

Multifractality in Asset Returns: Theory and Evidence*

Laurent Calvet

Department of Economics, Harvard University

Adlai Fisher

Finance Division, University of British Columbia

This Draft: February 16, 2001

First Draft: October 1996

*Littauer Center, Cambridge, MA 02138-3001, lcalvet@kuznets.harvard.edu; Faculty of Commerce, 2053 Main Hall, Vancouver, BC, Canada V6T 1Z2, adlai.fisher@commerce.ubc.ca. We are very grateful to Benoit Mandelbrot for inspirational mentoring and exceptional advice during our years at Yale and beyond. We also received helpful comments from T. Andersen, D. Andrews, D. Backus, J. Campbell, P. Clark, T. Conley, R. Deo, F. Diebold, J. Geanakoplos, W. Goetzmann, O. Hart, O. Linton, A. Lo, E. Maskin, R. Merton, P. C. B. Phillips, M. Richardson, S. Ross, J. Scheinkman, R. Shiller, C. Sims, J. Stein, J. Stock, B. Zame, S. Zin, anonymous referees, and seminar participants at numerous institutions, including Yale, Boston University, the Federal Reserve Board, Harvard, Johns Hopkins, MIT, NYU, UC Riverside, the University of Colorado at Boulder, the University of Miami, the University of Utah, the Santa-Fe Institute, the Berkeley Program in Finance, and the 1998 Summer Meeting of the Econometric Society.

Abstract

This paper investigates the *Multifractal Model of Asset Returns*, a class of continuous-time processes that incorporate the thick tails and volatility persistence exhibited by many financial time series. The simplest version of the model compounds a Brownian Motion with a multifractal time-deformation process. Prices follow a semi-martingale, which precludes arbitrage in a standard two-asset economy. Volatility has long memory, and the highest finite moments of returns can take any value greater than two. The local variability of the process is highly heterogeneous, and is usefully characterized by the local Hölder exponent at every instant. In contrast with earlier processes, this exponent takes a *continuum* of values in any time interval. The model also predicts that the moments of returns vary as a power law of the time horizon. We confirm this property for Deutsche Mark/U.S. Dollar exchange rates and several equity series. We then develop an estimator, and infer a parsimonious generating mechanism for the exchange rate series. The moment-scaling rule in the data is replicated by simulated samples from the estimated model.

JEL Classification: G0, C5.

Keywords: Multifractal Model of Asset Returns, Compound Stochastic Process, Time Deformation, Scaling, Self-Similarity, Multifractal Spectrum, Stochastic Volatility.

1. Introduction

The Multifractal Model of Asset Returns (MMAR) is a continuous-time process that captures the thick tails and long-memory volatility persistence exhibited by many financial time series.¹ It is constructed by compounding a standard Brownian motion with a random time-deformation process, which is specified to be multifractal. The time deformation produces clustering and long memory in volatility, and implies that the moments of returns vary as a power law of the time horizon. We empirically confirm this property for the Deutsche Mark/U.S. Dollar exchange rate, a U.S. equity index, and several individual stocks.

The MMAR is characterized by a form of time-invariance called multiscaling, which combines extreme returns with long-memory in volatility. This specification improves on traditional models with scaling features in several ways. First, the MMAR is consistent with economic equilibrium. The simplest version implies uncorrelated returns and semi-martingale prices, thus precluding arbitrage in a standard two-asset economy. The model also permits significant flexibility in matching the data. Returns have a finite variance, and their highest finite moment can take any value greater than two. Consistent with many financial series, the unconditional distribution of returns displays thinner tails as the time scale increases. In contrast with earlier processes, however, the distribution need not converge to a Gaussian at low frequencies and never converges to a Gaussian at high frequencies. The MMAR thus captures the distributional nonlinearities exhibited by financial data, while retaining the parsimony and tractability of scaling models.

The time-deformation process is obtained as the limit of a simple iterative procedure called a multiplicative cascade. The construction begins with a uniform distribution of volatility at a suitably long time horizon, and randomly concentrates volatility into progressively smaller time intervals. The procedure follows the same rule at each stage of the cascade, which provides parsimony and implies moment-scaling. By construction, volatility clustering exists at all frequencies, which corresponds to the intuition that economic factors such as technological shocks, business cycles, earnings cycles, and liquidity shocks have different time

¹Long memory is conveniently defined by hyperbolically declining autocorrelations either for a process itself or functions of it. This property was first analyzed in the context of fractional integration of Brownian motion by Mandelbrot (1965, 1971, 2000), Mandelbrot and van Ness (1968) and Mandelbrot and Wallis (1969). It has been documented in squared and absolute returns for many financial data sets (Taylor, 1986; Ding, Granger, and Engle, 1993; Dacorogna *et al.*, 1993). Baillie (1996) provides a survey of long-memory in economics.

scales.² We anticipate that rational equilibrium models can generate the MMAR, either exogenously through multifractal shocks, or endogenously due to market incompleteness or informational cascades.

The MMAR provides a fundamentally new class of stochastic processes to financial economists. In particular, the multifractal model is a continuous diffusion that lies outside the class of Itô processes. While these traditional models locally vary as $(dt)^{1/2}$ along their sample paths, the MMAR generates the richer class $(dt)^{\alpha(t)}$, where the local scale $\alpha(t)$ can take a continuum of values. The relative occurrences of the local scales $\alpha(t)$ are conveniently summarized in a renormalized probability density called the multifractal spectrum. Given a specification of the model, we provide a general rule for calculating this function, and derive its closed-form expressions in a number of examples. The applied researcher can estimate the spectrum from the moments of the data, and then infer the specification of the multifractal generating process.

Our empirical work begins by examining the Deutsche Mark/U.S. Dollar (“DM/USD”) exchange rate. We use a high frequency data set of approximately 1.5 million quotes collected over one year, and a twenty-three year sample of daily prices. The exchange rate displays the moment-scaling property predicted by the model over a remarkable range of time horizons. We estimate the multifractal spectrum and infer a generating mechanism that replicates DM/USD scaling. Monte-Carlo simulations then show that GARCH and FIGARCH samples are less likely than the MMAR to reproduce these results. We find additional evidence of scaling in a U.S. equity index and five individual stocks.³

Volatility modelling has received considerable attention in finance, and the most common approaches currently include numerous variants of the ARCH/GARCH class (Engle, 1982; Bollerslev, 1986) and stochastic volatility models (Wiggins, 1987).⁴ Because early processes in this literature had difficulty capturing the outliers of financial series, researchers have proposed conditional dis-

²This idea is further elaborated in Calvet and Fisher (1999a).

³The moment-scaling properties of financial returns are also the object of a growing physics literature (Galluccio et al., 1997; Vandewalle and Ausloos, 1998; Pasquini and Serva, 1999, 2000; Richards, 2000). These contributions confirm that multiscaling is exhibited by many financial time series, and are thus complementary of the empirical work contained in FCM (1997) and further developed in this paper. While the physics literature focuses on these phenomenological regularities, the MMAR is a parsimonious stochastic process that allows a unified treatment of the theoretical and empirical properties of the price dynamics.

⁴See Ghysels, Harvey and Renault (1996) for a recent survey of the stochastic volatility literature.

tributions of returns with thicker tails than a Gaussian. In discrete time, these adaptations include the Student- t (Bollerslev, 1987) and non-parametric specifications (Engle and Gonzalez-Rivera, 1991). The problem of modeling thick tails is more acute in continuous time, and is typically addressed by incorporating an independent jump process. Bates (1995, 1996) thus finds that standard diffusions cannot produce tails sufficient to explain the implied volatility smile in option prices, and recommends the incorporation of jumps. Although a continuous diffusion, the MMAR incorporates enough bursts of extreme volatility to capture the fat tails of financial series. It also extends the characterization of volatility in continuous time by considering a multiplicity of local scales. In particular, the multifractal model can generate local oscillations that are intermediate between Itô diffusion and discontinuous jumps.

While early processes from the ARCH/GARCH literature have weak persistence, long-memory in squared returns is a characteristic feature of FIGARCH (Baillie, Bollerslev and Mikkelsen, 1996) and the Long Memory Stochastic Volatility (LMSV) approach (Breidt, Crato, and DeLima, 1997). The MMAR is reminiscent of the long-memory property of these models. In addition, the multifractal process is convenient to analyze under temporal aggregation, and parsimoniously consistent with the moment-scaling properties of financial data.

The multifractal model fundamentally differs from previous volatility models in its scaling properties. The emphasis on scaling originates in the work of Mandelbrot (1963) for extreme variations, and Mandelbrot (1965) and Mandelbrot and van Ness (1968) for long-memory. Multifractality is a form of generalized scaling that includes both extreme variations and long-memory, which was first developed in the context of turbulent dissipation (Mandelbrot, 1972, 1974). These developments are summarized in Mandelbrot (1997), where the use of multifractality in finance is also forcefully advocated.

Section 1.1 discusses the relation between the MMAR and earlier scaling models. Section 2 defines multifractals, and demonstrates their construction through a number of simple examples. Section 3 formalizes the MMAR by compounding a Brownian motion with a continuous time-deformation process. Section 4 shows that multifractal processes can take a continuum of local scales, whose distribution is conveniently characterized by the multifractal spectrum. Section 5 extends the model to permit long-memory in returns. This allows testing of the martingale hypothesis and may be useful in modelling economic series with persistence. In Section 6, we verify the moment-scaling rule for DM/USD exchange rates, and estimate the corresponding multifractal spectrum. We infer a data-

generating process and show that simulated samples replicate the scaling features of the data. Evidence of multifractal scaling is also found in a U.S. equity index and five individual stocks. Section 7 summarizes our results and discusses possible extensions.

This paper simplifies the discussion and extends the results of three earlier working papers ([70], [26], [41]). In the remainder of the text, we refer to the working papers as MFC, CFM, and FCM, signifying the various permutations of the authors. All proofs are contained in the Appendix.

1.1. Review of the Literature

The multifractal model combines several elements of previous research on financial time series. First, the MMAR generates fat tails in the unconditional distribution of returns, and is thus reminiscent of the L -stable processes of Mandelbrot (1963).⁵ The MMAR improves on this earlier model by generating returns with a finite variance, as seems to be empirically the case in most financial series. Furthermore, the L -stable model assumes that increments are independent through time, and have thus the same variability at every instant. In contrast, the MMAR helps model one of the main features of financial markets - fluctuations in volatility.

Second, the multifractal model has long memory in the absolute value of returns, but the returns themselves have a white spectrum. Long memory is the characteristic feature of *fractional Brownian motion* (FBM), introduced by Mandelbrot (1965) and Mandelbrot and van Ness (1968). A FBM, denoted $B_H(t)$, has continuous sample paths, as well as Gaussian and possibly dependent increments. The FBM is an ordinary Brownian motion for $H = 1/2$, is antipersistent when $0 < H < 1/2$, and displays persistence and long memory when $1/2 < H < 1$. Granger and Joyeux (1980) and Hosking (1981) introduced ARFIMA, a discrete-time counterpart of the FBM that helped the use of long memory in economics. FBM and ARFIMA do not disentangle volatility persistence from long-memory in returns.⁶ This has obvious limitations in financial applications and has led to the construction of processes such as FIGARCH and LMSV. Like these recent models, the MMAR separates persistence in volatility from persistence in returns, but in a parsimonious, continuous-time setting.

⁵Recent applications of the L -stable model to foreign exchange rates include Koedijk and Kool (1992) and Phillips, McFarland and McMahon (1996).

⁶Taqqu (1975) establishes that $B_H(t)$ has long memory in the absolute value of increments when $H > 1/2$.

The third essential component of the multifractal model is the concept of trading time, introduced by Mandelbrot and Taylor (1967).

Definition 1. *Let $\{B(t)\}$ be a stochastic process, and $\theta(t)$ an increasing function of t . We call*

$$X(t) \equiv B[\theta(t)]$$

a compound⁷ process. The index t denotes clock time, and $\theta(t)$ is called the trading time or time-deformation process.

When the directing process B is a martingale, the trading time speeds up or slows down the process $X(t)$ without influencing its direction. Compounding can thus separate the direction and the size of price movements, and has been used in the literature to model the unobserved natural time-scale of economic series (Mandelbrot and Taylor, 1967; Clark, 1973; Stock, 1987, 1988). More recently, this method has been used to build models integrating seasonal factors (Dacorogna *et al.*, 1993; Müller *et al.*, 1995), and measures of market activity (Ghysels, Gouriéroux, and Jasiak, 1996). The MMAR also incorporates compounding, and its primary innovation is to specify the trading time θ to be multifractal. While the MMAR is not a structural model of trade, future work may define the trading time θ to be a function of observable data.

Finally, the MMAR generalizes the concept of scaling, in the sense that a well-defined rule relates returns over different sampling intervals. Mandelbrot (1963) suggested that the shape of the distribution of returns should be the same when the time scale is changed, or more formally:

Definition 2. *A random process $\{X(t)\}$ that satisfies*

$$\{X(ct_1), \dots, X(ct_k)\} \stackrel{d}{=} \{c^H X(t_1), \dots, c^H X(t_k)\}$$

for some $H > 0$ and all $c, k, t_1, \dots, t_k \geq 0$, is called self-affine.⁸ The number H is the self-affinity index, or scaling exponent, of the process $\{X(t)\}$.

⁷Processes of this type have also been called subordinated in the recent economics literature. In mathematics, subordination differs from compounding, and requires that $\theta(t)$ have independent increments (Bochner, 1955; Feller, 1968). The economics literature has evolved to describe any generic time deformation process as a subordinator.

⁸Self-affine processes are sometimes called self-similar in the literature.

The Brownian motion, the L -stable process and the FBM are the main examples of self-affine processes in finance. Empirical evidence suggests that many financial series are not exactly self-affine, but instead have thinner tails and become less peaked in the bells when the sampling interval increases. The MMAR captures this feature, as well as a generalized version of self-affinity exhibited by the data. While maintaining the simplicity of self-affine processes, the MMAR is thus sufficiently flexible to model the nonlinearities, fat tails, and long-memory volatility persistence exhibited by many financial time series.

