonlinear changes in the return density with the time horizon ([**L01**]).

[3] A number $t \in [0, 1]$ is called *dyadic* if $t = 1$ or $t = \eta_1 2^{-1} + ... + \eta_k 2^{-k}$ for a finite k and $\eta_1, .., \eta_k \in \{0, 1\}$. A dyadic interval has dyadic endpoints.

Despite its appealing properties, the MMAR is unwieldy for econometric applications because of two features of the underlying measure: (a) the recursive reallocation of mass on an entire time-interval does not fit well with standard time series tools; and (b) the limiting measure contains a residual grid of instants that makes it non-stationary.

3. The Markov-Switching Multifractal (MSM)

The Markov-switching multifractal (MSM) resolves these difficulties by constructing a fully stationary volatility process that evolves stochastically through time ([**CF01**], [**CF04**], [**CF**]). MSM builds a bridge between multifractality and regime-switching, which permits the application of powerful statistical methods to a multifractal process.

3.1. Definition in Discrete Time. We assume that time is defined on the grid $t = 0, 1, \ldots, \infty$. We consider:

- a first-order Markov state vector $M_t = (M_{k,t})_{1 \leq k \leq \bar{k}} \in \mathbb{R}_+^{\bar{k}}$, and
- a distribution $M \geq 0$ with a unit mean: $\mathbb{E}(M) = 1$.

In this survey, we consider for expositional simplicity that the components $M_{1,t}, M_{2,t}, \ldots, M_{\bar{k},t}$ are independent across k. Each component $M_{k,t}$ is a Markov process in its own right, which is constructed as follows. Given $M_{k,t-1}$, the next-period component $M_{k,t}$ is

drawn from distribution M	with probability γ_k
set equal to its current level $M_{k,t-1}$	with probability $1 - \gamma_k$.

The construction generally requires no other assumptions about M. For simplicity, however, we will assume in the rest of this paper that M is a Bernoulli distribution taking values $m_0 \in [1, 2]$ or $2 - m_0 \in [0, 1]$ with equal probability.

The transition probabilities are tightly specified by

$$\gamma_k = 1 - (1 - \gamma_1)^{\left(b^{k-1}\right)}, \tag{3.1}$$

where $\gamma_1 \in (0, 1)$ and $b \in (1, \infty)$. The definition implies that $\gamma_1 <$ ... $< \gamma_{\bar{k}}$, so that components with a low index k are more persistent than higher-k components. If the parameter γ_1 is small compared to unity, the transition probability $\gamma_k \sim \gamma_1 b^{k-1}$ grows approximately at geometric rate b for low values of k. The rate of increase eventually slows down so that the parameter γ_k remains lower than unity.

We specify volatility a date t by:

$$\sigma(M_t) = \bar{\sigma} \left(\prod_{k=1}^{\bar{k}} M_{k,t} \right)^{1/2}, \tag{3.2}$$

where $\bar{\sigma} \in \mathbb{R}_{++}$ is a positive constant. We assume that returns $r_t = p_t - p_{t-1}$ are given by

$$r_t = \mu + \sigma(M_t)\varepsilon_t, \tag{3.3}$$

where $\mu \in \mathbb{R}$ and $\{\varepsilon_t\}$ are independent standard Gaussians. The parameter μ is the unconditional mean of the return process: $\mathbb{E}(r_t) = \mu$. Since $\mathbb{E}[(r_t - \mu)^2] = \bar{\sigma}^2$, the parameter $\bar{\sigma}$ is the unconditional standard deviation of the return r_t.

We call this construct the Markov-Switching Multifractal process, and denote by MSM($\overline{k}$) the version of the model with $\overline{k}$ frequencies. An MSM($\overline{k}$) process is fully parameterized by

$$\psi \equiv (m_0, \bar{\sigma}, b, \gamma_1, \mu) \in [1, 2] \times \mathbb{R}_{++} \times (1, +\infty) \times (0, 1) \times \mathbb{R},$$

where m_0 characterizes the distribution of the multipliers, $\bar{\sigma}$ is the unconditional standard deviation of returns, b and γ_1 define the set of switching probabilities, and μ the average return per period.

The multiplicative structure (3.3) is appealing to model the high variability and high volatility persistence exhibited by financial time series. The components have the same marginal distribution M but differ in their transition probabilities γ_k. When a low level multiplier changes, volatility varies discontinuously and has strong persistence. In addition, high frequency multipliers produce substantial outliers.

MSM parsimoniously specifies a high-dimensional state space. Since M is a binomial, the number of states is equal to $2^{\bar{k}}$. However, MSM is also parsimonious. In a general Markov chain, the size of the transition matrix is equal to the square of the number of states. For instance a general Markov chain with 2^8 states generally needs to be parameterized by $2^{16} = 65,536$ elements. In comparison, the MSM dynamics are fully characterized by five parameters.

Figure 4 illustrates the construction of binomial MSM. The top three panels represent the sample path of the volatility components $M_{k,t}$ for k varying from 1 to 3. We see that the number of switches tends to increase with k due to the geometric progression of arrival intensities (3.7). The fourth panel represents the variance $\sigma^2(M_t) \equiv \bar{\sigma}^2 \, M_{1,t}...M_{\bar{k},t}$, where $\bar{k} = 8$ and $\bar{\sigma} = 1$. The construction generates cycles of different frequencies, consistent with the economic intuition that there are volatile decades and less volatile decades, volatile years and less volatile years, and so on. The panel also shows pronounced peaks and intermittent bursts of volatility, which produce fat tails in returns. The last panel illustrates the impact of these various frequencies on the asset return series.

[Insert Figure 4 here.]

3.2. Filtering and Estimation. Since the distribution M is a binomial, the state vector M_t takes $d = 2^{\bar{k}}$ possible values $m^1, \ldots, m^d \in \mathbb{R}_+^{\bar{k}}$. Let $A = (a_{i,j})_{1 \le i,j \le d}$ denote the transition matrix with components[4] $a_{ij} = \mathbb{P}(M_{t+1} = m^j | M_t = m^i)$. Conditional on the volatility state, the return has Gaussian density $\omega_j(r_t) = n[(r_t - \mu)/\sigma(m^j)]/\sigma(m^j)$, where $n(\cdot)$ denotes the density of a standard normal.

The financial statistician observes the returns r_t but not the state vector M_t. She therefore computes the conditional probability distribution $\Pi_t = \left(\Pi_t^1, \ldots, \Pi_t^d\right) \in \mathbb{R}_+^d$, where for every $j \in \{1, \ldots, d\}$,

$$\Pi_t^j \equiv \mathbb{P}\left(M_t = m^j | r_1, \ldots, r_t\right).$$

The vector Π_t is computed recursively. Bayes' rule implies that

$$\Pi_t^j \propto f(r_t | M_t = m^j, r_1, \ldots, r_{t-1}) \mathbb{P}\left(M_t = m^j | r_1, \ldots, r_{t-1}\right), \tag{3.4}$$

or $\Pi_t^j \propto \omega_j(r_t) \sum_{i=1}^d a_{i,j} \mathbb{P}(M_{t-1} = m^i | r_1, \ldots, r_{t-1})$. The vector Π_t is therefore a function of its lagged value and the contemporaneous return r_t :

$$\Pi_t = \frac{\omega(r_t) \circ (\Pi_{t-1} A)}{[\omega(r_t) \circ (\Pi_{t-1} A)]\, \iota'}, \tag{3.5}$$

where $\omega(r_t) = [\omega_1(r_t), \ldots, \omega_d(r_t)]$, $\iota = (1, \ldots, 1) \in \mathbb{R}^d$, and $x \circ y$ denotes the Hadamard product $(x_1 y_1, \ldots, x_d y_d)$ for any $x, y \in \mathbb{R}^d$. These results are familiar in regime-switching models ([**H**]). In empirical applications, the initial vector Π_0 is chosen to be the ergodic distribution of the Markov process.

Let $L(r_1, ..., r_T; \psi)$ denote the probability density function of the time series $r_1, ..., r_T$ under the MSM model with parameter vector ψ. We easily check that:

$$\log L(r_1, ..., r_T; \psi) = \sum_{t=1}^{T} \log[\omega(r_t) \cdot (\Pi_{t-1} A)]. \tag{3.6}$$

For a fixed $\bar{k}$, the maximum likelihood estimator (ML),

$$\hat{\psi} = \operatorname{argmax}_{\psi} \log L(r_1, ..., r_T; \psi),$$

[4]We note that $a_{ij} = \prod_{k=1}^{\bar{k}} \left[(1 - \gamma_k) 1_{\{m_k^i = m_k^j\}} + \gamma_k \mathbb{P}(M = m_k^j) \right]$, where m_k^i denotes the mth component of vector m^i, and $1_{\{m_k^i = m_k^j\}}$ is the dummy variable equal to 1 if $m_k^i = m_k^j$, and 0 otherwise.

is consistent and asymptotically normal: $\sqrt{T}(\hat{\psi} - \psi) \xrightarrow{d} \mathcal{N}(0, V)$. Furthermore, ψ is asymptotically efficient, in the sense that no other estimator has a smaller asymptotic variance-covariance matrix V (see, e.g., [**C**]). The maximum likelihood estimator also performs well in finite samples ([**CF04**]). Filtering and parameter estimation are therefore remarkably convenient with MSM.

3.3. Empirical Estimation and Forecasting. We apply maximum likelihood estimation to the Yen / U.S. dollar exchange rate series shown in Figure 1. The daily logarithmic returns are calculated from exchange rates beginning in June 1973, extending to the end of our sample at the end of March, 2012.

