

Multifractality of Deutschemark / US Dollar Exchange Rates*

Adlai Fisher and Laurent Calvet[†]
Department of Economics, Yale University

Benoit Mandelbrot[‡]
*Department of Mathematics, Yale University and
IBM T. J. Watson Research Center*

Cowles Foundation Discussion Paper No. 1165

This Draft: September 15, 1997
First Draft: October 1996

THIS VERSION TEXT ONLY
Download figures at <http://www.econ.yale.edu/~fisher/papers.html>

*We are grateful to Olsen and Associates, Zürich, Switzerland, for kindly providing the foreign exchange data.

[†]28 Hillhouse Avenue, New Haven, CT 06520-1972. e-mail: lcalvet@minerva.cis.yale.edu, fisher@econ.yale.edu

[‡]10 Hillhouse Avenue, New Haven, CT 06520-8283. e-mail: fractal@watson.ibm.com

Abstract

This paper presents the first empirical investigation of the Multifractal Model of Asset Returns (“MMAR”). The MMAR, developed in Mandelbrot, Fisher, and Calvet (1997), is an alternative to ARCH-type representations for modelling temporal heterogeneity in financial returns. Typically, researchers introduce temporal heterogeneity through time-varying conditional second moments in a discrete time framework, or time-varying volatility in a continuous time framework. Multifractality introduces a new source of heterogeneity through time-varying local regularity in the price path. The concept of local Hölder exponent describes local regularity. Multifractal processes bridge the gap between locally Gaussian (Itô) diffusions and jump-diffusions by allowing a multiplicity of Hölder exponents. This paper investigates multifractality in Deutschemark / US Dollar currency exchange rates. After finding evidence of multifractal scaling, we show how to estimate the multifractal spectrum via the Legendre transform. The scaling laws found in the data are replicated in simulations. Further simulation experiments test whether alternative representations, such as FIGARCH, are likely to replicate the multifractal signature of the Deutschemark / US Dollar data. On the basis of this evidence, the MMAR hypothesis appears more likely. Overall, the MMAR is quite successful in uncovering a previously unseen empirical regularity. Additionally, the model generates realistic sample paths, and opens the door to new theoretical and applied approaches to asset pricing and risk valuation. We conclude by advocating further empirical study of multifractality in financial data, along with more intensive study of estimation techniques and inference procedures.

Keywords: Multifractal Model of Asset Returns, Multifractal Process, Compound Stochastic Process, Trading Time, Time Deformation, Scaling Laws, Multiscaling, Self-Similarity, Self-Affinity

1 Introduction

Multifractality is a promising new tool for economists. Recent theoretical research proposes the Multifractal Model of Asset Returns (“MMAR”) as a method of modelling financial prices. (Mandelbrot, Fisher and Calvet, 1997; Calvet, Fisher and Mandelbrot, 1997; Mandelbrot, 1997). The MMAR is thus an alternative to ARCH and its variants. Multifractal processes incorporate long-tailed asset returns and long memory in “volatility”. Additionally, the MMAR predicts a form of scaling in the moments of returns that is new to both theoretical and empirical finance.

Current finance literature notes a discrepancy between financial theory, which is largely set in continuous time, and empirical research that tends towards discrete formulations. (Rossi, 1997; Drost and Werker, 1996; Melino, 1994) This discrepancy is sometimes viewed as a “gap” to be remedied by improvements in empirical models. Multifractality offers a different perspective, showing that a large class of interesting stochastic processes are not covered by the Itô diffusion paradigm.¹ Importantly, the MMAR contains a restriction, scaling, that is strong enough to generate empirical tests. “Scaling” describes a specific relationship between data samples of different time scales, e.g. daily, weekly, or monthly returns.² We envision a growing set of econometric methods that explicitly test or model the behavior of data at different time scales. Established research along these lines includes the literature on decreasing sequences of ARCH representations (Nelson, 1990; Nelson and Foster, 1994; Drost and Werker, 1996). This work assumes a decreasing sequence of sampling frequencies, and therefore matches the natural continuous time framework of the MMAR.

The present paper conducts the first empirical investigation of the MMAR. The empirical work uses an estimation technique that is new to economics, and is validated through simulation. Simulations also are used to illuminate important differences in a core group of extant financial models. Overall, the empirical results are quite promising for the MMAR, showing strong regularities in an important data set that are consistent with restrictions imposed by multifractality. We hope

¹Itô diffusions are locally quadratic, allowing arbitrage arguments and other mathematical conveniences that would require distributional assumptions in discrete time. Theoretical results that depend critically upon Itô diffusions are generally understood to be idealized, and problematic because of the presence of discontinuities. The MMAR reinforces reservations about the “continuous hedge” by postulating a more generally diverse variety of local behavior. Even when price paths are continuous, they in general are not locally quadratic.

²A similar term that has been used in these papers is “scale-consistent”, which has the same meaning as “closed under temporal aggregation”. By this, we mean that the class includes a representation of the process at each discrete time scale. Thus, GARCH is not scaling, but under certain assumptions, is scale-consistent. (See Drost and Nijman, 1993).

to encourage further research by demonstrating the empirical content of the model, and to lay groundwork that suggests which approaches will be most fruitful.

The structure of the paper is as follows: Section 2 presents the primary estimation methodology. Multifractal processes are reviewed and their scaling properties are related to restrictions upon moments of the data. This leads to an informal test of multifractality, and an estimation method for the *scaling function*, which describes growth rates of moments as the time scale increases. Section 3 applies these methods to Deutschemark / US Dollar exchange rate data. We find strong evidence of scaling, consistent with the predictions of the MMAR.

Section 4 reviews the multifractal spectrum, which is a renormalized density of points in time with different degrees of local regularity. We show that the multifractal spectrum can be estimated by Legendre transform of the scaling function. Section 5 carries out the Legendre transform on the estimated scaling function and discusses these results. The MMAR components are recovered from the estimated multifractal spectrum. Finally, Monte Carlo simulations show a striking resemblance to the DM/USD data, although consistent replication of scaling requires large simulated sample sizes. This result is discussed in Section 5.

Section 6 simulates GARCH and FIGARCH processes, using parameters obtained from previously published research on exchange rates. Lack of long memory is made intuitively clear by graphical representations of GARCH over long time intervals. Similar FIGARCH simulations show long memory, but lack variability at high frequencies when compared to the data (or the MMAR). Neither GARCH nor FIGARCH captures the scaling properties of the data with consistency, even at large sample sizes. We discuss how to interpret these results.

Section 7 conducts a variety of robustness tests. We apply several seasonal adjustment filters to the high frequency data, split the data to test for model stationarity, and perform similar tests with the Japanese Yen/ US Dollar series. Scaling properties are robust to reasonable transformations of the DM/USD data. Scaling behavior in the Yen/ Dollar series is more complicated, and is discussed at the end of Section 7.

Section 8 concludes. Overall, we find that the MMAR captures a multitude of important features of the data. These include long memory in volatility, long tails, scaling, and high variability in returns at high frequencies. This line of research should therefore be extended in several directions, including theory, empirical methods, and simulation methods. There is also a need for more

comprehensive empirical work using different data sets. Finally, the most important applications of this model will be in the valuation of risk.