2. Multifractal Measures and Processes

The MMAR is constructed in Section 3 by compounding a Brownian motion $B(t)$ with a random increasing function $\theta(t)$:

$$\ln P(t) - \ln P(0) = B[\theta(t)].$$

The trading time $\theta(t)$ will be specified as the cumulative distribution function (c.d.f.) of a multifractal measure μ , a concept which we now present.

2.1. The Binomial Measure

Multifractal measures can be built by iterating a simple procedure called a *multiplicative cascade*. We first present one of the simplest examples, the binomial measure⁹ on $[0, 1]$. Consider the uniform probability measure μ_0 on the unit interval, and two positive numbers m_0 and m_1 adding up to 1. In the first step of the cascade, 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, as illustrated in Figure 1a.

In the second stage of the cascade, 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 of $\mu_1[0, 1/2]$, while the right subinterval $[1/4, 1/2]$ receives a fraction m_1 . Applying a similar procedure to $[1/2, 1]$, we obtain a measure μ_2 such that:

$$\begin{aligned} \mu_2[0, 1/4] &= m_0 m_0, & \mu_2[1/4, 1/2] &= m_0 m_1, \\ \mu_2[1/2, 3/4] &= m_1 m_0, & \mu_2[3/4, 1] &= m_1 m_1. \end{aligned}$$

⁹The binomial is sometimes called the Bernoulli or Besicovitch measure.

Iteration of this procedure generates an infinite sequence of measures (μ_k) that weakly converges to the binomial measure μ . Figure 1b illustrates the density of the measure μ_4 obtained after $k = 4$ steps of the recursion.

Since $m_0 + m_1 = 1$, each stage of the construction preserves the mass of split dyadic intervals.¹⁰ Consider the interval $[t, t + 2^{-k}]$, where $t = \overline{0.\eta_1\dots\eta_k} = \sum_{i=1}^k \eta_i 2^{-i}$ for some $\eta_1, \dots, \eta_k \in \{0, 1\}$. We denote by φ_0 and φ_1 the relative frequencies of 0s and 1s in (η_1, \dots, η_k) . The measure of the dyadic interval then simplifies to $\mu[t, t + 2^{-k}] = m_0^{k\varphi_0} m_1^{k\varphi_1}$. This illustrates that the construction creates large and increasing heterogeneity in the allocation of mass. As a result, the binomial, like many multifractals, is a continuous but singular probability measure that has no density and no point mass.

This construction is easily generalized. For instance, we can uniformly split intervals into an arbitrary number $b \geq 2$ of cells at each stage of the cascade. Subintervals, indexed from left to right by β ($0 \leq \beta \leq b - 1$), then receive fractions m_0, \dots, m_{b-1} of the measure. We preserve mass in the construction by imposing that these fractions, also called *multipliers*, add up to one: $\sum m_\beta = 1$. This defines the class of *multinomial* measures, which are discussed in Mandelbrot (1989a).

A more significant extension randomizes the allocation of mass between subintervals. The multiplier of each cell is then a discrete random variable M_β taking values m_0, m_1, \dots, m_{b-1} with probabilities p_0, \dots, p_{b-1} . We preserve mass in the construction by imposing the additivity constraint: $\sum M_\beta = 1$. This modified algorithm generates a *random* multifractal measure. Figure 1c shows the random density obtained after $k = 10$ iterations with parameters $b = 2$, $p = p_0 = 0.5$ and $m_0 = 0.6$. This density, which represents the flow of trading time, begins to show the properties we desire in modeling financial volatility. The occasional bursts of trading time generate thick tails in the compound price process, and their clustering generates volatility persistence. Because the reshuffling of mass follows the same rule at each stage of the cascade, volatility clustering is present at all time scales.

2.2. Multiplicative Measures

We can also consider non-negative multipliers M_β ($0 \leq \beta \leq b - 1$) with *arbitrary* distributions. Assume for simplicity that all multipliers are identically distributed

¹⁰A number $t \in [0, 1]$ is called *dyadic* if $t = 1$ or $t = \eta_1 2^{-1} + \dots + \eta_k 2^{-k}$ for a finite k and $\eta_1, \dots, \eta_k \in \{0, 1\}$. A dyadic interval has dyadic endpoints.

($M_\beta \stackrel{d}{=} M \ \forall \beta$), and that multipliers defined at different stages of the construction are independent. The limit *multiplicative measure* is called *conservative* when mass is conserved exactly at each stage: $\sum M_\beta \equiv 1$, and *canonical* when it is preserved only *on average*: $\mathbb{E}(\sum M_\beta) = 1$ or equivalently $\mathbb{E}M = 1/b$. A canonical measure can be conveniently generated by choosing *independent* multipliers M_β within each stage of the cascade.

The moments of multiplicative measures have interesting scaling properties. To show this, first consider the generating cascade of a *conservative* measure μ . Stage 1 uniformly splits the unit interval into cells of length b^{-1} , and allocates random masses M_0, \dots, M_{b-1} to each cell. Similarly, the measure of a b -adic cell of length $\Delta t = b^{-k}$, starting at $t = \overline{0.\eta_1\dots\eta_k} = \sum \eta_i b^{-i}$, is the product of k multipliers:

$$\mu(\Delta t) = M_{\eta_1} M_{\eta_1, \eta_2} \dots M_{\eta_1, \dots, \eta_k}. \quad (2.1)$$

Since multipliers defined at different stages of the cascade are independent, we infer that $\mathbb{E}[\mu(\Delta t)^q] = [\mathbb{E}(M^q)]^k$ or equivalently

$$\mathbb{E}[\mu(\Delta t)^q] = (\Delta t)^{\tau(q)+1}, \quad (2.2)$$

where $\tau(q) = -\log_b \mathbb{E}(M^q) - 1$. The moment of an interval's measure is thus a power function of the length Δt . This important scaling rule characterizes multifractals.

Scaling relation (2.2) easily generalizes to a *canonical* measure μ , which by definition is generated by a cascade that only conserves mass *on average*: $\mathbb{E}(\sum M_\beta) = 1$. The mass of the unit interval is then a random variable $\Omega = \mu[0, 1] \geq 0$. More generally, the measure of a b -adic cell satisfies

$$\mu(\Delta t) = M_{\eta_1} M_{\eta_1, \eta_2} \dots M_{\eta_1, \dots, \eta_k} \Omega_{\eta_1, \dots, \eta_k}, \quad (2.3)$$

where $\Omega_{\eta_1, \dots, \eta_k}$ has the same distribution as Ω . This directly implies the scaling relationship

$$\mathbb{E}[\mu(\Delta t)^q] = \mathbb{E}(\Omega^q) (\Delta t)^{\tau(q)+1}, \quad (2.4)$$

which generalizes (2.2).

The right tail of the measure $\mu(\Delta t)$ is determined by the way mass is preserved at each stage of the construction. When μ is conservative, the mass of the cell is bounded above by the deterministic mass of the unit interval:

$0 \leq \mu(\Delta t) \leq \mu[0, 1] = 1$, and has therefore finite moments of every order. On the other hand, consider a canonical measure generated by independent multipliers M_β . We assume for simplicity that $\mathbb{E}(M^q) < \infty$ for all q . Guivarc'h (1987) shows that the random mass $\Omega \geq 0$ of the unit interval then has a Paretian right tail:

$$\mathbb{P}\{\Omega > \omega\} \sim C_1 \omega^{-q_{crit}} \text{ as } \omega \rightarrow +\infty,$$

where $C_1 > 0$ and the critical moment q_{crit} is finite and larger than one: $1 < q_{crit} < \infty$.¹¹ By (2.3), the mass of every cell has the same critical moment q_{crit} as the random variable Ω . The property $q_{crit} > 1$ will prove particularly important because it implies that returns have a finite variance in the MMAR.

The multiplicative measures constructed so far are *grid-bound*, in the sense that the scaling rule (2.4) holds only when $t = \overline{0.\eta_1, \dots, \eta_k}$ and $\Delta t = b^{-l}$, $l \geq k$. Let \mathcal{D} denote the set of couples $(t, \Delta t)$ satisfying scaling rule (2.4). \mathcal{D} has interesting topological properties that are summarized in Condition 1 of Appendix 8.1. Alternatively, we can consider *grid-free* random measures that satisfy scaling rule (2.4) for all admissible values of $(t, \Delta t)$ (Mandelbrot, 1989a). This leads to the following

Definition 3. *A random measure μ defined on $[0, 1]$ is called multifractal if it satisfies*

$$\mathbb{E}(\mu[t, t + \Delta t]^q) = c(q)(\Delta t)^{\tau(q)+1} \quad \text{for all } (t, \Delta t) \in \mathcal{D}, q \in \mathcal{Q},$$

where \mathcal{D} is a subset of $[0, 1] \times [0, 1]$, \mathcal{Q} is an interval, and $\tau(q)$ and $c(q)$ are functions with domain \mathcal{Q} . Moreover, $[0, 1] \subseteq \mathcal{Q}$, and \mathcal{D} satisfies Condition 1, which is defined in the Appendix.

Maintaining the distinction between grid-bound and grid-free measures would prove cumbersome and lead to unnecessary technicalities. We will therefore neglect the difference between the two classes in the remainder of this paper. The interested reader can refer to Calvet and Fisher (1999a) for a detailed treatment of grid-free multifractals.

2.3. Multifractal Processes

Multifractality is easily extended from measures to stochastic processes:

¹¹The cascade construction also implies that Ω satisfies the invariance relation $\sum_{i=1}^b M_i \Omega_i \stackrel{d}{=} \Omega$, where $M_1, \dots, M_b, \Omega_1, \dots, \Omega_b$ are independent copies of the random variables M and Ω .

Definition 4. A stochastic process $\{X(t)\}$ is called multifractal if it has stationary increments and satisfies

$$\mathbb{E}(|X(t)|^q) = c(q)t^{\tau(q)+1}, \text{ for all } t \in \mathcal{T}, q \in \mathcal{Q}, \quad (2.5)$$

where \mathcal{T} and \mathcal{Q} are intervals on the real line, $\tau(q)$ and $c(q)$ are functions with domain \mathcal{Q} . Moreover, \mathcal{T} and \mathcal{Q} have positive lengths, and $0 \in \mathcal{T}$, $[0, 1] \subseteq \mathcal{Q}$.

The function $\tau(q)$ is called the *scaling function* of the multifractal process. Setting $q = 0$ in condition (2.5), we see that all scaling functions have the same intercept $\tau(0) = -1$. In addition, it is easy to show

Proposition 1. The scaling function $\tau(q)$ is concave.

We will see that the distinction between linear and nonlinear scaling functions $\tau(q)$ is particularly important.

A self-affine process $\{X(t)\}$ is multifractal and has a linear function $\tau(q)$, as is now shown. Denoting by H the self-affinity index, we observe that the invariance condition $X(t) \stackrel{d}{=} t^H X(1)$ implies $\mathbb{E}(|X(t)|^q) = t^{Hq} \mathbb{E}(|X(1)|^q)$. Scaling rule (2.5) therefore holds with $c(q) = \mathbb{E}(|X(1)|^q)$ and

$$\tau(q) = Hq - 1.$$

In this special case, the scaling function $\tau(q)$ is linear and fully determined by its index H . More generally, linear scaling functions $\tau(q)$ are determined by a unique parameter, their slope, and the corresponding processes are called *uniscaling* or *unifractal*.

Uniscaling processes, which may seem appealing for their simplicity, do not satisfactorily model asset returns. This is because most financial data sets have thinner tails and become less peaked in the bell when the sampling intervals Δt increases. In this paper, we focus on *multiscaling* processes, which have a nonlinear $\tau(q)$. They provide a parsimonious framework with strict moment conditions, and enough flexibility to model a wide range of financial prices.

3. The Multifractal Model of Asset Returns

We now formalize construction of the MMAR. Consider the price of a financial asset $P(t)$ on a bounded interval $[0, T]$, and define the *log-price* process

$$X(t) \equiv \ln P(t) - \ln P(0).$$

We model $X(t)$ by compounding a Brownian motion with a multifractal trading time:

Assumption 1. $X(t)$ is a compound process

$$X(t) \equiv B[\theta(t)]$$

where $B(t)$ is a Brownian motion, and $\theta(t)$ is a stochastic trading time.

Assumption 2. Trading time $\theta(t)$ is the c.d.f. of a multifractal measure μ defined on $[0, T]$.

Assumption 3. The processes $\{B(t)\}$ and $\{\theta(t)\}$ are independent.

This construction generates a large class of multifractal processes.

We will show that the price process is a semi-martingale, which implies the absence of arbitrage in simple economies. A straightforward generalization of this model allows the broader class of fractional Brownians $B_H(t)$ in Assumption 1, as developed in Section 5. In Assumption 2, the multifractal measure μ can be multinomial or multiplicative, which implies a continuous trading time $\theta(t)$ with non-decreasing paths and stationary increments. Assumption 3 ensures that the unconditional distribution of returns is symmetric. Weakening this assumption allows leverage effects, as in EGARCH (Nelson, 1991) and Glosten, Jagannathan and Runkle (1993), and is a promising direction for future research.

Under the above assumptions,

Theorem 1. The log-price $X(t)$ is a multifractal process with stationary increments and scaling function $\tau_X(q) \equiv \tau_\theta(q/2)$.

Trading time controls the tails of the process $X(t)$. As shown in the proof, the q -th moment of X exists if (and only if) the process θ has a moment of order $q/2$. In particular if $\mathbb{E} |X(t)|^q$ is finite for some instant t , then it is finite for all t . We therefore drop the time index when discussing the critical moment of the multifractal process.