Table 1 reports the maximum likelihood estimates. For convenience and to focus attention on volatility, we set the drift parameter $\mu = 0$. In Panel A, following [**CF04**] and [**CFT**], we estimate the four remaining parameters m_0, $\bar{\sigma}$, $\gamma_{\bar{k}}$, and b, for the number of frequencies $\bar{k}$ varying from 1 to 12. The first column is therefore a standard Markov-switching model with only two possible values for volatility. As $\bar{k}$ increases, the number of states increases at the rate $2^{\bar{k}}$. There are thus over four thousand states when $\bar{k} = 12$.

[Insert Table 1 here.]

The results in Panel A show that the multiplier parameter $\hat{m}_0$ tends to decline with $\bar{k}$: with a larger number of components, less variability is required in each M_{kt} to match the fluctuations in volatility exhibited by the data. The estimates $\hat{\sigma}$ vary across $\bar{k}$ with no particular pattern. Standard errors of $\hat{\sigma}$ increase with $\bar{k}$, consistent with the idea that long-run averages are difficult to identify in models permitting long volatility cycles. We next examine the frequency parameters $\hat{\gamma}_{\bar{k}}$ and $\hat{b}$. When $\bar{k} = 1$, the single multiplier has a duration of about one week. As $\bar{k}$ increases, the switching probability of the highest frequency multiplier increases until a switch occurs almost once a day for large $\bar{k}$. At the same time, the estimate $\hat{b}$ decreases steadily with $\bar{k}$. The increasing number of frequencies permits the range of frequencies to spread out to include both low- and high-frequency components, while the spacing between frequencies becomes tighter.

We finally examine the behavior of the log-likelihood function as the number of frequencies $\bar{k}$ increases from 1 to 12. The likelihood rises substantially as we add components for low $\bar{k}$, and continues to rise at a decreasing rate as we add components until approximately flattening in the range of $\bar{k} = 10$ to $\bar{k} = 12$ components. This behavior of the likelihood confirms one of the main premises of the multifractal approach: fluctuations in volatility occur with heterogeneous degrees of persistence, and explicitly incorporating a larger number of frequencies results in a better fit.

In Panel B, we restrict two of the MSM parameters. Consistent with the idea that the long-run mean of volatility is poorly identified, we set the unconditional volatility $\bar{\sigma}$ equal to the sample standard deviation of returns. Since the lowest-frequency volatility component is difficult to identify even in a long data sample, we set the value of γ_1 such that a switch in this component is expected to occur once in a sample four times the available sample size. With these restrictions, the model is fully identified by estimating only the two parameters m_0 and b. Empirically, these restrictions reduce the likelihood substantially when $\bar{k}$ is small, but for large values of $\bar{k}$ the restricted likelihood is very close to the unrestricted values shown in Panel A. These results suggest that restricting the values of σ and γ_1 can be a pragmatic empirical approach that further simplifies the estimation of the MSM model.

It is natural to compare the in-sample likelihood results of MSM with estimates for a standard volatility process. Generalized Auto-Regressive Conditional Heteroskedasticity ("GARCH", [**E82**], [**Bo**]) has the form $r_t = h_t^{1/2} e_t$, where h_t is the conditional variance of r_t at date $t-1$, and $\{e_t\}$ are i.i.d. Student innovations with unit variance and ν degrees of freedom. The standard GARCH(1,1) specification assumes the recursion $h_{t+1} = \omega + \alpha r_t^2 + \beta h_t$, and therefore has four parameters, ν, ω, α, and β. We estimate the GARCH model on the Yen / U.S. dollar exchange rate data, and find a likelihood of -8299.20, almost 100 points lower then MSM.

The MSM model also produces good out-of-sample forecasts. For both MSM and GARCH we estimate the models in-sample using data from the beginning of the sample until the end of 1995. We then use the data from the beginning of 1996 to the end of our data to evaluate out of sample performance. For each model we evaluate ability to forecast realized volatility $RV_{n,t} = \sum_{s=t-n+1}^{t} r_s^2$, for forecasting horizons ranging from 1 to 100 days. The out-of-sample forecasting R^2 for each model at each horizon is given by $R^2 = 1 - MSE/TSS$, where TSS is the out-of-sample variance of squared returns: $TSS = L^{-1} \sum_{t=T-L+1}^{T} (r_t^2 - \sum_{i=T-L+1}^{T} r_i^2/L)^2$, and the mean squared error (MSE) quantifies forecast errors in the out-of-sample period: $L^{-1} \sum_{t=T-L+1}^{T} [r_t^2 - \mathbb{E}_{t-1}(r_t^2)]^2$.

Table 2 reports summary forecasting results for horizons of 1, 5, 10, 20, 50, and 100 days. In addition to the Yen / dollar exchange rate, we consider three additional currencies.[5] The multifractal model shows robust good performance at all horizons and for all currencies, with particular strength appearing over longer horizons of 50 and 100 days. Additional evidence on the in- and out-of-sample performance of MSM in a variety of settings is provided in [**CF04**], [**CFT**], [**L08**], [**BKM**], [**CDS**], [**BSZ**], and [**I**].

[Insert Table 2 here.]

3.4. Long Memory in Volatility and Moment-Scaling. MSM generates a hyperbolic decline in autocovariances for a range of lags. Consider two arbitrary numbers α_1 and α_2 in the open interval $(0,1)$. The set of integers $I_{\bar{k}} = \{n : \alpha_1 \log_b(b^{\bar{k}}) \leq \log_b n \leq \alpha_2 \log_b(b^{\bar{k}})\}$ contains a large range of intermediate lags. We show in [**CF04**]:

Theorem 1 (Hyperbolic autocorrelation in volatility). *Consider a fixed vector ψ and let $q > 0$. The autocorrelation in levels satisfies*

$$\lim_{\bar{k}\to+\infty}\left(\sup_{n\in I_{\bar{k}}}\left|\frac{\log \rho_q(n)}{\log n^{-\delta(q)}} - 1\right|\right) = 0,$$

where $\delta(q) = \log_b \mathbb{E}(M^q) - 2\log_b \mathbb{E}(M^{q/2})$.

MSM mimics the hyperbolic autocovariograms $\log \rho_q(n) \sim -\delta(q)\log n$ exhibited by many financial series (e.g., [**D**], [**DGE**], [**BBM**]).

MSM illustrates that a Markov-chain regime-switching model can theoretically exhibit one of the defining features of long memory, a hyperbolic decline of the autocovariogram. Fractionally integrated processes generate such patterns by assuming that an innovation linearly affects future periods at a hyperbolically declining weight; as a result, fractional integration tends to produce smooth processes ([**MV68**], [**BBM**]). By contrast, MSM generates long cycles with a switching mechanism that also gives abrupt volatility changes. The combination of long-memory behavior with sudden volatility movements has a natural appeal for financial modeling.

MSM also replicates the moment-scaling properties of financial series. Intuitively, MSM is a randomized version of the MMAR, and

[5] The Deutschemark-Euro / U.S. dollar series is obtained by splicing the Deutschemark exchange rate with the Euro exchange rate, using the official Deutschemark / Euro exchange rate instituted at the end of 1998. We also consider the British pound / U.S. dollar exchange rate and the Canadian dollar / U.S. dollar exchange rate.

therefore inherits the moment-scaling properties of its precursor. We refer the reader to [**CF**] for theoretical results and simulations.

3.5. Continuous-Time MSM. The MSM construction works just as well in continuous time. We consider a filtration $\{\mathcal{F}_t\}$ on the probability space $(\Omega, \mathcal{F}, \mathbb{P})$. The Markov state vector

$$M_t = (M_{1,t}; M_{2,t}; \ldots; M_{\bar{k},t}) \in \mathbb{R}_+^{\bar{k}}$$

is now defined for all $t \in \mathbb{R}_+$. Given the Markov state M_t at date t, the dynamics over an infinitesimal interval are defined as follows. For each $k \in \{1, \ldots, \bar{k}\}$, a change in $M_{k,t}$ may be triggered by a Poisson arrival with intensity λ_k. The component $M_{k,t+dt}$ is drawn from a fixed distribution M if there is an arrival, and otherwise remains at its current value: $M_{k,t+dt} = M_{k,t}$. The construction can be summarized as:

$$\begin{array}{ll} M_{k,t+dt} \text{ drawn from distribution } M & \text{with probability } \lambda_k dt \\ M_{k,t+dt} = M_{k,t} & \text{with probability } 1 - \lambda_k dt. \end{array}$$

The Poisson arrivals and new draws from M are independent across k and t. The sample paths of a component $M_{k,t}$ are càdlàg, i.e. are right-continuous and have a limit point to the left of any instant.[6]

The arrival intensities are specified by

$$\lambda_k = \lambda_1 b^{k-1}, \qquad k \in \{1, \ldots, \bar{k}\}. \tag{3.7}$$

The parameter λ_1 determines the persistence of the lowest frequency component, and b the spacing between component frequencies.

The price process follows the Itô diffusion

$$\frac{dP_t}{P_t} = \mu dt + \sigma(M_t) dZ_t, \tag{3.8}$$

where $\sigma(M_t)$ follows the maintained equation (3.2).

3.6. Limiting Process with Countably Many Frequencies. The MSM construction can accommodate an infinity of frequencies, as we now show. For given parameters $(\mu, \bar{\sigma}, m_0, \lambda_1, b)$, let $M_t = (M_{k,t})_{k=1}^{\infty} \in \mathbb{R}_+^{\infty}$ denote an MSM Markov state process with countably many components. Each component $M_{k,t}$ is characterized by the arrival intensity $\lambda_k = \lambda_1 b^{k-1}$. For a finite $\bar{k}$, stochastic volatility is defined as the product of the first $\bar{k}$ components of the state vector: $\sigma_{\bar{k}}(M_t) \equiv \bar{\sigma} \left(\prod_{k=1}^{\bar{k}} M_{k,t} \right)^{1/2}$.