2 Testing the MMAR: Methodology

Mandelbrot, Fisher, and Calvet (“MFC”, 1997) introduces multifractality to economics, focusing upon scaling properties in moments. The main result is the Multifractal Model of Asset Returns (“MMAR”), which shows that multifractal measures, when combined with the concept of trading time, generate continuous-time stochastic processes that have long memory in volatility, long tails, and are flexible enough to generate either martingale or long memory behavior in log prices.

Calvet, Fisher, and Mandelbrot (“CFM”, 1997) further develops the theory of multifractals, focusing upon local scaling properties. This paper develops the concept of Hölder exponent, which describes the local regularity of the price path, and the multifractal spectrum, which is a renormalized density of local Holder exponents. These concepts are then related back to the MMAR.

This section and Section 4 review multifractality, emphasizing concepts necessary for empirical work. Readers interested in further development of the theory should consult MFC, CFM, and the literature cited therein.

2.1 Review of the MMAR and Multifractal Processes

Let $X(t, \Delta t)$ generically denote the increments of a stochastic process $X(t)$, i.e.

$$X(t, \Delta t) \equiv X(t + \Delta t) - X(t), \quad 0 \leq t \leq T.$$

Also, let $X(\Delta t)$ stand for $X(0, \Delta t)$, which is distributionally invariant under translations when stationarity in increments is assumed. MFC defines a multifractal process as a continuous time process with stationary increments that satisfy

$$\mathbb{E}[|X(t, \Delta t)|^q] = c(q) (\Delta t)^{\tau(q)+1} \tag{1}$$

for all $t, \Delta t$ on which X is defined, and for all q such that $\mathbb{E}[|X(t, \Delta t)|^q] < \infty$. This expectation is assumed finite over an interval $[0, q_{\max})$, which may be open or closed on the right.

Equation (1) is a scaling law in the moments of $X(\Delta t)$. It specifies a form for their change as Δt varies. All information about the rate of growth is contained in the function $\tau(q)$, called the

scaling function. Thus, $\tau(q)$ is a focal concept in the theory of multifractals. The prefactor, $c(q)$, is not considered in this paper, except for the fact that it does not vary over t or Δt .

The MMAR posits a Fractional Brownian Motion, $B_H(t)$, compounded by a multifractal trading time, $\theta(t)$. Trading time and the Fractional Brownian Motion are assumed independent. The resulting multifractal price process is written³

$$\log P(t) - \log P(0) = B_H[\theta(t)]. \quad (2)$$

The strategy followed in this paper will be first to test the scaling law in equation (1). We then estimate $\tau(q)$, which is transformed into an estimate the multifractal spectrum. This, in turn, allows us to recover the individual components of the MMAR and simulate the price process (2). A review of the multifractal spectrum, which is not necessary to test the existence of a multifractal scaling law, is deferred until Section 4.1.

2.2 A Simple Test of Multifractality

Consider a price series $P(t)$ defined on all t belonging to the fixed time interval $[0, T]$. Denote the logarithm of price by $X(t) \equiv \log P(t) - \log P(0)$. Dividing $[0, T]$ into N intervals of length Δt , define the *sample sum* (also called *partition function*⁴ from origins in thermodynamics),

$$S_q(T, \Delta t) \equiv \sum_{i=0}^{N-1} |X(i\Delta t, \Delta t)|^q. \quad (3)$$

By the stationary increments property of multifractal processes, the addends are identically distributed, although they may have arbitrary correlation. If $X(t)$ is multifractal, the scaling law yields

$$\log \mathbb{E}[S_q(T, \Delta t)] = \tau(q) \log(\Delta t) + c(q) \log T \quad (4)$$

when the q^{th} moment exists.

The methodology employed in the remainder of the paper is straightforward. We plot $\log S_q(\Delta t)$ against $\log \Delta t$ for various values of q and various values of Δt . The linearity of these plots for given

³All of the empirical work in this paper uses base 10 logarithms, except for the initial transformation of prices, which is a natural logarithm. The companion theoretical papers work in natural logarithms exclusively, and use the \ln notation. This paper uses a single notation, \log , which should be interpreted according to whether the use is empirical or theoretical. All of the theory converts easily (but tediously) to base 10 (or vice versa).

⁴The logarithm of S_q is also frequently referred to as the partition function. We hope the reader will tolerate this ambiguity in order to expedite discussion.

values of q is proposed as a test of the MMAR. We check linearity by visual inspection. The slope of the lines, estimated by OLS regression, gives an estimate of the scaling function $\tau(q)$. The estimated scaling function is easily transformed into an estimated multifractal spectrum.

Despite the simplicity of these methods, theoretical justification beyond equation (4) is quite difficult. Mandelbrot (1995, Section 6.4) has conjectured consistency of estimated slopes for a large class of multifractals as $\Delta t \rightarrow 0$. In general, a theoretical proof of this conjecture has been elusive, except for a few simple cases such as the binomial and multinomial multiplicative cascades. Consistency, along with inference methods, remains an open and interesting area for future theoretical work.

Until additional theory is developed, we justify the above methodology on two counts. First, while inversion of logarithms and expectations will produce small sample biases in estimation, these need not prevent transmission of linearity in (4) to approximate linearity in the sample partition function.⁵ Secondly, estimation biases for $\tau(q)$ can be studied and understood for particular cases of interest.

Monte Carlo simulation validates the empirical methods used in this paper. After Section 3 shows linearity in DM/USD partition functions, Section 5 estimates $\tau(q)$ and uses this information to generate an estimated multifractal spectrum. Using results provided in MFC and CFM, we then recover a generating mechanism with theoretical multifractal spectrum similar to the data. Simulated partition functions of this process follow their predicted paths, with (approximately) linear partition functions of (approximately) the same slope as the DM/USD data.

Finally, the partition function methodology has some power to discriminate against reasonable alternatives. Section 5 simulates GARCH and FIGARCH processes with parameters from published research on exchange rates. Although the simulated alternatives mimic scaling behavior in individual samples, the MMAR replicates the scaling behavior found in the data more consistently. All of the simulation results should be considered preliminary investigations, and should be studied more thoroughly in future research.

⁵By Jensen's inequality, rewriting equation (4) as a regression yields errors with negative expectation. The effect on linearity of the partition function and estimated slope depends on both the magnitude of the bias and how the bias varies with Δt . Both of these problems are tractable with simulation methods, although beyond the scope of the present paper.

2.3 Unifractal Examples: Brownian Motion and Fractional Brownian Motion

Consider a (Wiener) Brownian Motion (“WBM”, to distinguish from Fractional Brownian Motion, “FBM”) with volatility σ . By self-affinity, $E[|X(t, \Delta t)|^q] = \sigma (\Delta t)^{q/2}$ for all t , and all $q, \Delta t > 0$. Independent increments and the strong law of large numbers immediately yield

$$\tau(q) = \frac{q}{2} - 1.$$

We simulate the increments at $T = 10^5$ uniform intervals, and calculate $S_q(\Delta t)$ for several values of q between 1.5 and 5 and numerous values of Δt between 1 and 1000. Figure 2 shows $\log \Delta t$ plotted against $\log S_q(\Delta t)$ for each value of q . The plotted values fall very nearly on lines of the predicted slope. The predicted values are shown as dotted lines in the figure.