The tails of $X(t)$ have different properties if the generating measure is conservative or canonical. This follows directly from the discussion of Section 2.2. If μ is *conservative*, trading time is bounded, and the process $X(t)$ has finite moments of all (non-negative) order. Conservative measures thus generate “mild” processes with relatively thin tails. Conversely, the total mass $\theta(T) \equiv \mu[0, T]$

of a *canonical* measure is a random variable with Paretian tails. In particular, there exists a *critical exponent* $q_{crit}(\theta) > 1$ for trading time such that $\mathbb{E}\theta^q$ is finite when $0 \leq q < q_{crit}(\theta)$, and infinite when $q \geq q_{crit}(\theta)$.¹² The log-price $X(t)$ then has infinite moments, and is accordingly called “wild”. Note however that $X(t)$ always has finite variance, since $q_{crit}(X) = 2q_{crit}(\theta) > 2$. Overall, the MMAR has enough flexibility to accommodate a wide variety of tail behaviors.

We can also analyze how the unconditional distribution of returns varies with the time horizon t . Consider for instance a conservative measure μ such as a random binomial. At the final instant T , the trading time $\theta(T)$ is deterministic, implying that the random variable $X(T)$ is normally distributed. As we move to a smaller horizon t , the allocation of mass becomes increasingly heterogeneous, as is apparent in Figure 1. The tails of returns thus become thicker at higher frequencies. The mass of a dyadic cell can be written as $\mu[t, t + 2^{-k}] = m_0^{k\varphi_0} m_1^{k(1-\varphi_0)}$, where $t = \overline{0.\eta_1\dots\eta_k}$ and φ_0 denotes the proportion of the multipliers $M_{\eta_1}, \dots, M_{\eta_1, \dots, \eta_k}$ equal to m_0 . By the law of large numbers, draws of φ_0 concentrate increasingly in the neighborhood of $1/2$ as k increases, implying that the bell of the distribution becomes thicker. The distribution of $X(t)$ thus accumulates more mass in the tails and in the bell as the time horizon decreases, while the middle of the distribution becomes thinner. This property, which is consistent with empirical observations, is further elaborated in Calvet and Fisher (1999b). In addition, when the measure μ is canonical, the random variables $\theta(T)$ and $X(T)$ have Paretian tails, thus illustrating that multifractal returns need not converge to a Gaussian at low frequency.

The model also has an appealing autocorrelation structure.

Theorem 2. *The price $\{P(t)\}$ is a semi-martingale (with respect to its natural filtration), and the process $\{X(t)\}$ is a martingale with finite variance and thus uncorrelated increments.*

The model thus implies that asset returns have a white spectrum, a property which has been extensively discussed in the market efficiency literature.¹³

¹²We also know that the scaling function $\tau_\theta(q)$ is negative when $0 < q < 1$, and positive when $1 < q < q_{crit}(\theta)$.

¹³See Campbell, Lo and MacKinlay (1997) for a recent discussion of these concepts. We also note that immediate extensions of the MMAR could add trends or other predictable components to the compound process in order to fit different financial time series.

The price $P(t)$ is a semi-martingale,¹⁴ which has important consequences for arbitrage.¹⁵ Consider for instance the *two asset economy* consisting of the multifractal security with price $P(t)$, and a riskless bond with constant rate of return r . Following Harrison and Kreps (1979), we can analyze if arbitrage profits can be made by frequently rebalancing a portfolio of these two securities. Theorem 2 directly implies

Theorem 3. *There are no arbitrage opportunities in the two asset economy.*

This suggests that future research may seek to embed the MMAR in standard financial models. Since the price $P(t)$ is a semi-martingale, stochastic integration can be used to calculate the gains from trading multifractal assets, which in future work will greatly help to develop portfolio selection and option pricing applications. Further research will also seek to integrate multifractality into equilibrium theory. We may thus obtain the MMAR in a general equilibrium model with *exogenous* multifractal technological shocks, in the spirit of Cox, Ingersoll and Ross (1985). Such a methodology is justified by the multifractality of many natural phenomena, such as weather patterns, and will help build new economic models of asset and commodity prices. Another line of research could also obtain multifractality as an *endogenous* equilibrium property, which might stem from the incompleteness of financial markets (Calvet, 1997) or informational cascades (Genotte and Leland, 1990; Bikhchandani, Hirshleifer and Welch, 1992; Jacklin, Kleidon and Pfleiderer, 1992; Bulow and Klemperer, 1994; Avery and Zemsky, 1998).

Recent research focuses not only on predictability in returns, but also on persistence in the size of price changes. The MMAR adds to this literature by proposing a continuous time model with long memory in volatility. Because the price process is only defined on a *bounded* time range, the definition of long memory seems problematic. We note, however, that for any stochastic process Z with stationary increments $Z(a, \Delta t) \equiv Z(a + \Delta t) - Z(a)$, the *autocovariance in levels*

$$\delta_Z(t, q) = Cov(|Z(a, \Delta t)|^q, |Z(a + t, \Delta t)|^q),$$

quantifies the dependence in the size of the process's increments. It is well-defined when $\mathbb{E}|Z(a, \Delta t)|^{2q}$ is finite. For a fixed q , we say that the process has *long*

¹⁴Since the process $X(t) = \ln P(t)$ is a martingale, Jensen's inequality implies that the price $P(t)$ is a submartingale but *not* a martingale. This result is of course not specific to the MMAR.

¹⁵See Dothan (1990) for a discussion of semi-martingales in the context of finance.

memory in the size of increments if the autocovariance in levels is hyperbolic in t when $t/\Delta t \rightarrow \infty$. When the process Z is multifractal, this concept does not depend on the particular choice of q .¹⁶ It is easy to show that when μ is a multiplicative measure,

Theorem 4. *Trading time $\theta(t)$ and log-price $X(t)$ have long memory in the size of increments.*

This result can be illustrated graphically. Figure 2 shows simulated first differences when $\theta(t)$ the c.d.f. of a randomized binomial measure with multiplier $m_0 = 0.6$. The simulated returns displays marked temporal heterogeneity at all time scales and intermittent large fluctuations.

The MMAR is thus a flexible continuous time framework that accommodates long memory in volatility, a variety of tail behaviors, and unpredictability in returns. Furthermore, the multifractal model contains volatility persistence at all time frequencies. Table 1 compares the MMAR with existing models of financial time series.

4. The Multifractal Spectrum

This section examines the geometric properties of sample paths in the MMAR. While we previously focused on *global* properties such as moments and autocovariances, we now adopt a more *local* viewpoint and examine the regularity of realized paths around a given instant. The analysis builds on a concept borrowed from real analysis, the local Hölder exponent. On a given path, the infinitesimal variation in price around a date t is heuristically of the form¹⁷

$$|\ln P(t + dt) - \ln P(t)| \sim C_t(dt)^{\alpha(t)},$$

where $\alpha(t)$ and C_t are respectively called the *local Hölder exponent* and the *prefactor* at t . As is apparent in this definition, the exponent $\alpha(t)$ quantifies the scaling properties of the process at a given point in time, and is also called the *local scale* of the process at t .

¹⁶Provided that $\mathbb{E}|Z(a, \Delta t)|^{2q} < \infty$, as is implicitly assumed in the rest of the paper.

¹⁷The expression $(dt)^{\alpha(t)}$ is an example of “non-standard infinitesimal”, as developed by Abraham Robinson.

In continuous Itô diffusions, the Hölder exponent takes the unique value $\alpha(t) = 1/2$ at every instant.¹⁸ For this reason, traditional research obtains time variations in market volatility through changes in the prefactor C_t . In contrast, the MMAR contains a *continuum* of local scales $\alpha(t)$ within any finite time interval. Thus, multifractal processes are *not* continuous Itô diffusions and cannot be generated by standard techniques. Fractal geometry imposes that in the MMAR, the instants $\{t : \alpha(t) < \alpha\}$ with local scale less than α cluster in clock time, thus accounting for the concentration of outliers in our model. The relative frequency of the local exponents can be represented by a renormalized density called the *multifractal spectrum*. For a broad class of multifractals, we calculate this spectrum by an application of Large Deviation Theory.

4.1. Local Scales

We first introduce

Definition 5. *Let g be a function defined on the neighborhood of a given date t . The number*

$$\alpha(t) = \text{Sup} \{ \beta \geq 0 : |g(t + \Delta t) - g(t)| = O(|\Delta t|^\beta) \text{ as } \Delta t \rightarrow 0 \}$$

is called the local Hölder exponent or local scale of g at t .

The Hölder exponent thus describes the local scaling of a path at a point in time, and lower values correspond to more abrupt variations. The exponent $\alpha(t)$ is non-negative when the function g is bounded around t , as is always the case in this paper. The definition readily extends to measures on the real line. At a given date t , a measure simply has the local exponent of its c.d.f.

We can easily compute Hölder exponents for many functions and processes. For instance, the local scale of a function is 0 at points of discontinuity, and 1 at (non-singular) differentiable points. Smooth functions thus have integral exponents almost everywhere. On the other hand, the unique scale $\alpha(t) = 1/2$ is observed on the jagged sample paths of a Brownian motion or of a continuous Itô diffusion. Similarly, a fractional Brownian $B_H(t)$ is characterized by a unique exponent $\alpha(t) = H$. Thus, the continuous processes typically used in finance

¹⁸More precisely, the set $\{t : \alpha(t) \neq 1/2\}$ of instants with a local scale different from $1/2$ has a Hausdorff-Besicovitch measure (and therefore a Lebesgue measure) equal to zero. This set can thus be neglected in our analysis. See Kahane (1997) for a recent survey of this topic.

each have a unique Hölder exponent. In contrast, multifractal processes contain a continuum of local scales.

The mathematics literature has developed a convenient representation for the distribution of Hölder exponents in a multifractal. This representation, called the multifractal spectrum, is a function $f(\alpha)$ that we now describe. From Definition 5, the Hölder exponent $\alpha(t)$ is the liminf of the ratio

$$\ln |g(t, \Delta t)| / \ln(\Delta t) \text{ as } \Delta t \rightarrow 0,$$

where, consistent with previous notation, $g(t, \Delta t) \equiv g(t + \Delta t) - g(t)$. This suggests estimating the distribution of the local scale $\alpha(t)$ at a random instant. For increasing $k \geq 1$, we partition $[0, T]$ into b^k subintervals $[t_i, t_i + \Delta t]$, where length $\Delta t = b^{-k}T$, and calculate for each subinterval the *coarse Hölder exponent*

$$\alpha_k(t_i) \equiv \ln |g(t_i, \Delta t)| / \ln \Delta t.$$

This operation generates a set $\{\alpha_k(t_i)\}$ of b^k observations. We then divide the range of α s into small intervals of length $\Delta\alpha$, and denote by $N_k(\alpha)$ the number of coarse exponents contained in $(\alpha, \alpha + \Delta\alpha]$. It would then be natural to calculate a histogram with the relative frequencies $N_k(\alpha)/b^k$, which converge as $k \rightarrow \infty$ to the probability that a random instant t has Hölder exponent α . Using this method, however, the histogram would degenerate into a spike and thus fail to distinguish the MMAR from traditional processes. This is because multifractals typically have a dominant exponent α_0 , in the sense that $\alpha(t) = \alpha_0$ at almost every instant. Mandelbrot (1974, 1989a) instead suggested

Definition 6. *The limit*

$$f(\alpha) \equiv \lim \left\{ \frac{\ln N_k(\alpha)}{\ln b^k} \right\} \text{ as } k \rightarrow \infty \quad (4.1)$$

represents a renormalized probability distribution of local Hölder exponents, and is called the multifractal spectrum.

For instance if $b = 3$ and $N_k(\alpha) = 2^k$, the frequency $N_k(\alpha)/b^k = (2/3)^k$ converges to zero as $k \rightarrow \infty$, while the ratio $\ln N_k(\alpha) / \ln b^k = \ln 2 / \ln 3$ is a positive constant. The multifractal spectrum thus helps to identify events that happen many times in the construction but at a vanishing frequency.

Frisch and Parisi (1985) and Halsey *et al.* (1986) interpreted $f(\alpha)$ as the fractal dimension of $T(\alpha) = \{t \in [0, T] : \alpha(t) = \alpha\}$, the set of instants having

local Hölder exponent α . For various levels of the scale α , Figure 1d illustrates the subintervals with coarse exponent $\alpha_k(t_i) < \alpha$. When the number of iterations k is sufficiently large, these “cuts” display a self-similar structure. Appendix 8.6 provides a more detailed discussion of this interpretation.

4.2. The Spectrum of Multiplicative Measures

We now use Large Deviation Theory to compute the spectrum of multiplicative measures. First consider a *conservative* measure μ defined on the unit interval $[0, 1]$. After k iterations, we know the masses $\mu[t, t + \Delta t] = M_{\eta_1} M_{\eta_1, \eta_2} \dots M_{\eta_1, \dots, \eta_k}$ in intervals of length $\Delta t = b^{-k}$. The coarse exponents $\alpha_k(t) = \ln \mu[t, t + \Delta t] / \ln \Delta t$ can thus be rewritten

$$\alpha_k(t) = -(\log_b M_{\eta_1} + \dots + \log_b M_{\eta_1, \dots, \eta_k}) / k. \quad (4.2)$$

The multifractal spectrum is obtained by forming renormalized histograms of these exponents. Letting $V_i \equiv -\log_b M_{\eta_1, \dots, \eta_i}$, we can interpret the coarse Hölder exponents as draws of the random variable

$$\alpha_k = \frac{1}{k} \sum_{i=1}^k V_i. \quad (4.3)$$

The spectrum $f(\alpha)$ then directly depends on the asymptotic distribution of α_k .

By the Strong Law of Large Numbers, α_k converges almost surely to¹⁹

$$\alpha_0 = \mathbb{E} V_1 = -\mathbb{E} \log_b M > 1. \quad (4.4)$$

As $k \rightarrow \infty$, almost all coarse exponents are contained in a small neighborhood of α_0 . The standard histogram $N_k(\alpha) / b^k$ thus collapses to a spike at α_0 as anticipated in Section 4.1. The other coarse exponents nonetheless matter greatly. In fact, most of the mass concentrates on intervals with Hölder exponents that are bounded away from α_0 .²⁰ Information on these “rare events” is presumably contained in the tail of the random variable α_k .