Since instantaneous volatility $\sigma_{\bar{k}}(M_t)$ depends on an increasing number of components, the differential representation (3.8) becomes unwieldy as $\bar{k} \to \infty$. In fact, the instantaneous volatility converges almost

[6]Càdlàg is a French acronym for *continue à droite, limite à gauche.*

surely to zero as $\bar{k} \to \infty$; since volatility is unbounded, however, the Lebesgue dominated convergence does not apply. We consider instead the time deformation

$$\theta_{\bar{k}}(t) \equiv \int_0^t \sigma_{\bar{k}}^2(M_s)ds. \tag{3.9}$$

Given a fixed instant t, the sequence $\{\theta_{\bar{k}}(t)\}_{\bar{k}=1}^{\infty}$ is a positive martingale with bounded expectation. By the martingale convergence theorem, the random variable $\theta_{\bar{k}}(t)$ converges to a limit distribution when $\bar{k} \to \infty$. A similar argument applies to any vector sequence $\{\theta_{\bar{k}}(t_1); \ldots; \theta_{\bar{k}}(t_d)\}$, guaranteeing that the stochastic process $\theta_{\bar{k}}$ has at most one limit point. As shown in [**CF01**], the sequence $\{\theta_{\bar{k}}\}_{\bar{k}}$ is tight under the following sufficient condition.[7]

CONDITION 1 (Tightness). $\mathbb{E}(M^2) < b$.

Intuitively, tightness prevents the time deformation from oscillating too wildly as $\bar{k} \to \infty$. Correspondingly, Condition 1 imposes that the volatility shocks are sufficiently small or have durations decreasing sufficiently fast to guarantee convergence.[8] Let $D[0,\infty)$ denote the space of càdlàg functions defined on $[0,\infty)$, and let d_∞° denote the Skohorod distance.

THEOREM 2 (Time deformation with countably many frequencies). *Under Condition 1, the sequence $(\theta_{\bar{k}})_{\bar{k}}$ weakly converges as $\bar{k} \to \infty$ to a measure θ_∞ defined on the metric space $(D[0,\infty), d_\infty^\circ)$. Furthermore, the sample paths of θ_∞ are continuous almost surely.*

The limiting process has a Markov structure analogous to MSM with a finite $\bar{k}$. In particular, $M_t = (M_{k,t})_{k=1}^{\infty}$ is the state vector of the limiting time deformation θ_∞.

The sample paths of the price process p are continuous but can be more irregular than a Brownian motion at some instants. Specifically, the local variability of a sample path at a given date t is characterized by the local Hölder exponent

$$\alpha(t) = \sup\{\beta \geq 0 : |p(t+\Delta t) - p(t)| = O(|\Delta t|^\beta) \text{ as } \Delta t \to 0\}.$$

---

[7] We refer the reader to [**Bi**] for a detailed exposition of weak convergence in function spaces.

[8] Because volatility exhibits increasingly extreme behavior as $\bar{k}$ goes up, the time deformation θ_∞ cannot be computed by taking the pointwise limit of the integrand $\sigma_{\bar{k}}^2(M_t)$ in equation (3.9). Specifically, $\sigma_{\bar{k}}^2(M_s)$ converges almost surely to zero as $\bar{k} \to \infty$ (by the Law of Large Numbers), suggesting that $\theta_\infty \equiv 0$. This conclusion would of course be misleading. For every fixed t, Condition 1 implies that $\sup_{\bar{k}} \mathbb{E}\left[\theta_{\bar{k}}^2(t)\right] < \infty$ ([**CF**]), and the sequence $\{\theta_{\bar{k}}(t)\}_{\bar{k}}$ is therefore uniformly integrable. Hence $\mathbb{E}\theta_\infty(t) = \mathbb{E}\theta_{\bar{k}}(t) = \bar{\sigma}^2\ t > 0$.

Heuristically, we can express the infinitesimal variations of the price process as being of order $(dt)^{\alpha(t)}$ around instant t. Lower values of $\alpha(t)$ correspond to more abrupt variations. Traditional jump diffusions impose that $\alpha(t)$ be equal to 0 at points of discontinuity, and to 1/2 otherwise. In a multifractal diffusion, however, the exponent $\alpha(t)$ takes a continuum of values in any time interval.

3.7. Extensions. MSM is a flexible framework that has been extended along several directions in recent years. ([**CFT**]) considers a multivariate model of asset returns, which captures the correlation both in levels and in volatility of the returns on several financial assets. Multivariate MSM preserves the tractability of univariate MSM, including analytical expressions for the likelihood function and the Bayesian filter. It captures well the joint dynamics of asset returns and provides accurate forecasts of the value at risk of a portfolio of assets.

[**I**] develops an extension of bivariate MSM that incorporates dynamic correlation in the Gaussian innovations. The new model, which the author coins MSMDCC, combines the multifrequency structure of bivariate MSM with the flexible correlation of Engle's Dynamic Conditional Correlation model [**E02**]. The likelihood and Bayesian filter of MSMDCC are available analytically. MSMDCC outperforms its two building blocks, MSM and DCC both in and out of sample.

[**CDS**] and [**BSZ**] introduce Markov-switching multifractal models of the time interval between two trades on a given security. Inter-trade durations are important in the microstructure literature and can help design trading strategies. The MSM duration models capture the key features of financial market inter-trade durations: long-memory dynamics and highly dispersed distributions. They also outperform their short-memory competitors in and out of sample.

4. Pricing Multifractal Risk

The integration of multifrequency models into asset pricing is now at the forefront of current research. We begin with an illustrative example drawn from [**CF08**].

4.1. A Continuous-Time Model of Stock Prices. The economy is specified by a standard Brownian motion $Z(t) \in \mathbb{R}$ and an MSM state vector $M_t \in \mathbb{R}_+^{\bar{k}}$, where $t \in [0, \infty)$ and $\bar{k}$ is a finite integer. The processes Z and M are mutually independent.

The stock is an asset that pays a continuous stream of dividend payments D_t. Let $\bar{g}_D \in \mathbb{R}_+$, and let $\sigma_D(M_t) = \bar{\sigma}_D(\prod_{k=1}^{\bar{k}} M_{k,t})^{1/2}$.

Condition 2 (Dividends). The dividend process satisfies

$$\log(D_t) \equiv \log(D_0) + \int_0^t \left[\bar{g}_D - \frac{\sigma_D^2(M_s)}{2}\right] ds + \int_0^t \sigma_D(M_s) dZ_D(s)$$

at every instant $t \in [0, \infty)$.

In the absence of arbitrage, the price of the stock at date t can be expressed as the present value of expected future dividends, where the expectation takes into account the risk aversion of investors ([**M**],[**DD**]).

The stock is priced by a collection of identical risk-averse agents, who observe the realization of the processes Z and M. Risk aversion is defined as follows. The agent ranks the desirability of a random consumption stream $\{C_t\}_{t \geq 0}$ according to the utility index

$$U(\{C_t\}) = \mathbb{E}\left[\int_0^{+\infty} e^{-\delta s} u(C_{t+s}) ds \,\middle|\, I_t\right],$$

where the discount rate δ is a strictly positive constant and the Bernoulli utility is:

$$u(C) \equiv \begin{cases} C^{1-\alpha}/(1-\alpha) & \text{if } \alpha \neq 1, \\ \log(C) & \text{if } \alpha = 1. \end{cases}$$

The agent strictly prefers the consumption stream $\{C_t\}$ to the consumption stream $\{C_t'\}$ if and only if $U(\{C_t\}) > U(\{C_t'\})$. We let $\rho = \delta - (1-\alpha)\bar{g}_D$, which we assume to be strictly positive. We use lower cases for the logarithms of all variables.

Theorem 3 (Equilibrium stock price). *The stock price is in logs the sum of the continuous dividend process and the price:dividend ratio:*

$$p_t = d_t + q(M_t),$$

where

$$q(M_t) = \log \mathbb{E}_t \left(\int_0^{+\infty} e^{-\rho s - \frac{\alpha(1-\alpha)}{2} \int_0^s \sigma_D^2(M_{t+h}) dh} ds \right). \tag{4.1}$$

The price process therefore follows a jump-diffusion. A price jump occurs when there is a discontinuous change in the Markov state M_t driving the continuous dividend process.

The endogenous price jumps contrast with the continuity of the dividend process. Over an infinitesimal time interval, the stock price changes by

$$d(p_t) = d(d_t) + \Delta(q_t),$$

where $\Delta(q_t) \equiv q_t - q_{t^-}$ denotes the finite variation of the price:dividend ratio in case of a discontinuous regime change. A Markov switch that increases the volatility of current and future dividends induces a negative realization of $\Delta(q_t)$ if $\alpha < 1$. Market pricing thus generates an endogenous negative correlation between volatility changes and price jumps.