We now repeat the example above for several Fractional Gaussian Noises (“FGN”), the discrete time analogues of Fractional Brownian Motion. Understanding these processes is important, since FBM is one of two main components in the MMAR. Mandelbrot and Wallis (1969) provides two discrete approximations to the FBM. A third approximation, which reduces computational complexity, appears in Mandelbrot (1971). These discrete FGN received widespread application in the natural sciences and economics.

Granger (1980), Granger and Joyeux (1981), and Hosking (1980) provide a fourth discrete approximation to the FBM,⁶ proposing the fractional differencing operator in discrete time as the binomial expansion of $(1 - L)^d$. This gives the familiar ARFIMA(0, d , 0) process,

$$(1 - L)^d x_t = u_t,$$

where $\{u_t\}$ is white noise and $d = H - 1/2$. The Wold decomposition is

$$\begin{aligned} x_t &= \sum_{k=0}^{\infty} \psi_k u_{t-k} \\ \psi_k &= \frac{\Gamma(k+d)}{\Gamma(k+1)\Gamma(d)} \sim \frac{1}{\Gamma(d)} k^{d-1} \quad \text{as } k \rightarrow \infty, \end{aligned} \tag{5}$$

where Γ denotes the Gamma function. Asymptotically, the decay rate of $d-1$ in the MA coefficients is identical to the previous discrete approximations of 1968 and 1971. This gives an asymptotic

⁶This approximation is called Fractional White Noise by some authors. The perceived conceptual distinction is that the FGN approximates increments of the FBM, while the FWN applies the fractional differencing operator to white noise.

decay rate in autocorrelations of $2d-1$, which (alternatively expressed in terms of ACF or spectrum) is the fundamental characteristic of FBM and discrete approximations to the FBM. Short-memory components can be added to fit high-frequency features of the data as suggested by Mandelbrot (1971), and formalized in the ARFIMA(p, d, q) process.

Figure 3 shows several simulations of Fractional Gaussian Noise, H ranging between .2 and .8. In addition, we show a near-Fractional Gaussian Noise that is just non-stationary, and several processes whose increments are FGN with $0 < H < 1$. Beginning with $d = .5$, the span of the vertical axis begins to increase dramatically as a consequence of non-stationarity.

These graphs are notably different from the data. The persistent values of $.5 < H < 1$ have long-memory in their increments, which is usually diagnosed by a slow rate of decay in the sample autocorrelation function. Long-memory can be detected visually when $H = .8$, and somewhat when $H = .6$, by apparent long-period trends in direction. More distinctive features of the DM/USD data are low frequency shifts in the size, or absolute value, of price increments. When simulating multifractal processes in Section 5, the distinction between long-memory in increments and long-memory in the absolute value of increments will again become apparent.

Figure 4 shows, in solid lines, the estimated partition functions for four of the simulated FGN's. The dotted lines are reference lines showing the expected values for Wiener Brownian Motion, and the dashed lines show theoretical expectations for FBM's of the same H . All simulated partition functions follow the prediction for FBM fairly closely.⁷ Even for $H = .53$, which is indistinguishable from Wiener Brownian Motion by inspection of sample paths or increments, differences in slope are apparent when using the partition function methodology.

3 Testing the MMAR: Empirical Results

3.1 The Data: Deutschemark / US Dollar Exchange Rates

The Deutschemark / US Dollar (“DM/USD”) currency market is the most thickly traded exchange rate market in the world (three times the size of the Yen/Dollar market on the basis of quotes), and high frequency (quote by quote) data is readily available. Olsen and Associates, a currency research and trading firm located in Zürich, Switzerland, provided data for a one year period (October 31,

⁷As the distance between H and $1/2$ increases, some biases appear. A likely explanation is approximation error in the simulations. We use the infinite MA representation in (5) with a cutoff at 30,000 lags. Alternative approximations to the FBM will be explored in future work.

1992 – September 1, 1993). The data set consists of all bid/ask quotes, with transmittal times, collected by Olsen & Associates during that year. Quotes are collected from banks and other trading institutions world-wide, creating a data set with round-the-clock observations. For the DM/USD series, this amounts to 1,472,241 quotes. Previous researchers using the Olsen data have noted that traders have strong incentives not to transmit spurious quotes, and Olsen provides a filter for quotes which they believe to be erroneous.⁸

A second data source is a nearly twenty-four year series of daily DM/USD exchange rates. This data set also was provided by Olsen and Associates, and spans from June 1973, shortly after the permanent demise of the Bretton-Woods system of exchange rate pegs, to December 1996. The reported price is the geometric mean of bid and ask quotes.⁹ The prices were recorded at 16:00 UK time as a linear interpolation of the two closest (in time) quotes. Data are provided for every weekday except New Years and Christmas. Figure 1 shows the DM/USD daily series, along with increments and log increments, over the length of this sample period.

Combining these two data sets allows calculation of the partition function over three orders of magnitude of Δt . This permits investigation of scaling behavior over a broad range of sampling frequencies.

3.2 Seasonality

Figure 5 shows weekly patterns in two measures of market activity. These are, respectively, average number of quotes per fifteen minutes and average absolute returns per fifteen minutes. Trading activity falls significantly during weekends, from the close of trading in the U.S. on Friday to the opening of the Australian and Asian markets on Monday. Weekday market activity subsides from approximately 3:00-4:30 AM GMT, when it is lunchtime in the major Asian economies and other major markets are closed.

Seasonal components are non-scaling, and preclude that trading time is multifractal as defined in MFC.¹⁰ While this claim may not be intuitively obvious, it depends upon viewing multifractal

⁸Olsen & Associates provide a flag for quotes they believe to be either erroneous or not representative of actual willingness to trade. We eliminate the flagged observations, which comprise .36% of the dataset.

⁹The logarithm of the geometric mean is the logarithmic mid-point. This is a standard approximation to price when only bid and ask quotes are available. We test sensitivity to this approximation in Section 7, which uses buying rates provided by the Federal Reserve.

¹⁰This set of papers has presented a fairly strict definition of multifractality, based upon the exact scaling equation (1). Alternatively, we have discussed statistical self-similarity of random measures at times, as in Appendix 9.1.

trading time as the cumulative distribution function (“c.d.f.”) of a statistically self-similar random measure. For interested readers, this argument is presented in Appendix 9.1. Otherwise, seasonal adjustment can be viewed simply as a stationarity requirement.

This complication is addressed by applying a seasonal adjustment filter¹¹ that reduces the seasonal component in the data. The seasonally modified version of the MMAR is written

$$\ln P(t) = B_H \{ \theta [SEAS_n(t)] \},$$

where $SEAS_n$ denotes the n^{th} seasonal adjustment filter among the several that we employ.¹² The basic idea is that time is expanded during typically active periods of the week, and contracted during typically slow periods. The expansions and contractions are restricted to maintain a constant amount of total time, i.e. 604,800 seconds in a week.