¹⁹The relation $-\mathbb{E} \log_b M > 1$ follows from Jensen’s inequality and $\mathbb{E} M = 1/b$.

²⁰Let T_k denote the set of b -adic cells with local exponents greater than $(1 + \alpha_0)/2$. When k is large, T_k contains “almost all” cells, but its mass:

$$\sum_{t \in T_k} (\Delta t)^{\alpha_k(t)} \leq b^k (\Delta t)^{(\alpha_0+1)/2} = b^{-k(\alpha_0-1)/2}$$

vanishes as k goes to infinity.

Tail behavior is the object of Large Deviation Theory. In 1938, H. Cramér established the following theorem under conditions that were gradually weakened.

Theorem 5. *Let $\{X_k\}$ denote a sequence of iid random variables. Then as $k \rightarrow \infty$,*

$$\frac{1}{k} \ln \mathbb{P} \left\{ \frac{1}{k} \sum_{i=1}^k X_i > \alpha \right\} \rightarrow \text{Inf}_q \ln [\mathbb{E} e^{q(\alpha - X_1)}],$$

for any $\alpha > \mathbb{E}X_1$.

Proofs can be found in Billingsley (1979) and Durrett (1991). The theorem implies

Theorem 6. *The multifractal spectrum $f(\alpha)$ is the Legendre transform*

$$f(\alpha) = \text{Inf}_q [\alpha q - \tau(q)] \tag{4.5}$$

of the scaling function $\tau(q)$.

This result holds for both conservative and canonical measures. It provides the foundation of the empirical work developed in Section 6, where an estimation procedure for the scaling function $\tau(q)$ is obtained and the Legendre transform yields an estimate of the multifractal spectrum $f(\alpha)$.

The theorem allows us to derive explicit formulae for the spectrum in a number of useful examples. To aid future reference, we denote by $f_\theta(\alpha)$ the spectrum common to a measure μ and its c.d.f. θ . Begin by considering a measure generated by a log-normal multiplier M with distribution $-\log_b M \sim \mathcal{N}(\lambda, \sigma^2)$. Conservation of mass imposes that $\mathbb{E} M = 1/b$ or equivalently $\sigma^2 = 2 \ln b / (\lambda - 1)$. It is easy to show that the scaling function $\tau(q) \equiv -\log_b(\mathbb{E} M^q) - 1$ has the closed-form expression $\tau(q) = \lambda q - 1 - q^2 \sigma^2 (\ln b) / 2$. We infer from Theorem 6 that the multifractal spectrum is a quadratic function

$$f_\theta(\alpha) = 1 - (\alpha - \lambda)^2 / [4(\lambda - 1)]$$

parameterized by a unique number $\lambda > 1$. More generally, Table 2 reports the spectrum when the random variable V is binomial, Poisson or Gamma (see CFM for detailed derivations). We note that the function $f_\theta(\alpha)$ is very sensitive to the distribution of the multiplier, which suggests that the MMAR has enough flexibility to model a wide range of financial prices. In the empirical work, this allows us to identify a multiplicative measure from its estimated spectrum.

4.3. Application to the MMAR

We now examine the spectrum of price processes generated by the MMAR. Denoting by $f_Z(\alpha)$ the spectrum of a process $Z(t)$, we show

Theorem 7. *The price $P(t)$ and the log-price $X(t)$ have identical multifractal spectra: $f_P(\alpha) \equiv f_X(\alpha) \equiv f_\theta(2\alpha)$.*

The log-price $X(t)$ contains a continuum of local exponents, and thus *cannot* be generated by an Itô diffusion process. Let $\alpha_0(Z)$ denote the most probable exponent of a process Z . Since $\alpha_0(\theta) > 1$, the log-price has a local scale $\alpha_0(X) \equiv \alpha_0(\theta)/2$ larger than $1/2$ at almost every instant. Despite their apparent irregularity, the MMAR's sample paths are almost everywhere *smoother* than the paths of a Brownian motion. Section 4.2 indicates that the variability of the MMAR is in fact explained by the “rare” local scales $\alpha < \alpha_0(X)$. While jump diffusions permit negligible sets to contribute to the total variation, multifractal processes are notable for combining *continuous* paths with variations dominated by rare events.

Although the local scale is larger than $1/2$ almost everywhere, Theorem 1 implies that the standard deviation of the process

$$\sqrt{\mathbb{E} \{ [X(t + \Delta t) - X(t)]^2 \}} = c_X(2)^{1/2} \sqrt{\Delta t}$$

is of the order $(\Delta t)^{1/2}$. Thus while most shocks are of order $(\Delta t)^{\alpha_0(X)}$, the exponents $\alpha < \alpha_0(X)$ appear frequently enough to alter the scaling properties of the variance. This contrasts with the textbook analysis that a standard deviation in $(\Delta t)^{1/2}$ implies that most shocks are of the same order.²¹ We expect these findings to have interesting consequences for decision and equilibrium theory.

²¹Merton (1990, ch. 3) provides an interesting discussion of multiple local scales and “rare events” in financial processes. Assume that the price variation over a time interval Δt is a discrete random variable taking values $\varepsilon_1, \dots, \varepsilon_m$ with probability p_1, \dots, p_m , and assume moreover that $p_i \sim (\Delta t)^{q_i}$, $\varepsilon_i \sim (\Delta t)^{r_i}$ and $r_i > 0$ for all i . Denote by I the events i such that $p_i \varepsilon_i^2 \sim \Delta t$. When the variance of the process $\sum_{i=1}^m p_i \varepsilon_i^2$ is of the order Δt , only events in I contribute to the variance. If all events belong to I , Merton establishes that only events of the order $(\Delta t)^{1/2}$ matter. The MMAR shows that events outside I can play a crucial role in the statistical properties of the price process, a property previously overlooked in the literature.

5. An Extension with Autocorrelated Returns

The multifractal model presented in Section 3 is characterized by long memory in volatility but the absence of correlation in returns. While there is little evidence of fractional integration in stock returns (Lo, 1991), long memory has been identified in the first differences of many economic series,²² including aggregate output (Adelman, 1965; Diebold and Rudebusch, 1989; Sowell, 1992), the Beveridge (1925) Wheat Price Index, the US Consumer Price Index (Baillie, Chung and Tieslau, 1996), and interest rates (Mandelbrot, 1971; Backus and Zin, 1993).²³ This has led authors to model these series with the FBM or the discrete-time ARFIMA specification. We note, however, that these economic series have volatility patterns which seem closer to the multifractal model than to the fractional Brownian motion. This suggests using the fractional Brownian motion of multifractal time.

We model an economic series $X(t)$ by replacing Assumption 1 in Section 3 with

Assumption 1a. $X(t)$ is a compound process

$$X(t) \equiv B_H[\theta(t)]$$

where $B_H(t)$ is a Fractional Brownian motion, and $\theta(t)$ is a stochastic trading time.

In addition, we maintain the multifractality of trading time (Assumption 2) and the independence of the processes $B_H(t)$ and $\theta(t)$ (Assumption 3). Note that this coincides with the earlier model if $H = 1/2$. For other values of the index H , the increments of $X(t)$ display either antipersistent ($H < 1/2$) or positive autocorrelations and long memory ($H > 1/2$). The more general model is fully developed in MFC, CFM and FCM.

The self-similarity of $B_H(t)$ implies

Theorem 8. *The process $X(t)$ is a multifractal process with stationary increments, scaling function $\tau_X(q) = \tau_\theta(Hq)$, and multifractal spectrum $f_X(\alpha) = f_\theta(\alpha/H)$.*

²²Maheswaran and Sims (1992) suggest potential applications in finance for processes lying outside the class of semi-martingales.

²³See Baillie (1996) for a review of this literature.

The proof of these results is provided in MFC. We observe that $\tau_X(1/H) = \tau_\theta(1) = 0$, which allows the estimation of the index H in the empirical work. The generalized construction has scaling properties analogous to the model explored earlier, and provides a useful additional tool for empirical applications.

6. Empirical Evidence

6.1. Multifractal Moment Restrictions

Consider a price series $P(t)$ on the time interval $[0, T]$, and the log-price $X(t) \equiv \ln P(t) - \ln P(0)$. Partitioning $[0, T]$ into integer N intervals of length Δt , we define the *sample sum* or *partition function*

$$S_q(T, \Delta t) \equiv \sum_{i=0}^{N-1} |X(i\Delta t + \Delta t) - X(i\Delta t)|^q. \quad (6.1)$$

When $X(t)$ is multifractal, the addends are identically distributed, and the scaling law (2.5) yields $\mathbb{E}[S_q(T, \Delta t)] = N c_X(q) (\Delta t)^{\tau_X(q)+1}$ when the q^{th} moment exists. This implies

$$\ln \mathbb{E}[S_q(T, \Delta t)] = \tau_X(q) \ln(\Delta t) + c_X^*(q) \quad (6.2)$$

where $c_X^*(q) = \ln c_X(q) + \ln T$. For each admissible q , equation (6.2) provides testable moment conditions describing how the partition function varies with increment size Δt . Various methods can be used to estimate $\hat{\tau}_X(q)$ from the sample moments of the data. By (4.5), its Legendre transform $\hat{f}(\alpha)$ provides an estimate of the multifractal spectrum, and can then be mapped back into a distribution for the multipliers.

The scaling function $\tau_X(q)$ is specified either parametrically or nonparametrically. We can for instance choose a parametric family for the distribution of the multiplier M . The multifractal spectrum $f(\alpha)$ then belongs to a specific class of functions (Table 2), a constraint that can be imposed in estimation. On the other hand, a non-parametric approach places fewer restrictions on the underlying process. Since $\tau_\theta(q) = -\log_b \mathbb{E}(M^q) - 1$, the sample moments provide all the finite moments of M , and thus a great deal of information on its underlying distribution.²⁴

²⁴The distribution of M may not be uniquely determined by its moments. See Feller (1971) and Durrett (1991) for good discussions of the uniqueness problem in moments.

This paper uses a very simple estimation procedure. Given a set of positive moments q and time scales Δt , we calculate the partition functions $S_q(T, \Delta t)$ of the data. The partition functions are then plotted against Δt in logarithmic scales. By (6.2), the multifractal model implies that these plots should be approximately linear when the q^{th} moment exists. Regression estimates of the slopes then provide the corresponding scaling exponents $\hat{\tau}_X(q)$. This procedure will reveal striking visual evidence of moment-scaling in DM/USD data. Simulation experiments are then conducted to assess the joint performance of the multifractal model and the estimation methodology.

6.2. Deutsche Mark/US Dollar Exchange Rates

We begin by investigating the multifractality of the Deutsche Mark/US Dollar (“DM/USD”) exchange rate. We use two data sets provided by Olsen and Associates, a currency research and trading firm based in Zürich. The first data set (“daily”) consists of a twenty-four year series of daily data spanning June 1973 to December 1996. Olsen collects price quotes from banks and other institutions through several electronic networks. A price quote is converted to a single *price observation* by taking the geometric mean of the concurrent bid and ask. The reported price in the daily data is then calculated by linear interpolation of the price observations closest to 16:00 UK on each side.²⁵ Figure 3 shows the daily data, which exhibits volatility clustering at all time scales and intermittent large fluctuations.

The second data set (“high-frequency”) contains all bid/ask quotes and transmittal times collected over the one year period from October 31, 1992 to September 1, 1993. We convert quotes to price observations using the same methodology as Olsen, and obtain a round-the-clock data set of 1,472,241 observations. Olsen provides a flag for quotes believed to be erroneous or not representative of actual willingness to trade. We eliminate these observations, which constitute 0.36% of the dataset. Combining the daily data and the high-frequency data allows us to calculate partition functions over three orders of magnitude for Δt .

The high-frequency data show strong patterns of daily seasonality. In continuous time, seasonality is a smooth transformation that does not affect local Hölder exponents. Since our data is discrete, however, we may expect seasonality to introduce noise. To reduce this effect, we can write a seasonally modified version of

²⁵An earlier working paper (FCM) also uses noon buying rates provided by the Federal Reserve, and finds no significant difference in reported results.

the MMAR:

$$\ln P(t) - \ln P(0) = B_H \{ \theta [SEAS(t)] \},$$

where the seasonal transformation $SEAS(t)$ is a differentiable function of clock time. In this paper, we use a prefilter that smoothes variation in average absolute returns over fifteen minute intervals of the week. Except for the reduction in noise, there are no systematic differences in reported results for filtered and unfiltered data. An earlier working paper (FCM) provides details on this and three other seasonal prefilters, and finds small, predictable differences in results depending on the deseasonalizing method.

6.3. Main Results

Figures 4 and 5 illustrate the partition functions of the two DM/USD data sets. Values of Δt are chosen to increase multiplicatively by a factor of 1.1 from minimum to maximum. Since we focus on the slopes $\hat{\tau}_X(q)$ but not the intercepts, plots for each q are renormalized by vertical displacement to begin at zero for the lowest value of Δt in each graph. This allows plots for many q to be presented simultaneously. The daily and high frequency plots are presented in the same graph to highlight the similarity in their slopes. This is achieved by a second vertical displacement of the daily data that provides the best linear fit under OLS, restricting both lines to have the same slope.

Figure 4 shows the full range of calculated Δt , from fifteen seconds to six months, and five values of q ranging from 1.75 to 2.25. This allows us to estimate the self-affinity index H in the extended model presented in Section 5. Since $\tau_X(1/H) = 0$ and the standard Brownian specification $H = 1/2$ has previous empirical support, we expect to find $\tau_X(q) = 0$ for a value of q near two.

We first note the approximate linearity of the partition functions beginning at $\Delta t = 1.4$ hours and extending to the largest increment used, $\Delta t = 6$ months. In this range, the slope is zero for a value of q slightly smaller than two, and we report

$$\hat{H} \approx 0.53,$$

which implies very slight persistence in the DM/USD series. It is not immediately clear whether this result is sufficiently close to $H = 1/2$ to be consistent with the martingale version of the MMAR, but we will return to this issue in the following section using simulation methods.