The size of the jumps depends on the persistence of the component that changes. Low-frequency multipliers deliver persistent and discrete switches, which have a large price impact. By contrast, higher frequency components have no noticeable effect on prices, but give additional outliers in returns through their direct effect on the tails of the dividend process. The price process is thus characterized by a large number of small jumps (high frequency $M_{k,t}$), a moderate number of moderate jumps (intermediate frequency $M_{k,t}$), and a small number of very large jumps. Earlier empirical research suggests that this is a good characterization of the dynamics of stock returns. The model thus avoids the difficult choice of a unique frequency and size for rare events, which is a common issue in specifying traditional jump-diffusions.[9]

We illustrate in Figure 5 the endogenous multifrequency pricing dynamics of the model, in the case where consumption is IID. The top two panels present a simulated dividend process, in growth rates and in logarithms of the level respectively. The middle two panels then display the corresponding stock returns and log prices. The price series exhibits much larger movements than dividends, due to the presence of endogenous jumps in the price-dividend ratio, $e^{q(M_t)}$. To see this clearly, the bottom two panels show consecutively: 1) the "feedback" effects, defined as the difference between log stock returns and log dividend growth, and 2) the price:dividend ratio. Consistent with Theorem 3, we observe a few infrequent but large jumps in prices, with smaller but more numerous small discontinuities. The simulation demonstrates that the difference between stock returns and dividend growth can be large even when the price-dividend ratio varies in a plausible and relatively modest range (between 26 and 33 in the figure). The pricing model thus captures multifrequency stochastic volatility, endogenous

[9]In the simplest exogenously specified jump-diffusions, it is often possible that discontinuities of heterogeneous but fixed sizes and different frequencies can be aggregated into a single collective jump process with an intensity equal to the sum of all the individual jumps, and a random distribution of sizes. A comparable analogy can be made for the state vector M_t in our model, but due to the equilibrium linkages between jump size and the duration of volatility shocks, and the state dependence of price jumps, no such reduction to a single aggregated frequency is possible for the equilibrium stock price.

multifrequency jumps in prices, and endogenous correlation between volatility and price innovations.

[Insert Figure 5 here.]

4.2. Convergence to a Multifractal Jump-Diffusion. We now investigate how the price diffusion evolves as $\bar{k} \to \infty$, i.e. as components of increasingly high frequency are added into the state vector. By Theorem 2, $d_{\bar{k}}(t)$ converges to

$$d_\infty(t) \equiv d_0 + \bar{g}_D t - \theta_\infty(t)/2 + B[\theta_\infty(t)]$$

as $\bar{k} \to \infty$. By (4.1), the process $q_{\bar{k}}(t)$ is a positive submartingale, which also converges to a limit as $\bar{k} \to \infty$.

Theorem 4 (Jump-diffusion with countably many frequencies). *We assume that $\alpha < 1$ and that the maintained conditions 1 and 2 hold. When the number of frequencies goes to infinity, the log-price process weakly converges to*

$$p_\infty(t) \equiv d_\infty(t) + q_\infty(t),$$

where

$$q_\infty(t) = \log \mathbb{E}\left[\int_0^{+\infty} e^{-\rho s - \frac{\alpha(1-\alpha)}{2}[\theta_\infty(t+s)-\theta_\infty(t)]} ds \,\middle|\, (M_{k,t})_{k=1}^{\infty} \right]$$

is a pure jump process. The limiting price is thus a jump diffusion with countably many frequencies.

In an economy with countably many frequencies, the log-price process is the sum of: (i) the continuous multifractal diffusion $d_\infty(t)$; and (ii) the pure jump process $q_\infty(t)$. We correspondingly call $p_\infty(t)$ a *multifractal jump-diffusion.*

When $\bar{k} = \infty$, the state space is a continuum while the multifractal jump-diffusion is tightly specified by the seven parameters $(\bar{g}_D, \bar{\sigma}_D, m_0, \gamma_1, b, \alpha, \delta)$. The limiting process $q_\infty(t)$ exhibits rich dynamic properties. Within any bounded time interval, there exists almost surely (a.s.) at least one multiplier $M_{k,t}$ that switches and triggers a jump in the stock price. A jump in price occurs a.s. in the neighborhood of any instant. The number of switches is also countable a.s. within any bounded time interval, implying that the process $q_\infty(t)$ has infinite activity and is continuous almost everywhere.

The convergence results provide useful guidance on the choice of the number of frequencies in theoretical and empirical applications. On the one hand, the convergence of the price process implies that the marginal contribution of additional components is likely to be small in applications concerned with fitting the price or return series. It is then convenient to consider a number of frequencies $\bar{k}$ that is sufficiently large to capture the heteroskedasticity of financial series, but sufficiently small to remain tractable. On the other hand, countably many frequencies might prove useful in more theoretical contexts, in which the local behavior of the price process needs to be carefully understood. Examples could include the construction of learning models or the design of dynamic hedging strategies.

4.3. Other Work. Several other papers derive the pricing implications of multifractal risk. [**CF07**] develops a discrete-time model of stock returns in which the volatility of dividend news follows an MSM. The resulting variance of stock returns is substantially higher than the variance of dividends, as is the case with the data. The MSM dividend specification improves on the classic [**CH**] model, which generates more modest amplification effects with a GARCH dividend process. [**CF07**] also investigates the dynamics of returns when the agent is not fully informed about the state M_t but must sequentially learn about it from dividends and other signals; the implied return process exhibits substantial negative skewness, which is again consistent with the data.

Multifractal modeling is also helpful for option pricing. [**CFFL**] introduces an extension of MSM that can account for the variation in skewness and term-structure of option data. Jumps to the return process are triggered by changes in lower-frequency volatility components, and the "leverage effect" is generated by a negative correlation of high-frequency innovations to returns and volatility. Using S&P 500 index returns and a panel of options with multiple maturities and strikes, the latent volatility components enable the model to dynamically fit a wide range of option surfaces.

Parsimonious models with multiple components have a natural use in interest rate modeling. [**CFW**]) develops a class of dynamic term structure models that accommodates arbitrarily many interest-rate factors with a fixed number of parameters. The approach builds on a short-rate cascade, a parsimonious recursive construction that ranks the state variables by their rates of mean reversion, each revolving around the preceding lower-frequency factor. The cascade accommodates a wide range of volatility and risk premium specifications, and the forward curves implied by absence of arbitrage are smooth, dynamically consistent, and available in closed form. [**CFW**] provides conditions under which as the number of factors goes to infinity the construction converges to a well-defined, infinite-dimensional dynamic term structure. The cascade thus overcomes the curse of dimensionality associated with general affine models. Using a panel of 15 LIBOR and swap rates, [**CFW**] estimates specifications with a number of factors ranging from one to 15, all specified by only five parameters. In sample, the implied yield curve fits the data almost perfectly. Out of sample, interest rate forecasts significantly outperform prior benchmarks.

Overall, the results presented in this section show that multifractal risk has rich pricing implications that have already allowed researchers to overcome key shortcomings of standard financial models. These early successes suggest that multifractals are promising powerful tools for asset pricing.

5. Conclusion

Fifty years ago, Benoît Mandelbrot discovered that financial returns exhibit strong departures from Gaussianity and advocated the use of self-similar Lévy-stable processes for modeling market fluctuations. These two insights sparked the introduction of fractal methods in finance. Since then, fat tails, fractional integration, and multifractal scaling have become familiar tools to financial practitioners, econometricians, statisticians and econophysicists. Fractal methods are now routinely combined with more traditional approaches, and have given rise to popular hybrid models such as fractionally integrated GARCH ([**BBM**]) or long-memory stochastic volatility ([**HMS**]). These advances are testimony to the successful integration of fractal methods into mainstream finance.

In the past fifteen years, fractal research in finance has centered on the development of multifractal models of returns, which can jointly capture fat tails, long-memory volatility persistence, multifractal moment scaling, and nonlinear changes in the distribution of returns observed over various horizons. Multifractal models capture these empirical regularities with a remarkably small number of parameters, and as a result are strong performers both in- and out-of-sample, as the empirical section of this article illustrates.

These developments in financial research have led to advances in multifractal methodology itself. Multifractal measures can now be constructed dynamically through time ([**CF01**], [**BDM**]), and several classes of multifractal diffusions are now available ([**CFM**], [**CF01**], [**BDM**]). These innovations provide new intuitions about the emergence of multifractal behavior in economic and natural phenomena. For instance, MSM shows that multifractality can be generated by a Markov process with multiple components, each of which has its own degree of persistence. MSM also permits the application of efficient statistical methods, such as likelihood estimation and Bayesian filtering, to a multifractal process. These developments are new to the multifractal literature and are now spreading outside the field of finance (e.g., [**RR**]). Furthermore, incorporating multifractal risk into a pricing model generates multifractal jump-diffusions, an entirely new mathematical object that deserves further investigation.

Despite these successes, multifractal finance remains a young and under-researched field, and many challenges remain. The statistical methodology can be improved to incorporate finer features of financial returns, for instance along the lines of [**CFFL**]. Improvements in statistical inference are undoubtedly possible, for instance by using different distributions M, by exploring different transition probability specifications or by simplifying the estimation method. The integration of fractal risk into pricing models seems very powerful, as illustrated by recent work on option pricing and the term structure of interest rates. The applications of multifractal techniques to finance nonetheless remains in their infancy. Last but not least, the multifractal behavior of financial returns remains unexplained and invites the financial theorist to explore the economic mechanisms producing self-similar regime-switching in financial volatility.

References

[AB] V. Akgiray and G. G. Booth, *The stable-law model of stock returns*, Journal of Business and Economic Statistics **6** (1988), 51–57.

[ABDL] T. Andersen, T. Bollerslev, F. X. Diebold, and P. Labys, *The distribution of realized exchange rate volatility*, Journal of the American Statistical Association **96** (2001), 42–55.

[Ba] L. Bachelier, *Théorie de la spéculation*, Annales Scientifiques de l'Ecole Normale Supérieure **17** (1900), 21–86.

[BDM] E. Bacry, J. Delour, and J.-F. Muzy, *Multifractal random walks*, Physical Review E **64** (2001), 026103–06.

[BKM] E. Bacry, A. Kozhemyak, and J.-F. Muzy, *Continuous cascade models for asset returns,* Journal of Economic Dynamics and Control 32 (2008), 156–99.