Even for established time series models, seasonal adjustment is a delicate issue. (See the survey by Ghysels, 1994.) Optimal procedures, in the sense of Grether and Nerlove (1970), are thus beyond the scope of the present paper, and should be investigated separately. Instead, we choose a range of imperfect, but reasonable, seasonal adjustment filters. Some have the merit of predictability in whether they under-adjust or over-adjust for the seasonal component. This range allows us to test robustness of results to reasonable variation in the seasonal adjustment method.

To simplify presentation, the results in Sections 3 and 5 use a single seasonal adjustment filter. Section 6 separately investigates the issue of robustness, and presents results for the other four filters. The filter used for the main results is $SEAS_2$ which deforms time to smooth variation in average absolute returns over fifteen minute intervals of the week.

A broader view of multifractality is simply the presence of multiple Hölder exponents, as in CFM. While we have not fully developed this viewpoint, one would proceed by presenting multifractality as an asymptotic, and thus less restrictive, requirement. This definition would permit seasonality, since seasonality can typically be viewed as a smooth transformation that affects prefactors, but not Hölder exponents. As long as prefactors are well behaved, the scaling properties as $\Delta t \rightarrow 0$ are not affected. This idea may be further developed in the future, but does not seem to present significant advantages in empirical work at the present time.

¹¹The term “filter” as used here simply means the data are “treated in a particular way before they are analyzed” (Hamilton, 1990, p.63).

¹²Hereafter, we specify the seasonal filter in text and graphs, and drop this notation from future equations. Construction of the seasonal adjustment filters is discussed in Appendix 9.2.

3.3 Plots of the Partition Function

Figure 6 shows renormalized graphs of $\log S_q(\Delta t)$ against $\log(\Delta t)$ for values of q around 2.¹³ From bottom to top on the graph, the values of q increase from 1.75 to 2.25. The values of Δt range from 20 seconds to 180 days, and are spaced at intervals of approximately .0414 on a \log_{10} scale, or multiples of about 1.1 on a real scale. The renormalization for all q is $\log S_q(20 \text{ seconds}) = 0$. This renormalization is useful because the intercepts of the curves are not studied, and it allows presentation of several curves on the same graph. Sample sums for the high frequency data and the daily data are spliced together by the vertical displacement that provides the best linear fit, using OLS, under the restriction that both lines have the same slope.¹⁴

Figure 7 plots higher moments (up to $q = 5$), which are generally more sensitive to deviations from scaling. Format is the same as the Wiener Brownian Motion example from Figure 2. Dotted lines are provided for reference, and show expectations for data generated by WBM. The horizontal axis begins at $\log(2 \text{ hours})$ to focus attention on the region of principal interest. Each partition function stops when the number of independent observations in the sample falls approximately below 50. Thus, for the 365 day sample of high frequency data, the highest value of Δt represented is 7 days, and 180 days is the approximate maximum for the daily data. The following section explains what can be learned from inspection of these graphs.

3.4 Main Results

3.4.1 High Frequency Crossover

The most striking aspects of Figure 6 are the two clean breaks in the curves, highlighted by vertical dotted lines. The first dotted line corresponds to five minutes of deseasonalized clock time. The second dotted line corresponds to about two hours of deseasonalized clock time.¹⁵

The existence of lower and upper scale bounds for a model has been termed *crossover* in the physics literature. This seems an appropriate term to extend to the phenomenon evident in Figure

¹³For the high frequency data, the price at time t is approximated by the geometric mean of bid and ask for the most recently observed quote. This is consistent with the format the daily data was received in, explained in Section 3.1. Sensitivity to this approximation is addressed in Section 7.

¹⁴For canonical multifractal measures, vertical displacement is generally a random variable that depends only on q . We expect this trait to be passed on to the price process by $\theta(t)$, and thus place no restrictions on the intercepts.

¹⁵During weekdays when a major market is open, these amounts would decrease by about 1/2 to 2/3 in units of clock time. (Recall that seasonal adjustment expands active times of the week and contracts slow times.) Exact conversions from deseasonalized clock time to clock time at different times of the week can be made using the vertical axis in Figure 5a.

6 at about two hours. For the remainder of this paper, we consider primarily the long range of sampling intervals greater than two hours. These sampling intervals still span nearly three orders of magnitude of Δt . The absence of a sharp *upper* crossover in this region, at least up to 180 days, is promising.

We conjecture that high-frequency crossover can be related to market frictions such as bid-ask spreads, discreteness of trading units, and discontinuous trading. This topic will be investigated in a forthcoming paper.

3.4.2 The Self-Affinity Index for $B_H(t)$

A close up view of the slopes of the partition functions around the moments $q = 2$ is chosen in Figure 6 for a particular purpose. The self-affinity index, H , of the Fractional Brownian Motion $B_H(t)$ in the MMAR is identified by the condition that

$$\tau_P(1/H) = 0, \tag{6}$$

where τ_P is the scaling function for the price process.¹⁶

From Figure 6, the value of q with slope zero appears to be about 1.9, indicating $H \approx .53$. This implies small persistence in DM/USD exchange rates. At the same time, the estimated H is very close to $1/2$. Previous studies, based on well understood estimation methods, find white spectra (implying $H = 1/2$) for log increments of typical financial data. Thus, the martingale hypothesis for log prices seems to be a reasonable approximation. At the same time, Section 5 shows strong evidence of multifractality in trading time, implying long memory in squared returns.

3.4.3 Linearity of the Partition Functions

Linearity of the partition functions in Figure 6 indicates a multifractal scaling law that holds, at a minimum, for close to three orders of magnitude of Δt . Scaling is even more striking in Figure 7 because the moments are higher, reaching to $q = 5$. Once again, the lines are relatively straight, consistent with the MMAR. Variability of the lines around their apparent slopes increases as q increases, as is to be expected.

¹⁶In a partition with fixed endpoints where all increments are restricted to have the same sign, the sum of increments cannot depend on grid size. Intuitively, the total variation of an increasing process, estimated by finite increments, does not depend on how we cut it up. Therefore, since $\theta(t)$ is increasing (being the c.d.f. of a random measure), $\tau_\theta(1) = 0$. For a Fractional Brownian Motion, $\tau(q) = Hq - 1$, and $\tau_P(q) = \tau_\theta(Hq)$ by Theorem 4.2 in MFC(1997). This establishes that $\tau_P(1/H) = 0$. See CFM and Mandelbrot (1997) for further discussion.

Closer inspection of Figures 6 and 7 reveals that the lines are in fact composed of two distinct segments, corresponding to the high frequency data and the daily data. At $\Delta t = 1$ day, the daily partition function overlaps the high frequency partition function with nearly the same slope, and greater regularity because of a sample span nearly 24 times as large. Once again, as Δt increases, the variability around the apparent slope increases, until estimation ceases at 180 days.