The partition functions in Figure 4 also show breaks in linearity at high frequencies. These are consistent with microstructure effects such as bid-ask spreads, discreteness of quoting units, and discontinuous trading. In particular, these microstructure effects can be expected to induce a negative autocorrelation at high frequencies, as is well-understood in the case of bid-ask bounce (Roll, 1984). Negative autocorrelation effectively acts as an additional source of volatility, as previously explored in the variance ratio literature (e.g., Campbell and Mankiw, 1987; Lo and MacKinlay, 1988; Richardson and Stock, 1989; Faust, 1992). The results in Figure 4 are analogous to variance ratio tests, exactly so if we focus on the moment $q = 2$. As we move to the left on the graph and sampling frequency increases, microstructure induced negative autocorrelation increases, and the plots bend upwards corresponding to the increase in variability.

Descriptive statistics help to confirm that high frequency breaks in linearity are caused by microstructure effects. The departure from linearity begins at a frequency of approximately $\Delta t = 1.4$ hours, which is highlighted by the dotted line in Figure 4. We first note that the absolute change in the DM/USD rate averages 0.14 pfennig²⁶ over a time increment of 1.4 hours. Comparing this to the average spread of 0.07 pfennig,²⁷ we observe that the spread covers a significant proportion of average variation at this time horizon. It is thus sensible that microstructure effects begin to effect scaling properties at this frequency. To further confirm this intuition, observe that for time scales between 3.6 minutes and 1.4 hours, the partition function has approximate slope of zero for the moment $q = 2.25$. This implies $H \approx 0.44 < 1/2$, consistent with the explanation that microstructure induced negative autocorrelation causes the observed departure from scaling in this high frequency region.²⁸

There are several potential solutions to these high-frequency microstructure effects. First, we could view the MMAR as an underlying price process and overlay it with a structural model of discrete trade and bid-ask spreads that could then be related to the observed Olsen quotes. This would permit use of the data at all frequencies, but at the cost of introducing a new set of modelling issues. A second alternative is to use a reduced form model to prefilter the data to remove

²⁶One pfennig equals 0.01 *DM*.

²⁷The two most common spread sizes are 0.05 pfennig (38.25%), and 0.10 pfennig (52.55%), together comprising over 90% of all observed spreads.

²⁸To further confirm this explanation, one could compare across assets the frequency at which departures from moment scaling begin. If the microstructure explanation is correct, one would expect scaling to extend to higher frequencies when trading frictions (measured by variables such as bid-ask spreads or the average time interval between trades) are lower.

high frequency autocorrelations. This approach has recently been employed by Andersen, Bollerslev, Diebold, and Ebens (2000), who use an MA(1) filter to remove first order serial correlation from five-minute returns on the thirty stocks tracked in the Dow Jones Industrial Index.²⁹ While this would certainly remove high-frequency autocorrelations, the effect of this procedure on scaling properties would require further investigation. For simplicity, we choose to discard from further analysis all values of Δt less than 1.4 hours,³⁰ and have three orders of magnitude of sampling frequencies with which to test the scaling properties of the data.

With attention now restricted to values of Δt between 1.4 hours and 6 months, Figure 5 presents partition functions for a larger range of moments $1.5 \leq q \leq 5$. Higher moments capture information in the tails of the distribution of returns, and are thus generally more sensitive to deviations from scaling. All of the plots are nonetheless remarkably linear, and the overlapping values from the two data sets appear to have almost the same slope. Thus despite the apparent non-stationarity of the 24 year series, such as long price swings and long cycles of volatility, the moment restrictions imposed by the MMAR seem to hold over a broad range of sampling frequencies.

Estimates of the slopes in Figure 5 and additional moments q are then used to obtain estimated scaling functions $\hat{\tau}_X(q)$ for both data sets. We note the increasing variability of the partition function plots with the time scale Δt , which can be attributed to the shrinking number of addends in the partition function at low frequencies. This suggests a weighted least squares or generalized least squares approach. In practice, however, weighting the observations has little effect on the results because the plots are very nearly linear. Preferring simplicity, we thus report in Figure 6 the estimated scaling functions from OLS regressions. The estimated scaling functions are strictly concave, indicating multifractality, and are fairly similar except for very large moments.

Theorem 6 suggests to estimate the multifractal spectrum $f_X(\alpha)$ by taking the

²⁹In related work, Andersen, Bollerslev, Diebold, and Labys (2000) use five-minute DM/USD data from Olsen, but do not correct for first order autocorrelations. Our results suggest that their quadratic variation estimates could be biased upwards because of high-frequency autocorrelation.

³⁰As an approximation, the choice of $\Delta t = 1.4$ hours as the high-frequency cutoff for our analysis is justified by Figure 4, but the exact value was chosen by *ad hoc* rounding. In the future, there may be some gains to adopting a more formal test for the high-frequency cutoff value. There is already a substantial econometric literature (see, e.g., Andrews, 1993) that could be adapted to this purpose.

Legendre transform of $\widehat{\tau}_X(q)$. Following this logic, Figure 7 shows the estimated multifractal spectrum of the daily data.³¹ The estimated spectrum is concave, in contrast to the degenerate spectra of Brownian Motion and other unifractals. Using the estimated spectrum, we can recover a generating mechanism for trading time based on the canonical multiplicative cascades described in Section 2.2.

The spectrum of daily data is very nearly quadratic, and Section 4.2 has shown that quadratic spectra are generated by lognormally distributed multipliers M . We thus specify $-\log_b M \sim \mathcal{N}(\lambda, \sigma^2)$, giving trading time $\theta(t)$ with multifractal spectrum $f_\theta(\alpha) = 1 - (\alpha - \lambda)^2/[4(\lambda - 1)]$. The log-price process has most probable exponent $\alpha_0 = \lambda H$, and spectrum

$$f_X(\alpha) = 1 - \frac{(\alpha - \alpha_0)^2}{4H(\alpha_0 - H)}$$

Since $\widehat{H} = 0.53$, the free parameter α_0 is used to fit the estimated spectrum. We report

$$\widehat{\alpha}_0 = 0.589,$$

which produces the parabola shown in Figure 7. Choosing a generating construction with base $b = 2$,³² this immediately implies $\widehat{\lambda} = 1.11$ and $\widehat{\sigma}^2 = 0.32$.³³ It is also natural to consider the martingale version of the MMAR with the restriction $H = 1/2$. For this case, we estimate the single parameter $\widehat{\alpha}_0 = 0.545$.

In both cases, the estimated value of the most probable local Hölder exponent α_0 is greater than $1/2$. On a set of Lebesgue measure 1, the estimated multifractal process is therefore more regular than a Brownian Motion. However, the concavity of the spectrum also implies the existence of lower Hölder exponents that correspond to more irregular instants of the price process. These contribute disproportionately to volatility.

³¹The estimated multifractal spectrum of the high frequency data is similar in many respects, and is discussed in FCM.

³²The base b of the multifractal generating process is not uniquely identified by the spectrum alone, hence we assume the commonly used value $b = 2$. Calvet and Fisher (1999a) develop a likelihood based filter under which b can be estimated for the class of multinomial multifractals.

³³Mandelbrot (1989a, b) shows that the partition function methodology provides reasonable estimates of $\tau_X(q)$ only for moments $q < 1/\sqrt{\alpha_0(X)/H} - 1$, which is approximately equal to 5.66 in our estimated process.

6.4. Monte Carlo Simulations

We now present simulation experiments that provide a preliminary assessment of the new model and the estimation procedure. Figure 8 shows the levels and log-differences of a random price path generated by the limit lognormal MMAR estimated in Section 6.3.³⁴ The simulation shows a variety of large price changes, apparent trends, persistent bursts of volatility, and other characteristics found in the DM/USD series.

The following sections examine whether the inferred process captures the moment properties of the data. We sketch the simulation methodology, and then provide a synthetic discussion of the numerical results.

6.4.1. Methodology

We use three types of tests to analyze the MMAR's performance. First, visual evidence is provided on the moment properties of simulated data. Figure 9a thus illustrates the partition functions corresponding to four simulations of the MMAR. For comparison, we report in Figure 9b the partition functions of a GARCH(1, 1) process with the parameter estimates of Baillie and Bollerslev (1989). Figure 9c similarly considers the FIGARCH(1, d , 0) specification of Baillie, Bollerslev, and Mikkelsen (1996). Each plot in Figure 9 is based on a long sample of 100,000 observations. This sample length, which exceeds the sample size of 6,118 daily DM/USD returns, has the advantage of reducing the noisiness of the partition functions.

Second, we consider small samples and examine distributional evidence on the linearity and slopes of the partition functions. The analysis focuses on four processes: the extended MMAR (with arbitrary H), the martingale MMAR ($H = 1/2$), FIGARCH, and GARCH. These models are respectively indexed by $m \in \{1, \dots, 4\}$. For each model m , we simulate $J = 10,000$ paths with the same length $T = 6,118$ as the DM/USD data. We denote each path by $Y_j^m = \{Y_{j,t}^m\}_{t=1}^T$, ($1 \leq j \leq J$), and focus the analysis on the moments $q \in Q = \{0.5, 1, 2, 3, 5\}$. For each path and each q , an OLS regression provides a slope estimate $\hat{\tau}(q, Y_j^m)$ and the corresponding sum of squared errors $SSE(q, Y_j^m)$. The distributions of these statistics appear unimodal with smoothly declining tails. Tables 3 and 4 report the percentiles of these statistics.

We summarize these findings with several measures of global fit. For a given

³⁴The simulation of a multifractal path is discussed in Appendix 8.9.

model $m \in \{1, \dots, 4\}$, each path Y_j^m generates a column vector of slope and SSE estimates³⁵:

$$h(Y_j^m) = \{[\hat{\tau}(q, Y_j^m), \ln SSE(q, Y_j^m)]_{q \in Q}\}'.$$

It is convenient to denote the data by $X = \{X_t\}_{t=1}^T$, and to arrange the simulated paths in a $J \times T$ matrix $Y^m = [Y_1^m, \dots, Y_J^m]'$. We also consider

$$H(X, Y^m) = h(X) - \frac{1}{J} \sum_{j=1}^J h(Y_j^m),$$

The function H is useful to test how a particular model fits the moment properties of the data. In particular, we can define a global statistic $G = H'WH$ for any positive-definite matrix W . The empirical work considers four different weighting matrices W_m , $m \in \{1, \dots, 4\}$, each of which is obtained by inverting the simulated covariance matrix of moment conditions: $W_m = \left[\sum_{j=1}^J H(Y_j^m, Y^m) H(Y_j^m, Y^m)' / J \right]^{-1}$. The global statistics

$$G_{m,n}(X) = H(X, Y^m)' W_n H(X, Y^m), \quad m, n \in \{1, \dots, 4\}.$$

are indexed by the model m that generates the simulated data Y^m and the model n that generates the weighting matrix W_n . This gives a set of sixteen global statistics. Assuming that m is the true model, we can estimate the cumulative distribution function $\mathbb{F}_{m,n}$ of each statistic from the set $\{G_{m,n}(Y_j^m)\}_{1 \leq j \leq J}$, and then quantify the p -value $1 - \mathbb{F}_{m,n}[G_{m,n}(X)]$. The global statistics $G_{m,n}(X)$ and their associated p -values are reported in Table 5.

6.4.2. Results

The simulation results in Figure 9 and Tables 3 – 5 confirm that the MMAR replicates the main scaling features of the data. The partition function plots in Figure 9a are approximately linear and tend to follow their theoretically predicted slopes, which are nearly identical to the estimates from the DM/USD data. Tables 3 and 4 permit a more detailed assessment. The extended MMAR is very close to the the data in both its theoretically predicted slopes τ_0 and the mean slopes

³⁵We use the logarithm of the SSE in calculating the global statistics because Table 4 shows that the SSE are heavily right-skewed.

7. Furthermore, the estimated slopes from the DM/USD data are well within the central bells of the simulated slope distributions generated by the extended MMAR.

The martingale version is subtly different in both its theoretically predicted and mean slopes. For low moments, the estimated slopes from the DM/USD data are more towards the upper tails of the simulated distributions generated by the martingale MMAR. In both cases, but more so for the martingale version, there appears to be a slight downward bias in the average simulated slope relative to its theoretical value. Future work may thus correct the bias in our estimation method by matching simulated moments (Ingram and Lee, 1991; Duffie and Singleton, 1993).

Table 4 analyzes the variability of the simulated partition functions around their slopes. The extended and martingale versions of the MMAR yield nearly identical results. For low moments, the data falls well within the likely range of the SSE statistic for both models. For high moments, the partition functions are typically more variable for the simulated MMAR than for the data. This at first seems curious, because the model has been designed specifically to produce scaling. It is consistent, however, with the finding of a slight downward bias in the slopes of the estimated partition functions. Correcting this bias would give a slightly milder multifractal process, and thus reduce the variability of the partition functions. This presents a promising avenue for improving estimation. Overall, these results suggest that the MMAR is successful in matching the main scaling features of the data.

Another important question is whether other standard econometric models possess scaling properties. Figure 9b shows that GARCH(1, 1) partition functions are fairly linear, but their apparent slope is similar to the predicted slope of Brownian Motion rather than the data. This is symptomatic of the fact that GARCH models are short memory processes. Over long time periods, temporal clustering disappears and GARCH scales like a Brownian Motion. Tables 3 and 4 confirm this visual evidence. The SSE statistics show that GARCH tends to be as linear as the data, but for two of the five moments, the slopes from the data are in the extreme tails of their distributions simulated under GARCH. Because it contains long-memory in volatility, FIGARCH can be expected to scale differently than Brownian Motion at low frequencies. This is confirmed in Figure 9c; however, the same plots suggest that simulated FIGARCH partition functions are more irregular than the scaling plots generated by GARCH, the MMAR, and the data. Tables 3 and 4 again complement this visual evidence. The simulated FIGARCH

slopes improve over the GARCH slopes, but are not as close to the data as the MMAR. The SSEs from the data are also far in the tails of their distributions generated under FIGARCH.