[B] R. Baillie, *Long memory processes and fractional integration in econometrics*, Journal of Econometrics **73** (1996), 5–59.

[BBM] R. Baillie, T. Bollerslev, and H. O. Mikkelsen, *Fractionally integrated generalized autoregressive conditional heteroscedasticity*, Journal of Econometrics **74** (1996), 3–30.

[BNS] O. Barndorff-Nielsen and N. Shephard, *Power and bipower variation with stochastic volatility and jumps*, Journal of Financial Econometrics **2** (2004), 1–37.

[B06] R. Barro, *Rare disasters and asset markets in the twentieth century*, Quarterly Journal of Economics **121** (2006), 823–866.

[BSZ] J. Baruník, N. Shenaiz, and F. Žikeš, *Modeling and forecasting persistent financial durations,* Working paper, Academy of Sciences of the Czech Republic and Imperial College London Business School.

[Bi] P. Billingsley, *Convergence of Probability Measures*, 2nd edition, New York: Wiley, 1999.

[BS] F. Black and M. Scholes, *The pricing of options and corporate liabilities*, Journal of Political Economy **81** (1973), 637–654.

[BG] R. C. Blattberg and N. J. Gonedes, *A comparison of the Student and stable distributions as statistical models of stock prices*, Journal of Business **47** (1974), 244–80.

[Bo] T. Bollerslev, *Generalized autoregressive conditional heteroskedasticity*, Journal of Econometrics **31** (1986), 307–327.

[CF01] L. E. Calvet and A. J. Fisher, *Forecasting multifractal volatility*, Journal of Econometrics **105** (2001), 27–58.

[CF02] ______, *Multifractality in asset returns: theory and evidence*, Review of Economics and Statistics **84** (2002a), 381–406.

[CF04] ______, *How to forecast long-run volatility: Regime-switching and the estimation of multifractal processes*, Journal of Financial Econometrics **2** (2004), 49–83.

[CF07] ______, *Multifrequency news and stock returns*, Journal of Financial Economics **86** (2007), 178–212.

[CF08] ______, *Multifrequency jump-diffusions: an equilibrium approach*, Journal of Mathematical Economics **44** (2008), 207–26.

[CF] ______, *Multifractal volatility: Theory, forecasting, and pricing*, Burlington, MA, Academic Press, 2008.

[CFFL] L. E. Calvet, M. Fearnley, A. J. Fisher, and M. Leippold, *Accurate option pricing with skew-adjusted multifractal volatility,* Working paper, HEC Paris, University of British Columbia, and University of Zürich, 2012.

[CFM] L. E. Calvet, A. J. Fisher, and B. Mandelbrot, *A multifractal model of asset returns*, Cowles Foundation Discussion Papers 1164–1166, Yale University, 1997, available at http://www.ssrn.com.

[CFT] L. E. Calvet, A. J. Fisher, and S. B. Thompson, *Volatility comovement: a multifrequency approach*, Journal of Econometrics **131** (2006), 179–215.

[CFW] L. E. Calvet, A. J. Fisher, and L. Wu, *Dimension-invariant dynamic term structures*, Working paper, HEC Paris, University of British Columbia and Baruch College, 2011.

[CH] J. Y. Campbell and L. Hentschel, *No news is good news: an asymmetric model of changing volatility in stock returns*, Journal of Financial Economics **31** (1992), 281–318.

[CV] J. Y. Campbell and L. Viceira, *Strategic Asset Allocation*, Oxford University Press, 2002.

[CDS] F. Chen, F. X. Diebold, and F. Schorfheide, *A Markov-Switching Multi-Fractal Inter-Trade Duration Model, with Application to U.S. Equities,* manuscript, Huazhong University and University of Pennsylvania, 2012.

[C] H. Cramér, *Mathematical Method of Statistics,* Princeton University Press, 1946.

[D] M. Dacorogna, U. Müller, R. Nagler, R. Olsen, and O. Pictet, *A geographical model for the daily and weekly seasonal volatility in the foreign exchange market*, Journal of International Money and Finance **12** (1993), 413–38.

[DGE] Z. Ding, C. Granger, and R. Engle, *A long memory property of stock returns and a new model*, Journal of Empirical Finance **1** (1993), 83–106.

[DD] D. Duffie, *Dynamic Asset Pricing Theory*, Princeton University Press, Third Edition, 2001.

[E82] R. Engle, *Autoregressive conditional heteroskedasticity with estimates of the variance of UK inflation*, Econometrica **50** (1982), 987–1008.

[E02] R. Engle, *Dynamic conditional correlation - A simple class of multivariate GARCH models,* Journal of Business and Economic Statistics 20 (2002), 339–350

[FR] B. D. Fielitz and J. P. Rozelle, *Stable distributions and mixtures of stable distributions hypotheses for common stock returns*, Journal of the American Statistical Association 78 (1982), 28–36.

[G09] X. Gabaix, *Power laws in economics and finance*, Annual Review of Economics and Finance **1** (2009), 255–294.

[G12] X. Gabaix, *An exactly solved framework for ten puzzles in macro-finance*, Quarterly Journal of Economics **127** (2012), 645–700.

[G] S. Ghashgaie, W. Breymann, J. Peinke, P. Talkner, and Y. Dodge, *Turbulent cascades in foreign exchange markets*, Nature **381** (1996), 767–70.

[GJ] C. Granger and R. Joyeux, *An introduction to long memory time series models and fractional differencing*, Journal of Time Series Analysis **1** (1980), 15–29.

[H] J. D. Hamilton, *A new approach to the economic analysis of nonstationary time series and the business cycle*, Econometrica **57** (1989), 357–84.

[Hi] B. M. Hill, *A simple general approach to inference about the tail of a distribution*, Annals of Statistics **3** (1975), 1163–74.

[HMS] C. M. Hurvich, E. Moulines, and P. Soulier, *Estimating long memory in volatility,* Econometrica 73 (2005), 1283328.

[I] J. Idier, *Correlation: A Markov-switching multifractal model with time-varying correlations*, Working paper, Bank of France, 2010.

[JD] D. W. Jansen and C. G. DeVries, *On the frequency of large stock returns: Putting booms and busts into perspective*, Review of Economics and Statistics **73** (1991), 18–24.

[KK] K. G. Koedijk and C. J. M. Kool, *Tail estimates of East European exchange rates,* Journal of Business and Economic Statistics **10** (1992), 83–96.

[KSV] K. G. Koedijk, M. M. A. Schafgans and C. G. deVries, *The tail index of exchange rate returns,* Journal of International Economics **29** (1990), 93–108.

[K] A. Kolmogorov, *Wienersche Spiralen und einige andere interessante Kurven im Hilbertschen raum*, Comptes Rendus de l'Académie des Sciences de l'URSS **26** (1940), 115–18.

[LB] B. LeBaron, *Stochastic volatility as a simple generator of apparent financial power laws and long memory*, Quantitative Finance **1** (2001), 621–31.

[L24] P. Lévy, *Théorie des erreurs: la loi de Gauss et les lois exceptionnelles*, Bulletin de la Société Mathématique de France **52** (1924), 49–85.

[LP] M. Loretan and P. C. B. Phillips, *Testing covariance stationarity of heavy-tailed time series: An overview of the theory with applications to several financial datasets,* Journal of Empirical Finance **1** (1994), 211–248 .

[L01] T. Lux, *Turbulence in financial markets: the surprising explanatory power of simple cascade models*, Quantitative Finance **1** (2001), 632–40.

[L08] T. Lux, *The Markov-switching multifractal model of asset returns: GMM estimation and linear forecasting of volatility*, Journal of Business and Economic Statistics **26** (2008), 194–210.

[MS] S. Maheswaran and C. Sims, *Empirical implications of arbitrage-free asset markets*, in Models, Methods and Applications of Econometrics, ed. P. Phillips, Basil Blackwell, Oxford, 1993.

[M63] B. B. Mandelbrot, *The variation of certain speculative prices*, Journal of Business **36** (1963), 394–419.

[M65] ———, *Une classe de processus stochastiques homothétiques à soi*, Comptes Rendus de l'Académie des Sciences de Paris **260** (1965), 3274–77.

[M67] ———, *The variation of some other speculative prices*, Journal of Business **40** (1967), 393–413.

[M74] ———, *Intermittent turbulence in self-similar cascades: divergence of high moments and dimension of the carrier*, Journal of Fluid Mechanics **62** (1974), 331–58.

[M82] ———, *The Fractal Geometry of Nature*, Freeman, New York, NY, 1982.

[M97] ———, *Fractals and scaling in finance*, New York, NY, Springer, 1997.

[MV68] B. B. Mandelbrot and J. Van Ness, *Fractional Brownian motion, fractional noises and applications*, SIAM Review **10** (1968), 422–37.

[MK] H. Markowitz, *Portfolio selection*, Journal of Finance **7** (1952), 77–91.

[M] R. C. Merton, *Continuous-Time Finance*, Basil Blackwell, Oxford, 1990.

[PMM] P. C. B. Phillips, J. W. McFarland, and P. C. McMahon, *Robust tests of forward exchange market efficiency with empirical evidence from the 1920s*, Journal of Applied Econometrics **11** (1996), 1–22.

[R] J. Regnault, *Calcul des Chances et Philosophie de la Bourse*, Mallet-Bachelier and Castel, Paris, 1863.

[R88] T. A. Rietz, *The equity risk premium: A solution,* Journal of Monetary Economics **22** (1988), 117–131.

[RR] M. Rypdal and K. Rypdal, *Discerning a linkage between solar wind turbulence and ionospheric dissipation by a method of conned multifractal motions*, Journal of Geophysical Research **116**, A02202.