The strong similarity in slopes between the high frequency data and the daily data suggests model stationarity. At least for moments up to $q = 5$, the scaling law for 1992-1993 is similar to the scaling law for 1973-1996. Returning to the daily price series of DM/USD returns presented in Figure 1, note several long swings in price as the Dollar appreciates and depreciates against the Mark. Working with the variety of GARCH, FIGARCH, ARFIMA, stochastic trend, and stochastic volatility models that are currently available, few researchers would attempt to capture long term behavior in the DM/USD series over the twenty-five year sample span with the same model that explains relatively high frequency behavior in 1992-1993.

This highlights a major innovation in the MMAR approach. By modelling restrictions upon the relationship between distributions or moments at different time scales, we can generate new predictions and new econometric methods for testing these restrictions.

This section has shown evidence of a multifractal scaling law for DM/USD exchange rates. In Figures 6 and 7 the eye is unable to discriminate differences in $\hat{\tau}(q)$ (where $\hat{\tau}$ indicates the estimated scaling function) between the high frequency data set and the daily data set. Section 4 estimates $\tau(q)$ separately for each data set, and employs Legendre transforms of $\hat{\tau}(q)$ to obtain estimates of the multifractal spectrum.

4 Review of the Multifractal Spectrum

The scaling law in moments (1) is useful in generating testable restrictions, but not an intuitive description of multifractality. Other intrinsic characteristics of multifractals might serve as definitions, though with greater technical complications. For example, Mandelbrot (1972, 1974) begins with the distributional scaling of statistically self-similar random measures and arrives at a scaling law in moments as an implication.

The deepest, though probably most difficult, understanding of multifractals comes from studying local scaling properties. This is explored in CFM, and reviewed here in order to prepare for the

empirical work in Section 5. The concept of Hölder exponent is essential. The local Hölder exponent, $\alpha(t)$, measures the local regularity of a process:

$$X(t, \Delta t) \sim (\Delta t)^{\alpha(t)} \quad \text{for all } t, \text{ and as } \Delta t \rightarrow 0. \quad (7)$$

In a standard Itô diffusion, the local variation is always proportional to $(\Delta t)^{1/2}$. The most common way to break this local scaling law is to introduce a jump component, which is akin to a local scaling rule of the form $(\Delta t)^0$.

Multifractal processes generate variety in local regularity while filling the gaps between locally $(\Delta t)^{1/2}$ and $(\Delta t)^0$ behavior. Multifractal processes contain a multiplicity of local Hölder exponents. One might expect this diversity to result in unstructured behaviors. On the contrary, knitting together sets of points with different Hölder exponents requires a particular kind of structure, which is the topic of the two theoretical papers MFC and CFM.

The multifractal spectrum is the main tool for describing the distribution of Hölder exponents. The following section reviews three subtly different interpretations. Section 4.2 provides a heuristic for estimating the multifractal spectrum via a Legendre transform of the scaling function. Finally, Section 4.3 reviews some basic facts about the multifractal spectrum that are useful in empirical work.

4.1 Three Interpretations of the Multifractal Spectrum

The multifractal spectrum is a function $f(\alpha) \leq 1$, that may be defined on any subset of the real line. The domain can be thought of as a space of Hölder exponents when $\alpha \geq 0$. (This is the only case considered in the present paper.)¹⁷ The goal of the multifractal spectrum is to describe the distribution of local Hölder exponents in a multifractal process. The most natural place to begin this investigation is with the standard frequency renormalization of probability theory. Partitioning time into a grid with increments of size $\Delta t = b^{-k}$, we define the *coarse Hölder exponents*,

$$\alpha_k(t) = \frac{\ln X[t, \Delta t]}{\ln \Delta t}.$$

By increasing k we proceed to the limit $\Delta t \rightarrow 0$, and may expect that the frequency histogram of coarse Hölder exponents will converge to the frequency distribution of Local Hölder exponents. This

¹⁷The case $\alpha < 0$ cannot correspond to a Hölder exponent for measures or bounded functions, but has gained an alternative interpretation as a *virtual* α . This topic is at the frontiers of research in mathematics, and is beyond the scope of the present paper.

intuition turns out to be correct, but misses a great deal of information. Mandelbrot (1989) shows that an alternative renormalization is more informative. CFM neatly categorizes this interpretation with two other common interpretations for $f(\alpha)$:

(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 $k^{-1} \log_b \mathbb{P} \{ \alpha_k > \alpha \} + 1$ provided by Large Deviation Theory.

The second interpretation is a statement about *fractal dimension*, or Hausdorff-Besicovitch dimension. In a partition of size Δt , the number of intervals characterized by α takes the form

$$N_\alpha(\Delta t) \sim (\Delta t)^{-f(\alpha)}.$$

This form justifies calling $f(\alpha)$ a fractal dimension.¹⁸ On the real line, fractal dimensions are constrained to lie between zero and one. Fractal dimension and (unit renormalized) Lebesgue measure coincide when both are either zero or one. On the other hand, all sets with positive Lebesgue measure have a fractal dimension of one, and all sets with fractal dimension less than one have zero Lebesgue measure. Thus, fractal dimension provides more information about some types of sets. Standard continuous-time financial models have components with fractal dimension of one, which applies to typical Itô diffusions, or zero, which applies to jump components. Thus, fractal sets fill space in a manner that is midway between the standard cases of continuous time finance.

Definition (D3) is a statement about applying Large Deviation Theory (“LDT”) to multifractals generated via multiplicative cascades. In the case of multifractals, the LDT provides far more information about the generating mechanism than do the Strong Law of Large Numbers or Central Limit Theorem.

These three definitions coincide when $f(\alpha) > 0$, but clearly (D1) precludes negative values of $f(\alpha)$.¹⁹ Negative fractal dimensions are known as *latent* α 's, deriving from the fact that sets of

¹⁸Note that the quantity $-\log N_\alpha(\Delta t)/\log(\Delta t)$, which is $f(\alpha)$, does not depend on Δt .

¹⁹Definition (D2) is compatible with negative values for $f(\alpha)$ since Mandelbrot (1988, 1995) introduced the concept of negative dimension as “degree of emptiness of an empty set.” The idea is troubling at first, but has an interesting probabilistic justification related to the embedding of a process or measure in a higher-dimensional space. For example, in a one dimensional section drawn from a three dimensional space populated by stellar matter, the fractal dimension of stellar matter is negative. An individual draw almost always has no matter, but by the generation of numerous low-dimensional sections, called *supersampling*, one obtains information about the (positive) dimension of stellar matter in the embedding space.

negative fractal dimension will typically not be found in a random sample. Estimating latent α 's requires *supersampling*, or obtaining many samples from the same higher dimensional space.²⁰ This paper focuses on *manifest* α 's, for which $f(\alpha) \geq 0$. The positive part of the multifractal spectrum can be obtained from a single long sample, and describe the typical behavior of the process.