The previous analysis has separately assessed ten moment conditions that capture different scaling features of the data. We now consider the evidence provided by the global statistics, which are quadratic functions of these ten moment conditions. Each column of the results in Table 5 is obtained by a different weighting of the set of quadratic terms, so that within column comparisons provide four separate views of ability to fit scaling features of the data. Each of the weighting matrices of course has different power against a given model, and asymptotic theory suggests that the most powerful weighting matrix for each model is provided by the inverse of its own covariance matrix of moment conditions. Thus, we expect the diagonal entries of the table to provide the greatest power to reject each model, and this is consistent with our results. Whether evaluated column-wise or by the diagonal elements, the results confirm that the MMAR is best able to replicate the scaling properties of the data.

These simulations are of course only a preliminary step to evaluating the usefulness of the multifractal model. Nonetheless, our results demonstrate that scaling properties contain important information for estimating and discriminating between models. A natural path for future work will be to incorporate this information in broader estimation and testing procedures. As these techniques develop (e.g., Calvet and Fisher, 1999a) and lead to more rigorous evaluations, it will be interesting to discover whether the promise of these early simulation results is fulfilled.

6.5. Equity Data

After observing multifractal properties in DM/USD exchange rates, it is natural to test the model on other financial data. This section presents evidence of moment-scaling in a sample of five major U.S. stocks and one equity index.³⁶

The Center for Research in Security Prices (CRSP) provides daily stock returns for 9,190 trading days from July 1962 to December 1998. We present results for the value weighted NYSE-AMEX-NASDAQ index (“CRSP Index”) and five

³⁶The multifractal model offers a flexible framework that may be amenable to many types of financial prices. In particular, equity data requires additional consideration for the relationship between volatility and mean returns. Since this has not been incorporated in the version of the model presented in this paper, the results in this section should be interpreted as a model-free investigation of scaling.

stocks: Archer Daniels Midland (ADM), General Motors (GM), Lockheed-Martin, Motorola, and United Airlines (UAL). The individual stocks are issued by large, well-known corporations from various economic sectors, and have reported data for the full CRSP sample span.³⁷ For each series, we convert the daily return data into a renormalized log-price series X_t , and then apply the partition function methodology described in Section 6.2.³⁸

Figure 10 shows results for the CRSP index and GM. In the first two panels, the full data sets are used with increments Δt ranging from one day to approximately one year. The partition functions for moments $q = \{1, 2, 3\}$ are approximately linear for both series, with little variation around the apparent slope. The slope for the moment $q = 2$ is noticeably positive for the CRSP index, indicating persistence. This characteristic is very atypical of individual securities, although short horizon persistence has been observed previously in index returns, and is typically attributed to asynchronous trading (e.g., Boudoukh, Richardson, and Whitelaw, 1993). In contrast to the results for low moments, the partition functions from both series vary considerably for the moment $q = 5$. This suggests investigation of the tails of the data. We find that the behavior of the fifth moment is dominated by volatility surrounding the stock market crash of October 1987. This is demonstrated by the second two panels of Figure 10, which show striking linearity after simply removing the crash day from both data sets.

Since discarding outliers seems an unsatisfactory approach to volatility modelling,³⁹ we reexamine the full data sets. The partition functions $S_{q=5}(T, \Delta t)$ for both series drop considerably from $\Delta t = 2$ days to $\Delta t = 3$ days. In the raw data, the CRSP index falls 17% on the day of the crash, but rebounds more than 8% two days later. GM loses 21% in the crash, recovering almost all its losses over the next two days. When $\Delta t = 3$ days, aggregation of these returns within a single interval contributes to the severe decline in the $q = 5$ partition functions.⁴⁰ As

³⁷Choosing stocks with full samples allows testing of the moment-scaling restrictions over a larger range of frequencies.

³⁸The CRSP holding period returns $r_t = (P_t - P_{t-1} + d_t)/P_{t-1}$ include cash distributions d_t . We construct the series $\{X_t\}_{t=0}^T$ by $X_0 = 0$, $X_t = X_{t-1} + \ln(1 + r_t)$.

³⁹While discarding data may be justified in specific circumstances, our approach in this paper has been to build a stationary model flexible enough to accommodate a wide range of changing economic circumstances. This includes both long-range structural shifts, and extreme tail events such as the 1987 crash.

⁴⁰When calculating the partition functions for $\Delta t = 2$ days, the crash day and the two following days belong to separate intervals, each of which make large contributions to $S_5(T, \Delta t = 2 \text{ days})$. When $\Delta t = 3$ days, however, the crash and the two following days belong to a

Δt grows, the crash and the two following days occasionally fall in separate intervals and the partition functions spike. More often, however, these three days lie within a single interval when Δt is large. Moreover, when a spike does occur for this reason with large Δt , its size is smaller because of the diminishing influence of the crash at low frequencies. This discussion explains why the partition functions $S_5(T, \Delta t)$ show large variability, but primarily for low values of Δt . Because of this variability, the estimated slopes of the partition function would appear to be relatively imprecise.

After removing the crash from both data sets, the same partition functions appear to give a more precise fit, but at a cost. Both slopes increase to appear more Brownian or “mild,” suggesting that important information has been lost. Additionally, removing the crash does not necessarily improve the fit of the MMAR since the theoretically predicted slopes constrain only the expectations of the partition functions, not their variability.⁴¹ In fact, the simulations in the previous section indicate that multifractal paths often have partition functions that vary considerable around their expected slope. Thus, removing the crash gives a false impression of improved model fit and alters scaling properties to imply a much milder process.

The other four stocks in our sample scale remarkably well despite the crash, as shown in Figure 11. Consistent with the martingale hypothesis for returns, three of the four stocks have almost exactly flat partition functions for $q = 2$, while ADM has a slight negative slope. The difference between Brownian scaling and multiscaling becomes perceptible for $q = 3$, and for the fifth moment, this difference is pronounced. UAL appears the most variable, with lower slopes at higher moments and thus a wider multifractal spectrum.

While not exhaustive, our empirical analysis indicates that moment-scaling is a prominent feature of many financial series. Using DM/USD data, we confirm this property across three orders of magnitude of frequencies and twenty-three years of daily returns. A simple estimation procedure helps to provide a specification of the multifractal model that reproduces scaling patterns found in the data. Finally, our analysis of equity data shows that the partition function plots summarize a great deal of information in a convenient form. This new tool may thus be useful in uncovering empirical regularities and building new financial models.

single interval. Their returns partially cancel when aggregated, explaining the decreases in the partition functions $S_5(T, \Delta t)$ at $\Delta t = 3$ days.

⁴¹The simulations in the previous section do, however, suggest that the variability of partition function plots can usefully be incorporated into estimation.

7. Conclusion

This paper has investigated the Multifractal Model of Asset Returns, a continuous time stochastic process that incorporates the outliers and volatility persistence exhibited by many financial time series. The model compounds a standard Brownian Motion with an independent multifractal time-deformation process that produces volatility clustering. We show how to construct a class of candidate time-deformations as the limit of a simple iterative procedure, called a multiplicative cascade. The cascade provides parsimonious modelling, and results in a generalized scaling rule that restricts return moments to vary as power laws of the time increment. The price process is a semi-martingale with uncorrelated returns, and thus precludes arbitrage in a standard two-asset setting.

The MMAR offers a fundamentally new class of processes to both finance and mathematics. Multifractal processes have continuous sample paths, but lie outside the class of Itô diffusions. Whereas standard processes can be characterized by a single local scale that describes the local growth rate of variation, sample paths of multifractal processes contain a *continuum* of local Hölder exponents within any time interval. The distribution of these exponents is conveniently quantified by a renormalized density, the multifractal spectrum $f(\alpha)$. For a large class of multifractal processes, the spectrum can be explicitly derived from Cramér's Large Deviation Theory. We demonstrate through a number of examples the sensitivity of the multifractal spectrum to the generating mechanism. The applied researcher may thus relate an empirical estimate of the spectrum back to a particular construction of the process, and is permitted considerable flexibility in modelling different types of data.

We find evidence of multifractality in the moment-scaling behavior of Deutsche Mark/US Dollar exchange rates. Over a range of observational frequencies from approximately two hours to 180 days, and over a range of time from 1973 to 1996, moments of the data grow approximately like a power law. We obtain an estimate of the multifractal spectrum by a Legendre transform of the moments' growth rates. From the shape of the estimated spectrum, we infer a lognormal distribution as the primitive of the generating mechanism, and estimate its parameters. We simulate the process, and confirm that the multifractal model replicates the moment behavior found in the data. We also demonstrate scaling behavior in an equity index and five major U.S. stocks.

Our results indicate several directions for future research. Using our simulation results as a guide, the moment-scaling features of the data can be incorporated

into broader estimation and testing procedures. Risk analysis, forecasting, and option pricing are promising applications that are currently being developed in other papers. Further research will also seek to derive the MMAR as an equilibrium process of economies with fully rational agents. In such frameworks, multifractality is expected to arise in equilibrium either exogenously, for instance as a consequence of multifractal technological shocks, or endogenously because of market incompleteness or informational cascades. The early empirical success of the MMAR thus offers new challenges in econometrics, finance, and economic theory.

8. Appendix

8.1. Scaling Rule

This Appendix analyzes the set \mathcal{D} defined by a multiplicative measure with parameter $b \geq 2$. Consider a fixed instant $t \in [0, 1]$. For all $\varepsilon > 0$, there exists a dyadic number t_n such that $|t_n - t| < \varepsilon$. We can then find a number $\Delta_n = b^{-k_n} < \varepsilon$ for which $(t_n, \Delta_n) \in \mathcal{D}$. In the plane \mathbb{R}^2 , the point $(t, 0)$ is thus the limit of the sequence $(t_n, \Delta_n) \in \mathcal{D}$, which establishes

Property 1. *The closure of \mathcal{D} contains the set $[0, 1] \times \{0\}$.*

The scaling relation (2.4) thus holds “in the neighborhood of any instant”.

8.2. Proof of Proposition 1

Consider two exponents q_1, q_2 , and two positive weights w_1, w_2 adding up to one. Hölder’s inequality implies

$$\mathbb{E}(|X(t)|^q) \leq [\mathbb{E}(|X(t)|^{q_1})]^{w_1} [\mathbb{E}(|X(t)|^{q_2})]^{w_2},$$

where $q = w_1 q_1 + w_2 q_2$. Taking logarithms and using (2.5), we obtain

$$\ln c(q) + \tau(q) \ln t \leq [w_1 \tau(q_1) + w_2 \tau(q_2)] \ln t + [w_1 \ln c(q_1) + w_2 \ln c(q_2)]. \quad (8.1)$$

We divide by $\ln t < 0$, and let t go to zero:

$$\tau(q) \geq w_1 \tau(q_1) + w_2 \tau(q_2), \quad (8.2)$$

which establishes the concavity of τ . This proof also contains additional information on multifractal processes. Assuming that relation (2.5) holds for $t \in [0, \infty)$, we divide inequality (8.1) by $\ln t > 0$ and let t go to infinity. We obtain the reverse of inequality (8.2), and conclude that $\tau(q)$ is linear. Thus exact multiscaling can only hold for bounded time intervals \mathcal{T} .

8.3. Proof of Theorem 1

Since the trading time and the Brownian motion $B(t)$ are independent, conditioning on $\theta(t)$ yields

$$\begin{aligned} \mathbb{E} \{ |X(t)|^q \mid \theta(t) = u \} &= \mathbb{E} [|B(u)|^q \mid \theta(t) = u] \\ &= \theta(t)^{q/2} \mathbb{E} [|B(1)|^q], \end{aligned}$$

and thus $\mathbb{E}[|X(t)|^q] = \mathbb{E}[\theta(t)^{q/2}] \mathbb{E}[|B(1)|^q]$. The process $X(t)$ satisfies the multiscaling relation (2.5), with $\tau_X(q) \equiv \tau_\theta(q/2)$ and $c_X(q) \equiv c_\theta(q/2) \mathbb{E}[|B(1)|^q]$.

8.4. Proof of Theorem 2

Let \mathcal{F}_t and \mathcal{F}'_t denote the natural filtrations of $\{X(t)\}$ and $\{X(t), \theta(t)\}$. For any t, T, u , the independence of B and θ implies

$$\begin{aligned} \mathbb{E}\{X(t+T) \mid \mathcal{F}'_t, \theta(t+T) = u\} &= \mathbb{E}\{B(u) \mid \mathcal{F}'_t\} \\ &= B[\theta(t)], \end{aligned}$$

since $\{B(t)\}$ is a martingale. We now infer that $\mathbb{E}[X(t+T) \mid \mathcal{F}_t] = X(t)$, which establishes that $X(t)$ is a martingale and has thus uncorrelated increments. The price $P(t)$ is a smooth function of $X(t)$ and therefore a semi-martingale, which precludes arbitrage opportunities in the two asset economy.