[ST] G. Samorodnitsky and M. Taqqu, *Stable Non-Gaussian Random Processes*, Chapman and Hall, 1994.

[S] W. Sharpe, *Capital asset prices: A theory of market equilibrium under conditions of risk*, Journal of Finance **19** (1964), 425–42.

[T] J. Tobin, *Liquidity preference as behavior towards risk*, Review of Economic Studies **25** (1958), 68–85.

[VA] N. Vandewalle and M. Ausloos, *Multi-affine analysis of typical currency exchange rates*, European Physical Journal B **4** (1998), 257–61.

[W] J. Wachter, *Can time-varying risk of rare disasters explain aggregate stock market volatility?*, Journal of Finance (2012), forthcoming.

Department of Finance, HEC Paris, 1 avenue de la Libération, 78351 Jouy-en-Josas, France
E-mail address: calvet@hec.fr

Department of Finance, Sauder School of Business, Vancouver, BC, V6T 1Z2, Canada
E-mail address: adlai.fisher@sauder.ubc.ca

TABLE 1. – Maximum Likelihood Results

	$\bar{k}=1$	2	3	4	5	6	7	8	9	10	11	12
						Japanese Yen / US Dollar						
						A. Four Parameters						
$\hat{m}_0$	1.732	1.730	1.663	1.625	1.563	1.541	1.493	1.488	1.435	1.405	1.400	1.376
	(0.012)	(0.014)	(0.010)	(0.011)	(0.009)	(0.012)	(0.009)	(0.010)	(0.010)	(0.009)	(0.009)	(0.009)
$\hat{\sigma}$	0.658	0.563	0.564	0.483	0.493	0.592	0.599	0.497	0.519	0.516	0.448	0.442
	(0.009)	(0.014)	(0.015)	(0.024)	(0.016)	(0.025)	(0.020)	(0.016)	(0.022)	(0.009)	(0.012)	(0.024)
$\hat{\gamma}_{\bar{k}}$	0.192	0.352	0.309	0.719	0.856	0.932	0.992	0.989	1.000	1.000	1.000	1.000
	(0.022)	(0.041)	(0.059)	(0.080)	(0.053)	(0.050)	(0.015)	(0.013)	(0.000)	(0.000)	(0.000)	(0.000)
$\hat{b}$	-	53.96	13.57	16.81	11.14	8.86	7.07	6.48	4.80	4.03	3.81	3.20
		(23.65)	(2.08)	(3.34)	(1.39)	(0.91)	(0.80)	(0.60)	(0.48)	(0.46)	(0.02)	(0.67)
ln L	-8887.13	-8520.07	-8339.72	-8269.27	-8233.90	-8217.42	-8208.84	-8203.64	-8200.27	-8199.34	-8197.46	-8196.81
						B. Two Parameters						
$\hat{m}_0$	-	1.716	1.710	1.640	1.612	1.552	1.506	1.471	1.453	1.427	1.403	1.382
		(0.008)	(0.012)	(0.010)	(0.008)	(0.009)	(0.008)	(0.008)	(0.010)	(0.009)	(0.008)	(0.008)
$\hat{b}$	-	5460	86.84	42.32	15.40	10.12	7.73	6.53	4.43	3.94	3.61	3.33
		(522)	(7.97)	(3.88)	(0.67)	(0.42)	(0.38)	(0.31)	(0.17)	(0.13)	(0.11)	(0.11)
ln L		-8573.37	-8420.14	-8308.86	-8236.95	-8222.57	-8213.22	-8214.04	-8208.38	-8202.43	-8200.51	-8199.23

Notes: This table shows maximum likelihood estimation results for the binomial MSM model and Japanese yen / U.S. dollar data. Panel A shows unrestricted estimation of the four parameters m_0, σ, $\gamma_{\bar{k}}$, and b. In Panel B, we set σ equal to the sample standard deviation, and γ_1 equal to the inverse of the sample size divided by four, which corresponds to a lowest frequency arrival rate that would occur on average once in a period four times the sample size. Columns correspond to the number of frequencies $\bar{k}$ in the estimated model. Asymptotic standard errors are in parenthesis.

TABLE 2. – Forecast Summary, Multiple Horizons

	Horizon (Days)					
	1	5	10	20	50	100
	Restricted R^2					
	Deutsche Mark/U.S. Dollar					
Binomial MSM	0.036	0.205	0.298	0.347	0.280	0.088
GARCH	0.045	0.197	0.285	0.328	0.174	-0.396
	Japanese Yen/U.S. Dollar					
Binomial MSM	0.052	0.120	0.166	0.206	0.166	0.103
GARCH	0.051	0.094	0.101	0.071	-0.172	-0.384
	British Pound/U.S. Dollar					
Binomial MSM	0.085	0.352	0.414	0.418	0.343	0.181
GARCH	0.117	0.410	0.485	0.489	0.452	0.118
	Canadian Dollar/U.S. Dollar					
Binomial MSM	0.100	0.270	0.324	0.316	0.257	0.219
GARCH	0.142	0.430	0.574	0.574	0.378	0.170

Notes: This table summarizes out-of-sample forecasting performance for MSM and GARCH across multiple forecasting horizons. We estimate the models in-sample using data from the beginning of the sample until the end of 1995. We then use the data from the beginning of 1996 to the end of our data to evaluate out of sample performance. For each model we evaluate ability to forecast realized volatility $RV_{n,t} = \sum_{s=t-n+1}^{t} r_s^2$, for forecasting horizons n ranging from 1 to 100 days. The out-of-sample forecasting R^2 for each model at each horizon is given by $R^2 = 1 - MSE/TSS$, where TSS is the out-of-sample variance of squared returns: $TSS = L^{-1}\sum_{t=T-L+1}^{T}(r_t^2 - \sum_{i=T-L+1}^{T} r_i^2/L)^2$, and the mean squared error (MSE) quantifies forecast errors in the out-of-sample period: $L^{-1}\sum_{t=T-L+1}^{T}[r_t^2 - \mathbb{E}_{t-1}(r_t^2)]^2$.

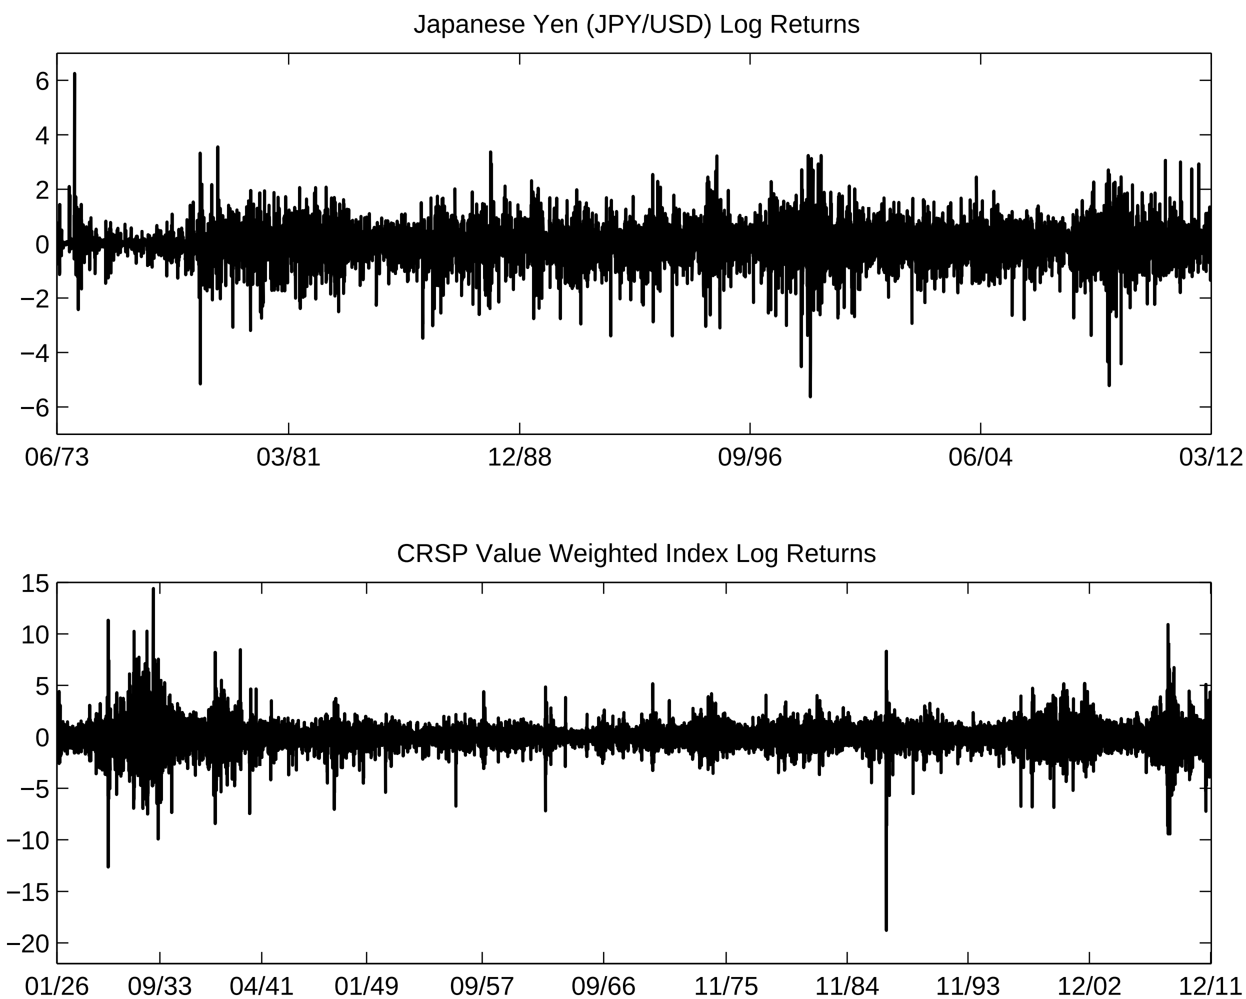

Figure 1: Financial Return Series. This figure shows daily logarithmic returns for the Japanese yen / U.S. dollar exchange rate series, and for the value weighted U.S. stock index compiled by the Center for Research in Securities Prices (CRSP). The yen series begins in June 1, 1973 and ends on March 30, 2012, and has 9751 return observations. The stock index return series begins on January 2, 1926 and ends on December 30, 2011 and has 22,780 return observations.