4.2 A Hölder Exponent Heuristic for the Estimation of $f(\alpha)$

This section provides a brief heuristic showing the relationship between the scaling law in moments, which is the focus of our empirical work, and the multiple Hölder exponents property of multifractal processes. We see through this heuristic that the multifractal scaling function $\tau(q)$ and asymptotic linearity of the partition function come from the fractal nature of sets of points in time with a given Hölder exponent (*D2*). In addition, the heuristic shows how to obtain the estimated multifractal spectrum of the process by applying a Legendre transform to the estimated scaling function $\hat{\tau}(q)$.²¹

Consider again a finite sample span $[0, T]$. Define

$$\hat{\tau}(q) = \log S_q(T, \Delta t) / \log \Delta t = \log \left[\sum_{i=0}^{N-1} |\log X(i\Delta t, \Delta t)|^q \right] / \log \Delta t$$

For a given realization of a sample path and small Δt ,

$$|X(t, \Delta t)|^q \sim (\Delta t)^{q\alpha(t)},$$

where $\alpha(t)$ is the local Hölder exponent of the path at t . At the scale Δt , the distribution of α is of the form $c(\alpha)(\Delta t)^{-f(\alpha)}$, where $c(\alpha)$ is a constant. Thus, $f(\alpha)$ is the fractal dimension of the set of points having local Hölder exponent α . Rewriting the sum over price increments as a sum over Hölder exponents yields

$$\sum_{i=0}^{N-1} |X(i\Delta t, \Delta t)|^q \sim \int c(\alpha) (\Delta t)^{q\alpha - f(\alpha)} d\alpha.$$

As Δt goes to zero, the main contribution to the integral comes from:

$$(\Delta t)^{Min[q\alpha - f(\alpha)]},$$

²⁰Supersampling circumvents non-ergodicity in long memory processes. This is conceptually possible by obtaining repeated long samples of the same generating mechanism, but probably has limited applicability economics. See the previous note.

²¹This heuristic is known as the Method of Steepest Descents in the thermodynamics literature, and dates at least back to Cramér (1936). It was later used by Halsey, et al. (1986), and Frisch and Parisi (1985).

and thus

$$\hat{\tau}(q) \rightarrow \underset{\alpha}{\text{Min}} [q\alpha - f(\alpha)] \quad \text{as } \Delta t \rightarrow 0.$$

We see that $\hat{\tau}(q)$ converges to the Legendre transform of f . Conversely, $f(\alpha)$ can be estimated by Legendre transform of $\hat{\tau}(q)$:

$$\hat{f}(\alpha) = \underset{\alpha}{\text{Min}} [q\alpha - \hat{\tau}(q)]. \quad (8)$$

This provides an estimation method for $f(\alpha)$.

4.3 Reference Points on the Multifractal Spectrum

The empirical work that follows estimates scaling functions and multifractal spectra for multifractal processes. Two essential results allow us to decompose the multifractal process into the two components postulated by the MMAR.

First, by Theorem 4.2 of MFC,

$$\tau_P(q) = \tau_\theta(Hq), \quad (9)$$

where P denotes the price process, θ denotes trading time, and H denotes the self-affinity index of the FBM. Since H is estimated by (6), both components are identified. The scaling functions of prices and trading time are each deducible from the other by (9). Similarly, Theorem 5.1 of CFM allows us to relate the multifractal spectra of the price process and trading time by the relation

$$f_P(\alpha) = f_\theta(\alpha/H). \quad (10)$$

Note that these two theorems imply one another by application of Legendre transforms.

An added benefit of these theorems is that simple results which have been derived for multifractal measures can be immediately applied to multifractal processes. For example, concavity of multiplicatively generated multifractal measures implies concavity of $f_P(\alpha)$. Also, the spectrum of a multifractal measure always lies below the line $f(\alpha) = \alpha$, and has a tangent to this line at exactly one point.²² Thus, $f_P(\alpha) \leq \alpha/H$, with equality at a single point. Additional notation and simple facts about the multifractal spectrum follow:

²²Violation of this would imply that total measure increases as we take finer and finer partitions of an interval. See Evertsz and Mandelbrot (1992) for further intuition.

- α_0 denotes the value of α for which $f(\alpha_0) = 1$. This is the maximum value of the multifractal spectrum. By strict concavity, α_0 is unique. This is also the most frequently occurring value of α , occupying a set of Lebesgue measure one (or T if the sample span is not renormalized to a value of one.)
- α_1 denotes the value of α that carries most of the mass of a multifractal measure, or most of the variation of a multifractal process. Multifractal measures satisfy $f(\alpha_1) = \alpha_1$. The spectrum of a multifractal process generated by the MMAR thus satisfies $f_P(\alpha_1) = 1/H$.
- α_{\min} denotes the smallest manifest α , i.e. $\alpha_{\min} = \left\{ \min_{\alpha} : f(\alpha) \geq 0 \right\}$.
- α_{\max} denotes the largest manifest α .

Each of these points is depicted in the multifractal diagram shown in Figure 8.

4.4 Unifractal Examples

While the Legendre transform (8) can be worked out mechanically, a graphical representation is more instructive. This is obtained by plotting, for each q , a line of slope q and vertical intercept $-\widehat{\tau}(q)$. Placing all of these lines upon the same graph, the lower envelope gives the estimated multifractal spectrum.

For each of the Fractional Gaussian Noises in Figure 4, the simulated partition functions follow lines of approximate slope $\tau(q) = Hq - 1$. Figure 9 shows Legendre transforms of scaling functions for FGN's with self-affinity indices .4, .53, .6, .8. Only the left side of each spectrum is shown. The right side of the multifractal spectrum comes from negative moments of the partition function, which are not used in the empirical work in this paper. Each of the lines drawn passes through the same point at $\{\alpha, f(\alpha)\} = \{H, 1\}$. The lower envelope is thus degenerate:

$$f(\alpha) = \begin{cases} 1, & \alpha = H \\ -\infty, & \alpha \neq H \end{cases}$$

By inspection of equation (8), this form derives from the linearity of $\tau(q)$. The multifractal spectrum is defined only for values of α that are tangency slopes for some value of the scaling function.

A single valued multifractal spectrum is an alternative way of saying that each of these processes contains only a single Hölder exponent. Thus α_0 , which carries all of the Lebesgue measure of the

process, is the same as α_1 , which carries all of the variation of the process. These two values of α are distinct for multifractal processes.

5 The Multifractal Spectrum of DM/USD Exchange Rates

5.1 Estimation of the Scaling Function

Section 3.3 claimed evidence of a multifractal scaling function for DM/USD exchange rates based upon linearity of partition functions. Estimating slopes of the partition functions for many q gives the estimated scaling function $\hat{\tau}(q)$. Figures 6 and 7 indicated the relevant, approximately linear, estimation region as Δt greater than two hours. With nearly linear partition functions, the estimation method has little effect. We proceed by a combination of OLS and weighted least squares.

Observations of $S_q(\Delta t)$ may be inversely weighted by some measure of their apparent variability. We tried several reasonable weighting matrices, including no weighting at all, and found the differences in estimated slopes to be very small. Likewise, dropping observations from the beginning or the end of either sample, i.e., starting with a Δt of four hours or using only values of Δt up to 90 days, had little effect. For all the results presented, observations $S_q(\Delta t)$ are unweighted.²³ All values of $S_q(\Delta t)$ with $\Delta t > 2$ hours are used.