8.5. Proof of Theorem 4

1. Trading Time

Consider a canonical cascade after $k \geq 1$ stages. Consistent with the notation of Section 2, the interval $[0, T]$ is partitioned into cells of length $\Delta t = b^{-k}T$, and $I_1 = [t_1, t_1 + \Delta t]$ and $I_2 = [t_2, t_2 + \Delta t]$ denote two distinct cells with lower endpoints of the form $t_1/T = \overline{0.\eta_1 \dots \eta_k}$ and $t_2/T = \overline{0.\zeta_1 \dots \zeta_k}$. Assume that the first $l \geq 1$ terms are equal in the b -adic expansions of t_1/T and t_2/T , so that $\zeta_1 = \eta_1, \dots, \zeta_l = \eta_l$, and $\zeta_{l+1} \neq \eta_{l+1}$. The distance $t = |t_2 - t_1|$ satisfies $b^{-l-1} \leq t/T < b^{-l}$, and the product $\mu(I_1)^q \mu(I_2)^q$, which is equal to

$$\begin{aligned} &\Omega_{\eta_1, \dots, \eta_k}^q \Omega_{\zeta_1, \dots, \zeta_k}^q (M_{\eta_1}^{2q} \dots M_{\eta_l, \dots, \eta_l}^{2q}) \\ &(M_{\eta_1, \dots, \eta_{l+1}}^q \dots M_{\eta_1, \dots, \eta_k}^q) (M_{\zeta_1, \dots, \zeta_{l+1}}^q \dots M_{\zeta_1, \dots, \zeta_k}^q), \end{aligned}$$

has mean $(\mathbb{E}\Omega^q)^2 [\mathbb{E}M^{2q}]^l [\mathbb{E}M^q]^{2(k-l)}$. We conclude that

$$\begin{aligned} Cov[\mu(I_1)^q; \mu(I_2)^q] &= (\mathbb{E}\Omega^q)^2 (\mathbb{E}M^q)^{2k} \left\{ \left[(\mathbb{E}M^{2q}) / (\mathbb{E}M^q)^2 \right]^l - 1 \right\} \\ &= C_1 (\Delta t)^{2\tau_\theta(q)+2} \left[b^{-l[\tau_\theta(2q)-2\tau_\theta(q)-1]} - 1 \right] \end{aligned}$$

is bounded by two hyperbolic functions of t .

2. Log-Price

Since $B(t)$ and $\theta(t)$ are independent processes, the conditional expectation

$$\mathbb{E}\{|X(0, \Delta t)X(t, \Delta t)|^q | \theta(\Delta t) = u_1, \theta(t) = u_2, \theta(t + \Delta t) = u_3\}, \quad (8.3)$$

simplifies to

$$\mathbb{E}[|B(u_1)|^q] \mathbb{E}[|B(u_3) - B(u_2)|^q] = |u_1|^{q/2} |u_3 - u_2|^{q/2} [\mathbb{E}|B(1)|^q]^2.$$

Taking expectations, we infer that

$$\mathbb{E}[|X(0, \Delta t)X(t, \Delta t)|^q] = \mathbb{E}\left[|\theta(0, \Delta t)\theta(t, \Delta t)|^{q/2}\right] [\mathbb{E}|B(1)|^q]^2$$

and therefore $\delta_X(t, q) = \delta_\theta(t, q/2) [\mathbb{E}|B(1)|^q]^2$.

8.6. Interpretation of $f(\alpha)$ as a Fractal Dimension

Fractal geometry considers irregular and winding structures that are not well described by their Euclidean length. For instance, a geographer measuring the length of a coastline will find very different results as she increases the precision of her measurement. In fact, the structure of the coastline is usually so intricate that the measured length diverges to infinity as the geographer's measurement scale goes to zero. For this reason, we cannot use the Euclidean length to compare two different coastlines, and it is natural to introduce a new concept of dimension. Given a precision level $\varepsilon > 0$, we consider coverings of the coastline with balls of diameter ε . Let $N(\varepsilon)$ denote the smallest number of balls required for such a covering. The approximate length of the coastline is defined by $L(\varepsilon) = \varepsilon N(\varepsilon)$. In many cases, $N(\varepsilon)$ satisfies a power law as ε goes to zero:

$$N(\varepsilon) \sim \varepsilon^{-D},$$

where D is a constant called the *fractal dimension*. Fractal dimension helps to analyze the structure of a fixed multifractal. For any $\alpha \geq 0$, we can define the set $T(\alpha)$ of instants with Hölder exponent α . As any subset of the real line, $T(\alpha)$ has a fractal dimension $D(\alpha)$, which satisfies $0 \leq D(\alpha) \leq 1$. It can be shown that for a large class of multifractals, the dimension $D(\alpha)$ coincides with the multifractal spectrum $f(\alpha)$.

In the case of measures, we can provide a heuristic interpretation of this result based on coarse Hölder exponents. Denoting by $N(\alpha, \Delta t)$ the number of intervals

$[t, t + \Delta t]$ required to cover $T(\alpha)$, we infer from Equation (4.1) that: $N(\alpha, \Delta t) \sim (\Delta t)^{-f(\alpha)}$. We then rewrite the total mass $\mu[0, T] = \sum \mu(\Delta t) \sim \sum (\Delta t)^{\alpha(t)}$, and rearrange it as a sum over Hölder exponents:

$$\mu[0, T] \sim \int (\Delta t)^{\alpha - f(\alpha)} d\alpha.$$

The integral is dominated by the contribution of the Hölder exponent α_1 that minimizes $\alpha - f(\alpha)$, and therefore

$$\mu[0, T] \sim (\Delta t)^{\alpha_1 - f(\alpha_1)}$$

Since the total mass $\mu[0, T]$ is positive, we infer that $f(\alpha_1) = \alpha_1$, and $f(\alpha) \leq \alpha$ for all α . When f is differentiable, the coefficient α_1 also satisfies $f'(\alpha_1) = 1$. The spectrum $f(\alpha)$ then lies under the 45° line, with tangential contact at $\alpha = \alpha_1$.

8.7. Large Deviation Theory and the Multifractal Spectrum

This Appendix sketches the proof of Theorem 6, and introduces the concepts of latent and virtual Hölder exponents.⁴² First consider a *conservative* multiplicative measure μ . Application of Large Deviation Theory (LDT) begins with the histogram method of Section 4.1. Subdivide the range of α s into intervals of length $\Delta\alpha$, and denote by $N_k(\alpha)$ the number of coarse Hölder exponents in the interval $(\alpha, \alpha + \Delta\alpha]$. For large values of k , we write

$$\frac{1}{k} \log_b \left[\frac{N_k(\alpha)}{b^k} \right] \sim \frac{1}{k} \log_b \mathbb{P} \{ \alpha < \alpha_k \leq \alpha + \Delta\alpha \}. \quad (8.4)$$

This relation holds exactly for multinomial measures, which have discrete coarse exponents α_k , but is postulated in more general cases. For any $\alpha > \alpha_0$, Cramér's theorem implies

$$k^{-1} \log_b \mathbb{P} \{ \alpha_k > \alpha \} \rightarrow \text{Inf}_q \log_b \left[\mathbb{E} e^{q(\alpha - V_1) \ln b} \right] \quad (8.5)$$

as $k \rightarrow \infty$. Using the definition of the scaling function, we simplify the limit to $\text{Inf}_q [\alpha q - \tau(q)] - 1$. Combining this with (4.1) and (8.4), it follows⁴³ that Theorem 6 holds.

⁴²We refer the reader to Mandelbrot (1989b), Peyrière (1991) and CFM for more detailed discussions.

⁴³See CFM for a more detailed proof.

These arguments easily extend to a *canonical* measure μ . Given a b -adic instant t , the coarse exponent $\alpha_k(t) = \ln \mu[t, t + \Delta t] / \ln \Delta t$ is the sum of a high frequency component, $-k^{-1} \log_b \Omega_{\eta_1, \dots, \eta_k}$, and of the familiar low frequency average

$$\alpha_{k,L}(t) = -[\log_b M_{\eta_1} + \dots + \log_b M_{\eta_1, \dots, \eta_k}] / k.$$

The exponent $\alpha_k(t)$ converges almost surely to $\alpha_0 = -\mathbb{E} \log_b M$, and the multifractal spectrum is again the Legendre transform of the scaling function $\tau(q)$.

Relation (8.5) also shows that $f(\alpha)$ is the limit of

$$\begin{aligned} k^{-1} \log_b \mathbb{P} \{ \alpha_{k,L}(t) > \alpha \} + 1 & \quad \text{if } \alpha > \alpha_0, \quad \text{and} \\ k^{-1} \log_b \mathbb{P} \{ \alpha_{k,L}(t) < \alpha \} + 1 & \quad \text{if } \alpha < \alpha_0. \end{aligned}$$

$f(\alpha)$ is therefore a hump-shaped function, reaching a maximum at the most probable exponent: $f(\alpha) \leq f(\alpha_0) = 1$.⁴⁴ We have successively viewed the spectrum $f(\alpha)$ as:

- (D1) the limit of a renormalized histogram of coarse Hölder exponents,
- (D2) the fractal dimension of the set of instants with Hölder exponent α ,
- (D3) the limit of $k^{-1} \log_b \mathbb{P} \{ \alpha_{k,L}(t) > \alpha \} + 1$ provided by LDT.

The three definitions coincide for multinomial measures, and (D1) and (D2) agree for a large class of multifractals (Peyrière, 1991). However, (D1) and (D2) imply that $f(\alpha) \geq 0$, while (D3) imposes no such restriction. When $f(\alpha) < 0$, the corresponding α s, called *latent*, are rare coarse exponents, which appear in few draws of the random measure and control high and low moments (Mandelbrot, 1989b). Similarly, since canonical measures allow M to be greater than 1, the low-frequency average $\alpha_{k,L}(t)$ can be negative with positive probability. (D3) thus defines the multifractal spectrum for negative, or *virtual*, values of α . This topic, further discussed in Mandelbrot (1989b), remains an active research area in mathematics.

⁴⁴It is easy to show that $\alpha_0 q - \tau(q)$ is minimal for $q = 0$. The set $T(\alpha_0)$ has therefore fractal dimension $f(\alpha_0) = -\tau(0) = 1$, and thus carries all of the Lebesgue measure. Moreover by the Central Limit Theorem, $f(\alpha)$ is locally quadratic around α_0 , as shown in CFM.

8.8. Proof of Theorem 7

Given a process Z , denote $\alpha_Z(t)$ as its local scale at date t , and $T_Z(\alpha)$ as the set of instants with scale α . At any date, the infinitesimal variation of the log-price $X(t + \Delta t) - X(t) = B[\theta(t + \Delta t)] - B[\theta(t)]$ satisfies

$$|X(t + \Delta t) - X(t)| \sim |\theta(t + \Delta t) - \theta(t)|^{1/2} \sim |\Delta t|^{\alpha_{\theta}(t)/2},$$

implying $\alpha_X(t) \equiv \alpha_{\theta}(t)/2$. The sets $T_X(\alpha)$ and $T_{\theta}(2\alpha)$ coincide, and in particular have identical fractal dimensions: $f_X(\alpha) \equiv f_{\theta}(2\alpha)$. Moreover since the price $P(t)$ is a differentiable function of $X(t)$, the two processes have identical local Hölder exponents and spectra.

8.9. Simulation of Multifractal Paths

The construction of the simulated multifractal price paths used in Section 6.4. has two basic components. First, a finite stage approximation to a canonical multifractal measure is constructed as suggested in Section 2.2. All simulations use the base $b = 2$, so that if a simulation of length T is desired, we choose the minimum integer number of stages k such that $2^k \geq T$. At each stage in the construction, we draw independent lognormal multipliers with identical distributions given by the results in Section 6.3. When k stages are completed, the measure μ_k is used as a discrete approximation to the quadratic variation of a multifractal path. Aggregating the increments of μ_k thus provides a simulated path from the trading time $\theta(t)$. The second part of the construction involves compounding, as suggested by Assumptions 1-3 of Section 3. For the martingale version of the MMAR, we simply calculate the standard deviation $[\mu_k(\Delta t)]^{1/2}$ over the discrete time interval Δt and multiply by an independent standard Gaussian. To simulate the extended MMAR, we first generate a discretized path from a FBM with parameter \hat{H} taken from the estimates in Section 6.3. Interpolation provides values of the path $B_H[\theta(t)]$ at the simulated values from the path $\theta(t)$.

References

- [1] Adelman, I. (1965), Long Cycles: Fact or Artefact?, *American Economic Review* **55**, 444-463.
- [2] Andersen, T. and T. Bollerslev (1997), Heterogeneous Information Arrivals and Return Volatility Dynamics: Uncovering the Long-Run in High Frequency Returns, *Journal of Finance* **52**, 975-1005.
- [3] Andersen, T., Bollerslev, T., Diebold, F., and Ebens, H. (2000), The Distribution of Stock Return Volatility, forthcoming in *Journal of Financial Economics*.
- [4] Andersen, T., Bollerslev, T., Diebold, F., and Labys, P. (2000), The Distribution of Realized Exchange Rate Volatility, forthcoming in *Journal of American Statistical Association*.
- [5] Andrews, D. (1993), Tests for Parameter Instability and Structural Change with Unknown Change Point, *Econometrica* **61**, 821-856.
- [6] Avery, C., and Zemsky, P. (1998), Multidimensional Uncertainty and Herd Behavior in Financial Markets, *American Economic Review* **88**, 724-748.
- [7] Backus, D. K., and Zin, S. E. (1993), Long Memory Inflation Uncertainty: Evidence from the Term Structure of Interest Rates, *Journal of Money, Credit and Banking* **25**, 681-700.
- [8] Baillie, R. T. (1996), Long Memory Processes and Fractional Integration in Econometrics, *Journal of Econometrics* **73**, 5-59.
- [9] Baillie, R. T., and Bollerslev, T. (1989), The Message in Daily Exchange Rates: A Conditional Variance Tale, *Journal of Business and Economic Statistics* **7**, 297-305.
- [10] Baillie, R. T., Bollerslev, T., and Mikkelsen, H. O. (1996), Fractionally Integrated Generalized Autoregressive Conditional Heteroskedasticity, *Journal of Econometrics* **74**, 3-30.
- [11] Baillie, R. T., Chung, C. F., and Tieslau, M. A. (1996), Analyzing Inflation by the Fractionally Integrated ARFIMA-GARCH Model, *Journal of Applied Econometrics* **11**, 23-40.