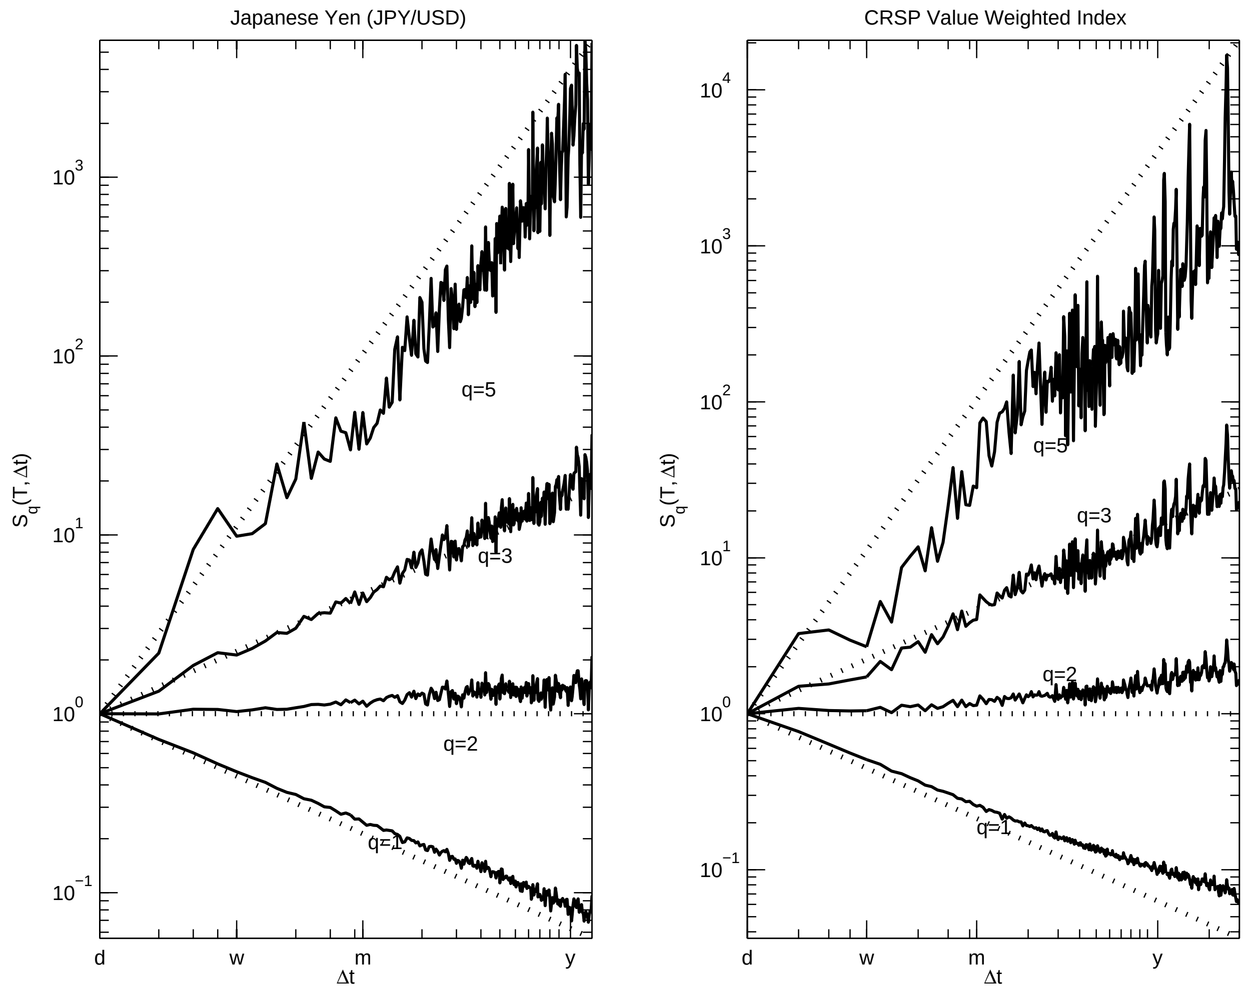

Figure 2: Moment Scaling of Financial Returns. This figure shows moment scaling in financial asset returns. Let $X(t)$ denote the cumulative logarithmic returns of a financial series. Partitioning $[0, T]$ into integer N intervals of length $\triangle t$, we define the partition function $S_q(T, \triangle t) \equiv \sum_{i=0}^{N-1} |X(i\triangle t + \triangle t) - X(i\triangle t)|^q$. When $X(t)$ is multifractal and has a finite q^{th} moment, then $\mathbb{E}\left[|X(\triangle t)|^q\right] = c_X(q)(\triangle t)^{\tau_X(q)+1}$, or equivalently $\ln \mathbb{E}[S_q(T, \triangle t)] = \tau_X(q) \ln(\triangle t) + c^*(q)$ where $c^*(q) = \ln c_X(q) + \ln T$. The figure provides log-log plots of $\triangle t$ against $S_q(T, \triangle t)$. When the data-generating process is multifractal these plots should be approximately linear.

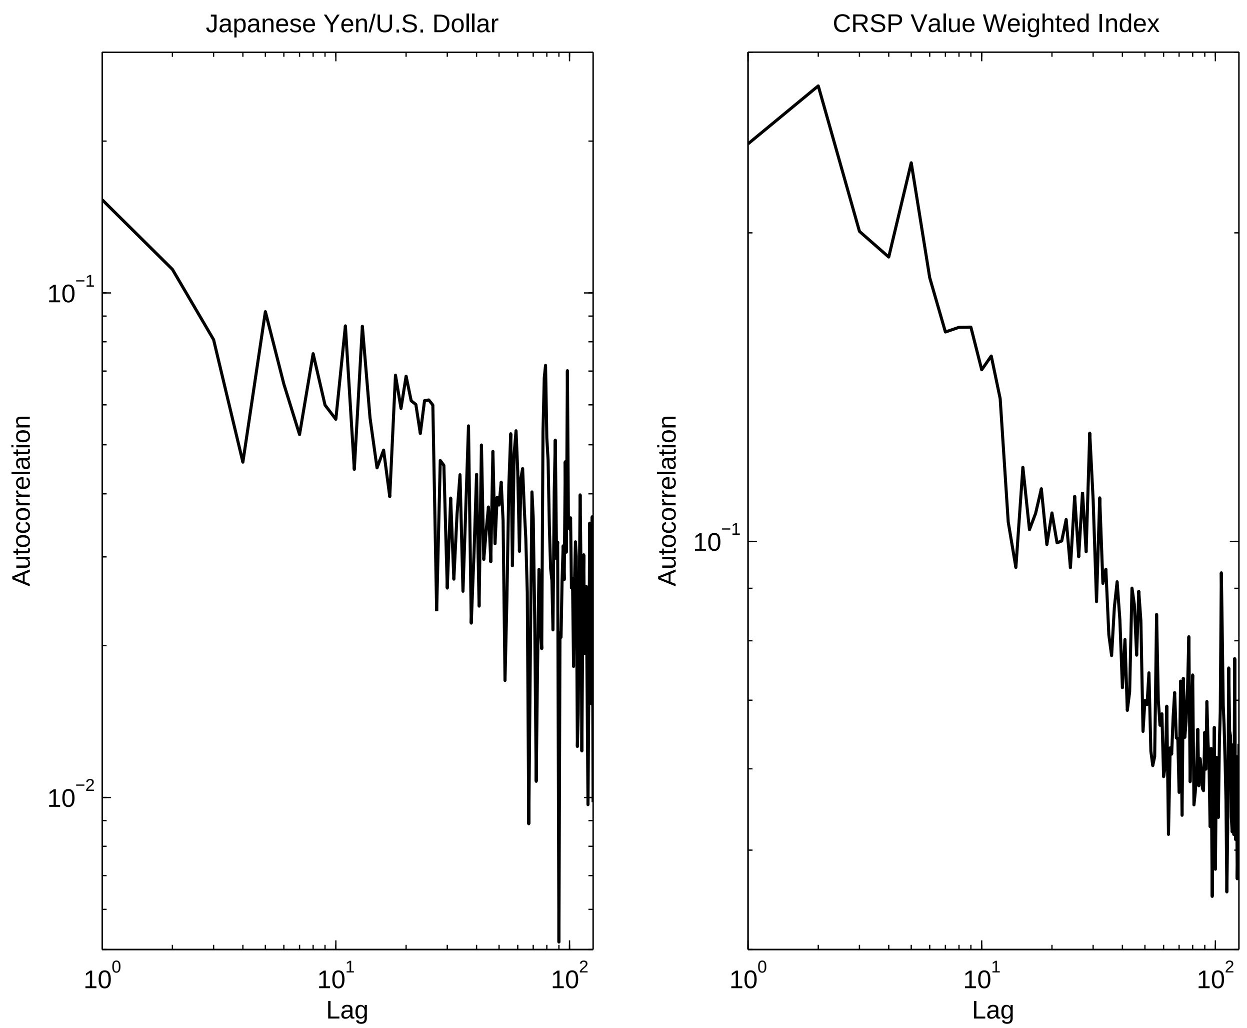

Figure 3: Long Memory in Financial Return Volatility. This figure shows the autocorrelations of squared logarithmic returns for the yen series (Panel A) and the CRSP stock index return series (Panel B). The autocorrelations are plotted on a log-log scale, so that a hyperbolic decay in autocorrelations, as occurs under long-memory, will appear as an approximately straight line in the figure.