Separate scaling functions were estimated for the daily data and the high frequency data. We also considered regression with the slopes restricted to be identical between data sets. These results are not reported since the separately estimated scaling functions contain more information.²⁴

Figure 10 shows the estimated scaling functions. The solid line represents the high frequency data and the dashed line represents the daily data. Values of the scaling function were estimated at 99 values of q between .01 and 30, with more concentration at low values of q and values around 2. The scaling function is only estimated for positive values of q . Negative exponents greatly amplify small fluctuations, and thus are unreliable with imprecisely measured data. Some degree of measurement error is typical of financial data. For similar reasons, the estimated slopes for higher

²³The variance of the partition function at each value of Δt can be directly estimated at the same time the partition functions are calculated through the variance of the summands. We omitted this step, instead checking robustness to changes in weighting, which holds over a broad range of diagonal weighting matrices. Estimation of covariances probably requires bootstrapping.

²⁴Any value between the two unrestricted estimates can be obtained by changing the relative weighting of observations from the two data sets.

values of q can be expected to be less robust.

The graphs show several characteristics that are typical of estimated scaling functions. First, $\hat{\tau}(0) = -1$, which derives from the definition of the partition function. Secondly, the scaling function is concave. This is a consequence of Hölder's inequality, as shown in MFC. In addition, each of the estimated scaling functions is non-decreasing.²⁵

Of particular interest in each graph is the value of q where $\hat{\tau}(q) = 0$. This value of q identifies H by (6). Finally, each of the scaling functions is asymptotically linear, with a slope approximately equal to α_{\min} . The minimum α corresponds to the most irregular instants on the price path, and thus the riskiest events for investors.

5.2 Estimation of the Multifractal Spectrum

The multifractal spectrum is obtained by applying the Legendre transform (8) to the estimated scaling function $\hat{\tau}(q)$. Figures 11 and 12 show the estimated multifractal spectra, on the left side of α_0 , for the high frequency data and the daily data respectively. The right side is not estimated because this requires estimation of the scaling function for negative values of q . As mentioned previously, calculation of $S_q(\Delta t)$ with $q < 0$ is not appropriate for financial data. Small measurement error amplify greatly when using negative moments, thus requiring a high degree of precision from the data.

While the estimated spectra of high frequency and daily data are not exactly identical, strong similarities indicate that the MMAR can be used to model both samples. Both figures show a strong degree of multifractality. Whereas the unifractal examples of Section 4.2 converged to a spike at $\alpha = 1/H$, both estimated spectra of the DM/USD data are concave. From left to right on the horizontal axis, values of α increase from zero, corresponding to discontinuities in the process, to .6, corresponding to a greater degree of regularity than the Brownian Motion. The most commonly occurring Hölder exponent, α_0 , is greater than $1/2$ in both figures. Thus, on a set of Lebesgue measure 1, the continuous-time process corresponding to either spectrum is more regular than Brownian Motion. Nonetheless, both spectra have Hölder exponents less than $1/2$, reaching down to values of about $1/4$. These are very irregular instants for the price process, and represent high

²⁵Theoretical scaling functions are not restricted to be nondecreasing. This is a case of *virtual* α 's, mentioned in note 17.

risk for the investor.²⁶

In terms of trading time, low values of α correspond to sets of points where the measure of trading time per unit of clock time is very high. This translates into a low Hölder exponent for the c.d.f. of trading time at these instants, i.e. very swift passage of trading time. The Fractional Brownian Motion component of the MMAR has a constant Hölder exponent, so irregularity of the price process at instants with low α is due solely to trading time.

Each of the two spectra can be decomposed into $B_H(t)$ and $\theta(t)$. Equation (10) is the key. Having determined $H \approx .53$ by (6), the multifractal spectrum of trading time is derived by expanding the price spectrum along the horizontal axis by a factor of $1/H$. The next section discusses recovery of the MMAR components more thoroughly.

Turning to some of the differences between the two figures, the high frequency spectrum is more curved than the spectrum of the daily data, and its slope converges to ∞ approximately as $\alpha \rightarrow \alpha_{\min}$ from above. This type of behavior is typical of binomial and multinomial cascades. The spectrum of daily data is more typical of a *canonical* trading time, which assumes that the multipliers in a multiplicative cascade sum to one in expectation, instead of exactly.²⁷ Canonical sample paths typically have an abrupt transitions at approximately α_{\min} . This can be seen in the spectrum of daily data where all of the high moments pass through, or pivot at, a point approximated by $\{\widehat{\alpha}_{\min}, f(\widehat{\alpha}_{\min})\}$.²⁸ Thus, at a single point the spectrum (approximately) has a finite slope to the right but is undefined to the left.

That $\widehat{f}(\alpha)$ has negative sample values for the daily data can be attributed to sample variability. In general, the behavior of the spectrum around $\widehat{f}(\alpha) = 0$ is most variable since it is determined by high moments of the data. Not surprisingly, this is where the differences in the two estimated spectra appear most strong.

²⁶Moreover, because of the particular structure of fractal sets, the set of instants with a particular value of α tends to be clumped together, simultaneously generating more risk and long memory in volatility. This is because uniform spacing, when viewed from a distance, becomes uniformly dense spacing and demands space-filling (or Lebesgue Measure of 1) behavior to maintain self-similarity. Only irregularly spaced sets on the real line, like the Cantor set, can maintain self-similarity without filling space.

²⁷A *microcanonical* trading time shifts time locally, but preserves mass exactly at each step. *Canonical* measures allow random variations in mass that on average are conservative. This results in a random total mass of the measure. See MFC and CFM for further discussion.

²⁸While a particular draw of the random measure is expected to pivot at $f(\alpha_{\min}) = 0$, there is some variability in samples, and we see some spectra that pivot above or below zero. Note also that the LDT interpretation of the multifractal spectrum allows $f(\alpha) < 0$. See note 19.

Section 6 shows that $SEAS_2$, used to generate Figure 11, overadjusts in removing seasonal components from the high frequency data. This causes underestimation of the fractal dimension for low values of α . By employing other seasonal filters in Section 6, we can posit a range of possible values for the “optimally seasonal adjusted” high-frequency spectrum. This range appears much closer to the spectrum of the daily data, and thus the similarity in scaling laws is stronger than suggested by Figures 11 and 12.

5.3 Recovering the MMAR Components

Each of the two spectra in Figures 11 and 12 can be compared to the shapes of known multifractal measures (see CFM, Section 4.2). The spectrum of daily data is approximately quadratic, suggesting the limit lognormal multifractal measure. The limit lognormal has a long history, and was in fact the first multifractal developed by Mandelbrot (1972).²⁹

The limit lognormal measure is a multiplicative cascade whose multipliers are lognormally distributed. Following MFC and CFM, the generating process starts with a uniform mass on $[0, 1]$ and divides into b equal sized intervals, each taking a proportion of the original mass equal to $M(\eta)$, where $0 \leq \eta < b$ indexes the interval, and M is distributed i.i.d. lognormal across intervals. The multipliers need not sum exactly to one, but do have an expected sum of 1, which makes the construction canonical. At the second stage, each interval is further subdivided in the same manner. In the limit these iterations lead to the lognormal measure.