- [12] Bates, D. S. (1995), Testing Option Pricing Models, *NBER Working Paper No. 5129*.
- [13] Bates, D. S. (1996), Jumps and Stochastic Volatility: Exchange Rate Process Implicit in Deutsche Mark Options, *Review of Financial Studies* **9**, 69-107.
- [14] Beveridge, W. H. (1925), Weather and Harvest Cycles, *Economic Journal* **31**, 429-452.
- [15] Bikhchandani, S., Hirshleifer, D., and Welch, I. (1992), A Theory of Fads, Fashion, Custom and Cultural Change as Informational Cascades, *Journal of Political Economy* **100**, 992-1027.
- [16] Billingsley, P. (1979), *Probability and Measure*, New York: John Wiley and Sons.
- [17] Bochner, S. (1955), *Harmonic Analysis and the Theory of Probability*, Berkeley: University of California Press.
- [18] Bollerslev, T. (1986), Generalized Autoregressive Conditional Heteroskedasticity, *Journal of Econometrics* **31**, 307-327.
- [19] Bollerslev, T. (1987), A Conditional Heteroskedastic Time Series Model for Speculative Prices and Rates of Return, *Review of Economics and Statistics* **69**, 542-547.
- [20] Boudoukh, J., Richardson, M., and Whitelaw, R. (1994), A Tale of Three Schools: Insights on Autocorrelations of Short-Horizon Returns, *Review of Financial Studies* **7**, 539-573.
- [21] Briedt, J., Crato, N., and DeLima, P. (1998), On the Detection and Estimation of Long-Memory in Stochastic Volatility, forthcoming, *Journal of Econometrics*.
- [22] Bulow, J., and Klemperer, P. (1994), Rational Frenzies and Crashes, *Journal of Political Economy* **102**, 1-24.
- [23] Calvet, L. (1997), Incomplete Markets and Volatility, forthcoming in *Journal of Economic Theory*.
- [24] Calvet, L., and Fisher, A. (1999a), Forecasting Multifractal Volatility, forthcoming in *Journal of Econometrics*.

- [25] Calvet, L., and Fisher, A. (1999*b*), Catastrophic Risk in Multifractal Returns, forthcoming in *Integrated Risk and Return Management for Insurance Companies*, E. Altman and I. Vanderhoof eds., Kluwer Academic Press.
- [26] Calvet, L., Fisher, A., and Mandelbrot, B. B. (1997), Large Deviation Theory and the Distribution of Price Changes, *Cowles Foundation Discussion Paper No. 1165*, Yale University, Paper available from the SSRN database: <http://www.ssrn.com>.
- [27] Campbell, J., and Mankiw, N. G. (1987), Are Output Fluctuations Transitory?, *Quarterly Journal of Economics* **102**, 857-880.
- [28] Campbell, J., Lo, A., and MacKinlay, A. C. (1997), *The Econometrics of Financial Markets*, Princeton University Press.
- [29] Clark, P. K. (1973), A Subordinated Stochastic Process Model with Finite Variance for Speculative Prices, *Econometrica* **41**, 135-156.
- [30] Cox, J. C., Ingersoll, J. E., and S. A. Ross (1985), An Intertemporal General Equilibrium Model of Asset Prices, *Econometrica* **53**, 363-84.
- [31] Dacorogna, M. M., Müller, U. A., Nagler, R. J., Olsen, R. B., and Pictet, O. V. (1993), A Geographical Model for the Daily and Weekly Seasonal Volatility in the Foreign Exchange Market, *Journal of International Money and Finance* **12**, 413-438.
- [32] Diebold, F. X., and Rudebusch, G. D. (1989), Long Memory and Persistence in Aggregate Output, *Journal of Monetary Economics* **24**, 189-209.
- [33] Ding, Z., Granger, C. W. J., and Engle, R. F. (1993), A Long Memory Property of Stock Returns and a New Model, *Journal of Empirical Finance* **1**, 83-106.
- [34] Dothan, M. (1990), *Prices in Financial Markets*, Oxford University Press.
- [35] Drost, F. C., and Nijman, T. E. (1993), Temporal Aggregation of GARCH Processes, *Econometrica* **61**, 909-927.
- [36] Duffie, D., and Singleton, K. J. (1993), Simulated Moments Estimation of Markov Models of Asset Prices, *Econometrica* **61**, 929-952.

- [37] Durrett, R. (1991), *Probability: Theory and Examples*, Pacific Grove, Calif.: Wadsworth & Brooks/Cole Advanced Books & Software.
- [38] Engle, R. F. (1982), Autoregressive Conditional Heteroscedasticity with Estimates of the Variance of United Kingdom Inflation, *Econometrica* **50**, 987-1007.
- [39] Engle, R. F., and Gonzalez-Rivera, G. (1991), Semiparametric ARCH Models, *Journal of Business and Economic Statistics* **9**, 345-359.
- [40] Faust, J. (1992), When are Variance Ratio Tests For Serial Dependence Optimal?, *Econometrica* **60**, 1215-1226.
- [41] Fisher, A., Calvet, L., and Mandelbrot, B. B. (1997), Multifractality of Deutsche Mark/US Dollar Exchange Rates, *Cowles Foundation Discussion Paper No. 1166*, Yale University. Paper available from the SSRN database: <http://www.ssrn.com>.
- [42] Frisch, U., and Parisi, G. (1985), Fully Developed Turbulence and Intermittency, in: M. Ghil ed., *Turbulence and Predictability in Geophysical Fluid Dynamics and Climate Dynamics*, 84-88, Amsterdam: North-Holland.
- [43] Galluccio, S., Caldarelli, G., Marsili, M., and Zhang, Y. C. (1997), Scaling in Currency Exchange, *Physica A* **245**, 423-436.
- [44] Gennotte, G., and Leland, H. (1990), Market Liquidity, Hedging and Crashes, *American Economic Review* **80**, 999-1021.
- [45] Ghysels, E., Gouriéroux, C., and Jasiak, J. (1995), Market Time and Asset Price Movements: Theory and Estimation, *CIRANO and CREST Discussion Paper*.
- [46] Ghysels, E., Gouriéroux, C., and Jasiak, J. (1996), Trading Patterns, Time Deformation and Stochastic Volatility in Foreign Exchange Markets, *CREST Working Paper n° 9655*.
- [47] Ghysels, E., Harvey, A. C., and Renault, E. (1996), Stochastic Volatility, in: G. S. Maddala and C. R. Rao Eds., *Handbook of Statistics*, Vol. 14, 119-191. Amsterdam: North-Holland.

- [48] Glosten, L. Jagannathan, R., and Runkle, D. (1993), On the Relation Between the Expected Value and the Volatility of the Nominal Excess Return on Stocks, *Journal of Finance* **48**, 1779-1801.
- [49] Granger, C. W. J., and Joyeux, R. (1980), An Introduction to Long Memory Time Series Models and Fractional Differencing, *Journal of Time Series Analysis* **1**, 15-29.
- [50] Guivarc'h, Y. (1987), Remarques sur les Solutions d'une Equation Fonctionnelle Non Linéaire de Benoit Mandelbrot, *Comptes Rendus (Paris)* **3051**, 139.
- [51] Halsey, T. C., Jensen, M. H., Kadanoff, L. P., Procaccia, I., and Shraiman, B. I. (1986), Fractal Measures and their Singularities: The Characterization of Strange Sets, *Physical Review Letters A* **33**, 1141.
- [52] Harrison, J. M., and Kreps, D. (1979), Martingales and Arbitrage in Multi-period Securities Markets, *Journal of Economic Theory* **20**, 381-408.
- [53] Hosking, J. R. M. (1981), Fractional Differencing, *Biometrika* **68**, 165-176.
- [54] Ingram, B. F., and Lee, B. S. (1991), Simulation Estimation of Time Series Models, *Journal of Econometrics* **47**, 197-250.
- [55] Jacklin, C., Kleidon, A., and Pfleiderer, P. (1992), Underestimation of Portfolio Insurance and the Crash of October 1987, *Review of Financial Studies* **5**, 35-63.
- [56] Kahane, J. P. (1997), A Century of Interplay Between Taylor Series, Fourier Series and Brownian Motion, *Bulletin of the London Mathematical Society* **29**, 257-279.
- [57] Koedjik, K. G., and Kool, C. J. M. (1992), Tail Estimates of East European Exchange Rates, *Journal of Business and Economic Statistics* **10**, 83-96.
- [58] Lo, A. W. (1991), Long Memory in Stock Market Prices, *Econometrica* **59**, 1279-1313.
- [59] Lo, A. W., and MacKinlay, A. C. (1988), Stock Market Prices Do Not Follow Random Walks: Evidence from a Simple Specification Test, *Review of Financial Studies* **1**, 41-66.

- [60] Mahesweran, S., and Sims, C. (1993), Empirical Implications of Arbitrage-Free Asset Markets, in *Models, Methods and Applications of Econometrics*, P. C. B. Phillips, ed., Basil Blackwell 1993.
- [61] Mandelbrot, B. B. (1963), The Variation of Certain Speculative Prices, *Journal of Business* **36**, 394-419.
- [62] Mandelbrot, B. B. (1965), Une Classe de Processus Stochastiques Homothétiques à Soi, *Comptes Rendus de l'Académie des Sciences de Paris* **260**, 3274-3277. Reprinted in Mandelbrot (2000).
- [63] Mandelbrot, B. B. (1971), Statistical Dependence in Prices and Interest Rates, *Fifty First Annual Report of the National Bureau of Economic Research*, 141-142. Excerpted in Mandelbrot (2000).
- [64] Mandelbrot, B. B. (1972), Possible Refinements of the Lognormal Hypothesis Concerning the Distribution of Energy Dissipation in Intermittent Turbulence, in: M. Rosenblatt and C. Van Atta eds., *Statistical Models and Turbulence*, New York: Springer Verlag.
- [65] Mandelbrot, B. B. (1974), Intermittent Turbulence in Self-Similar Cascades: Divergence of High Moments and Dimension of the Carrier. *Journal of Fluid Mechanics* **62**, 331-358.
- [66] Mandelbrot, B. B. (1989a), Multifractal Measures, Especially for the Geophysicist, *Pure and Applied Geophysics* **131**, 5-42.
- [67] Mandelbrot, B. B. (1989b), Examples of Multinomial Multifractal Measures that Have Negative Latent Values for the Dimension $f(\alpha)$, in: L. Pietronero ed., *Fractals' Physical Origins and Properties*, 3-29, New York: Plenum.
- [68] Mandelbrot, B. B. (1997), *Fractals and Scaling in Finance: Discontinuity, Concentration, Risk*, New York: Springer Verlag.
- [69] Mandelbrot, B. B. (2000), *Gaussian Self-Affinity and Fractals*, New York: Springer Verlag.
- [70] Mandelbrot, B. B., Fisher, A., and Calvet, L. (1997), The Multifractal Model of Asset Returns, *Cowles Foundation Discussion Paper No. 1164*, Yale University. Paper available from the SSRN database: <http://www.ssrn.com>.

- [71] Mandelbrot, B. B., and Ness, J. W. van (1968), Fractional Brownian Motion, Fractional Noises and Application, *SIAM Review* **10**, 422-437.
- [72] Mandelbrot, B. B. and Taylor, H. W. (1967), On the Distribution of Stock Price Differences, *Operations Research* **15**, 1057-1062.
- [73] Mandelbrot, B. B., and Wallis, J. R. (1969), Some Long Run Properties of Geophysical Records, *Water Resources Research* **5**, 321-340. Reprinted in Mandelbrot (2000).
- [74] Merton, R. C. (1990), *Continuous-Time Finance*, Cambridge, Mass. : Blackwell.
- [75] Müller, U. A., Dacorogna, M. M., Davé, R. D., Pictet, O. V., Olsen, R. B., and Ward, J. R. (1995), *Fractals and Intrinsic Time: A Challenge to Econometricians*, Discussion Paper Presented at the 1993 International Conference of the Applied Econometrics Association held in Luxembourg.
- [76] Nelson, D. (1991), Conditional Heteroskedasticity in Asset Returns: A New Approach, *Econometrica* **45**, 7-38.
- [77] Pasquini, M., and Serva, M. (1999), Multiscale Behavior of Volatility Autocorrelations in a Financial Market, *Economics Letters* **65**, 275-279.
- [78] Pasquini, M., and Serva, M. (2000), Clustering of Volatility as a Multiscale Phenomenon, *European Physical Journal B* **16**, 195-201.
- [79] Peyrière, J. (1991), Multifractal Measures, in *Proceedings of the NATO ASI "Probabilistic Stochastic Methods in Analysis, with Applications"*.
- [80] Phillips, P. C. B., McFarland, J. W., and McMahon, P. C. (1994), Robust Tests of Forward Exchange Market Efficiency with Empirical Evidence from the 1920's, *Journal of Applied Econometrics* **11**, 1-22.
- [81] Richards, G. (2000), The Fractal Structure of Exchange Rates: Measurement and Forecasting, *Journal of International Financial Markets, Institutions and Money* **10**, 163-180.
- [82] Richardson, M., and Stock, J. (1989), Drawing Inferences from Statistics Based on Multi-Year Asset Returns, *Journal of Financial Economics* **25**, 323-348.

- [83] Roll, R. (1984), A Simple Implicit Measure of the Effective Bid-Ask Spread in an Efficient Market, *Journal of Finance* **39**, 1127-1140.
- [84] Samorodnitsky, G., and Taqqu, M. S. (1994), *Stable Non-Gaussian Random Processes*, New York: Chapman and Hall.
- [85] Sowell, F. B. (1992), Modeling Long Run Behavior with the Fractional ARIMA Model, *Journal of Monetary Economics* **29**, 277-302.
- [86] Stock, J. H. (1987), Measuring Business Cycle Time, *Journal of Political Economy* **95**, 1240-1261.
- [87] Stock, J. H. (1988), Estimating Continuous-Time Processes Subject to Time Deformation, *Journal of the American Statistical Association* **83**, 77-85.
- [88] Taqqu, M. S. (1975), Weak Convergence to Fractional Brownian Motion and to the Rosenblatt Process, *Z. Wahrscheinlichkeitstheorie verw. Gebiete* **31**, 287-302.
- [89] Taylor, S. (1986), *Modelling Financial Time Series*, New York: Wiley.
- [90] Vandewalle, N., and Ausloos, M. (1998), Multi-Affine Analysis of Typical Currency Exchange Rates, *European Physical Journal B* **4**, 257-261.
- [91] Wiggins, J. B. (1987), Option Values under Stochastic Volatility: Theory and Empirical Estimates, *Journal of Financial Economics* **19**, 351-372.