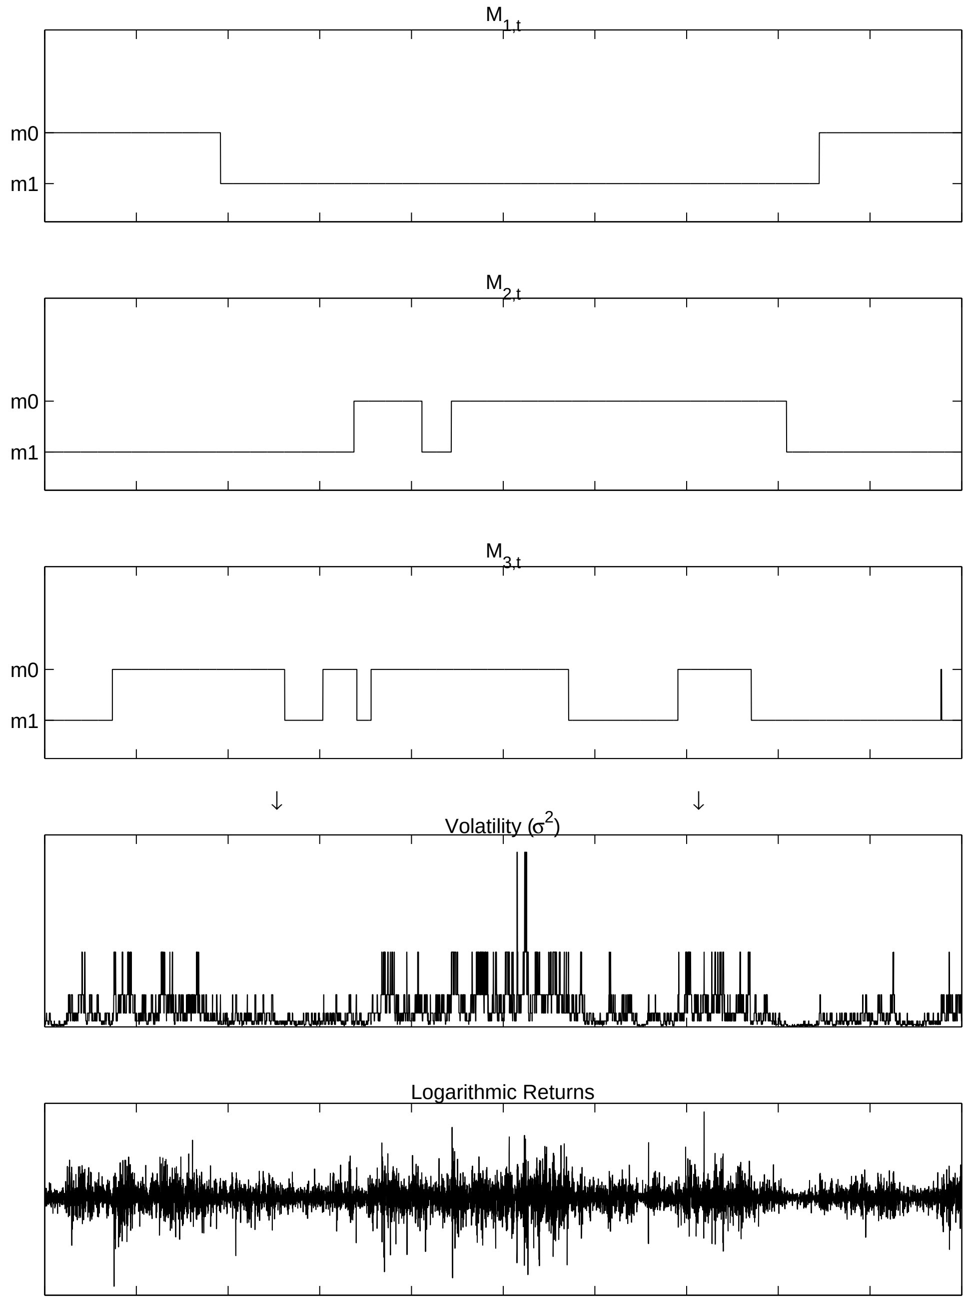

Figure 4: Construction of the Markov-Switching Multifractal. This figure shows the construction of multifractal volatility. The first three panels show the values of the three lowest-frequency volatility components $M_{1,t}$, $M_{2,t}$, and $M_{3,t}$. The fourth panel shows the variance σ_t^2, which is the product of all multipliers. The construction uses $\bar{k} = 8$ multipliers. The final panel shows the simulated logarithmic return series. The $T = 10,000$ period simulation uses the binomial MSM construction with $m1 = 1.4$, $b = 2$, and $\gamma_1 = b/T$.

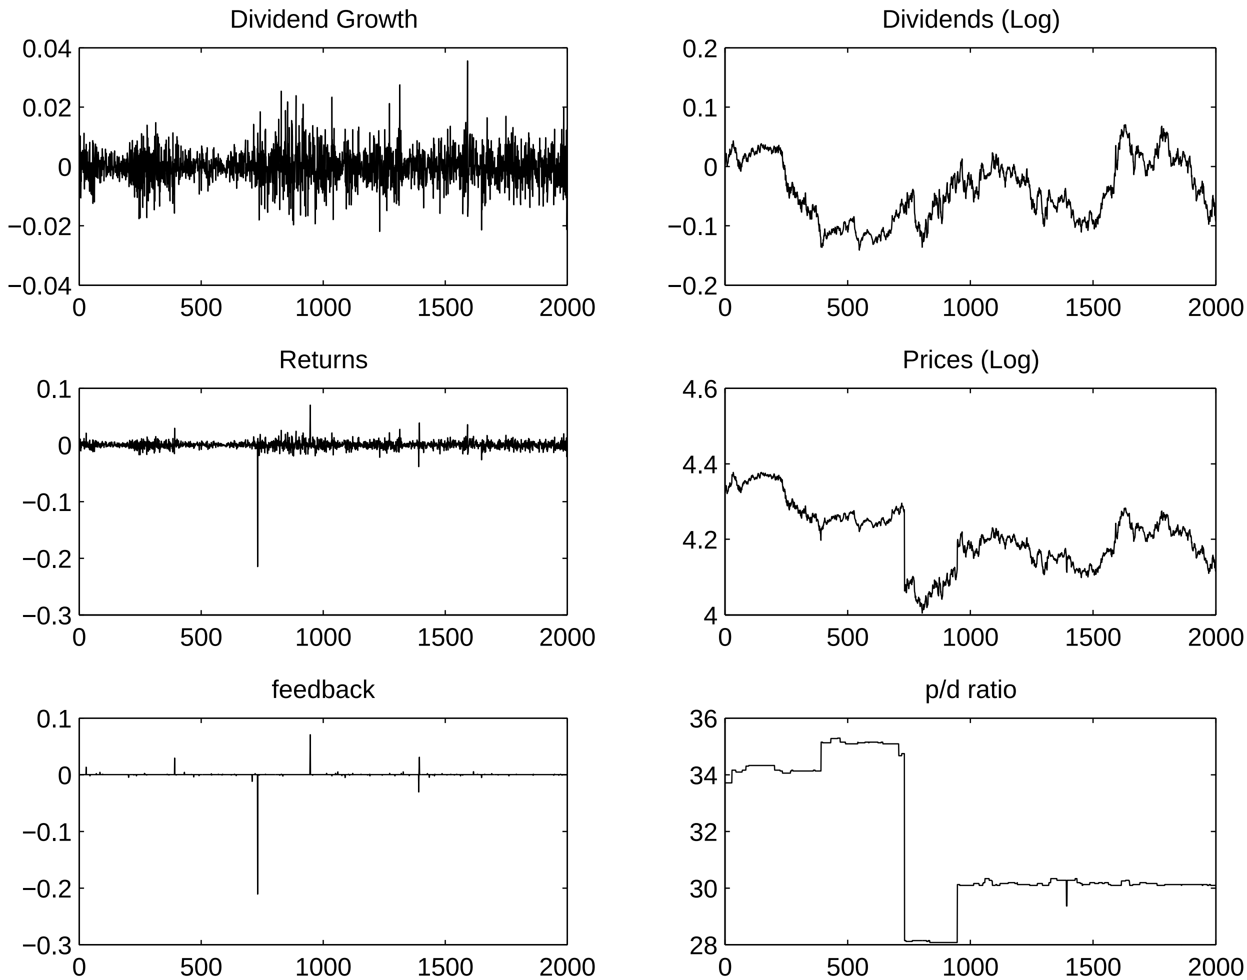

Figure 5: The Multifractal Jump-Diffusion. This figure shows the relation between exogenous dividends and equilibrium prices. The top two panels show dividend growth and dividend levels. The middle panel shows the equilibrium returns and price process, which are considerably more variable. The bottom panels isolate the endogenous portion of returns and prices. The left hand side displays the equilibrium feedback effects, defined as the diffierence between log returns and log dividend growth. The right hand side shows the price:dividend ratio, which varies due to changes in the state variables.