Following this construction, at the k^{th} stage of subdivision all intervals are of the size b^{-k} . We can denote an interval at the k^{th} stage by the b -adic expansion of its starting time

$$t = 0.\eta_1 \dots \eta_k.$$

²⁹Mandelbrot (1972) is a critique of theories put forward by Kolmogorov, Yaglom, and others concerning the distribution of turbulent dissipation over cubic eddies. The contending theories held that dissipation is identically lognormally distributed over cubic eddies of the same size. This parallels arguments put forward by Clark (1973) which proposed modelling trading time with i.i.d. lognormally distributed increments. As argued by Mandelbrot in 1972, this creates an inconsistent model, since the sum of independent lognormals is not lognormal. Hence, Clark’s model could only hold at a single time scale.

The arguments put forward in Mandelbrot (1972) are slightly more complicated. Previous theories had conjectured correlation patterns that allowed the sum of lognormals to be lognormal. Mandelbrot showed that these correlation patterns could not maintain self-similarity at different scales without introducing the concept of conservation of mass on the average, as opposed to local conservation, i.e. canonical vs. microcanonical. This observation began the study of multifractals.

The measure in an interval is the product

$$\mu(\eta_1 \dots \eta_k) = M_1(\eta_1) \dots M_k(\eta_k) \Omega(\eta_1 \dots \eta_k).$$

M_j denotes the multiplier at the j^{th} level of the cascade. The r.v. Ω is called the high-frequency component, and captures changes in total mass of the interval caused by stages beyond k . It can be shown that Ω has the same distribution for all levels of the cascade, and thus has the same distribution as the limiting mass in the interval $[0, 1]$.³⁰

For convenience, denote

$$V = -\log_b M \sim N(\lambda, \sigma^2).$$

CFM shows that the multifractal spectrum of a multiplicative cascade with lognormal multipliers is given by

$$f_\theta(\alpha) = 1 - \frac{(\alpha - \lambda)^2}{2\sigma^2 \log b}. \quad (11)$$

We can generate a further restriction by requiring average conservation of mass, so that

$$\frac{1}{b} = E(M) = E(e^{-V \log b}),$$

leading to

$$\log b = \frac{2(\lambda - 1)}{\sigma^2}.$$

Substituting this into (11) eliminates both b and σ^2 , giving

$$f_\theta(\alpha) = 1 - \frac{(\alpha - \lambda)^2}{4(\lambda - 1)}.$$

We now have a single parameter expression for the multifractal spectrum of trading time; however, we directly estimate the price spectrum which is obtained by applying equation (10) :

$$f_P(\alpha) = 1 - \frac{(\alpha - H\lambda)^2}{4H^2(\lambda - 1)}$$

Finally, using the substitution $\lambda = \alpha_0/H$ allows us to write the spectrum in terms of α_0 and H :

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

³⁰Convergence of mass to a random variable is guaranteed by the convergence of a positive martingale, or alternatively, by the ergodic theorem.

H is estimated by equation (6), which gives

$$\widehat{H} = 1.88^{-1}$$

for the daily data. This leaves us with only one parameter to fit.³¹ Using

$$\widehat{\alpha}_0 = .589$$

produces the parabola shown by circled stars in Figure 12. The fit with the actual data appears quite remarkable. The next section uses Monte Carlo methods to assess the variability of the spectrum.

5.4 Monte Carlo Simulation

The methods used for simulation are quite simple. (See MFC for other simulation examples). We approximate Fractional Brownian Motion with a Fractional Gaussian Noise, as in Section 2.3. For trading time, a finite multiplicative cascade with lognormal multipliers approximates a limit lognormal multifractal measure. The values of λ and σ^2 are recovered from \widehat{H} and $\widehat{\alpha}_0$ by the equations above.

Several simulated processes, with corresponding log first differences, are shown in Figure 13. The results are promising, with a variety of large asymmetric price changes, apparent trends, and other characteristics typical of real data.

Another goal of simulation is to test the partition function methodology. We calculated partition functions for twenty simulations with sample size identical to the data. Replication of the scaling properties of the data constitutes a joint test of the partition function methodology and simulation methods. About six or seven simulations followed a scaling law close to the data. Figure 14 shows four of these graphs. The plots are relatively straight, and vary somewhat around their predicted values. Some of the remaining simulations had nonlinear partition functions. The others had straight partition functions, but with different slope than the data.

Rejection of the MMAR is thus difficult, in a problematic sense, because of apparent low power. At the same time, the results are promising. Although variability is high, the simulated partition functions tend to follow their predicted lines. Improvements in either partition function

³¹Joint estimation of parameters will generally improve fit, by whatever metric is chosen. This point is relatively unimportant since sequential estimation produces good results.

methodology or simulation methods should improve power.³² Incorporating information from all moments of the data into a single statistic should have the same result.

At a more basic level, we may also consider that the data is better represented by an alternative scaling model with less variability. This is a reversal of the usual situation in economics, where data is typically much noisier than its first-approximation model. Strong empirical regularities of the type seen in the DM/USD data are worthy of further theoretical attempts at description.

Longer simulated records improve ability to distinguish between scaling in the MMAR and non-scaling in FIGARCH. A complete sample of twenty partition functions, each calculated from simulations with 10^5 increments, is shown in Figure 15.³³ Figure 15*a* shows simulated high moments, up to $q = 5$, and 15*b* shows moments around $q = 2$. The results are convincing. Most of the plots are approximately linear at high moments, (although there are a few notable exceptions). The graphs strongly tend to follow their predicted lines, perhaps with some inclination towards a downward bias. For low moments, the simulations are even more consistent with expectations.

Figure 16 shows estimated multifractal spectra for all twenty simulations. Dotted lines represent simulated data, and circles represent the theoretically prediction. The simulated spectra cluster around the predicted value, with more variability at low values of $f(\alpha)$, as expected.

From these experiments, we arrive at two conclusions. First, the MMAR theory, the simulation methods, and the estimation methods are all conceptually sound. Starting from the theory, going through simulation, and comparing estimates from simulation to their predicted values gives reasonable results. Secondly, there do appear to be convergence issues and possible biases, both of which must be studied more carefully. Overall, the results are positive. Both model and estimation method are in their infancy, and continued research should improve our ability to do inference. Future research should attempt to systematically describe, for various specifications of the model, convergence of the partition functions to their expected values as sample sizes increase.

³²A detailed discussion of simulation methods for multifractal trading times is beyond the scope of the present paper. In this work, we use a finite lognormal cascade, with a high frequency approximation to Ω . See MFC for further detail on this construction. Future work will treat simulation methods independently, and will explore alternative approximations to the limit lognormal measure. Better simulation methods may reduce variability.

³³In simulation, the value of H was $1/2$, rather than the value $.53$ from estimation. This allows us to forego computationally expensive simulation of the Fractional Brownian Motion. At lower sample sizes and this value of H , unreported tests indicate that the FBM does not add substantially to variability.

6 GARCH and FIGARCH Alternatives

This section simulates data from GARCH and FIGARCH representations. The parameter values come from previously published research on foreign exchange data. We find that these processes are distinct from the MMAR under simulation, and that the scaling properties of the DM/USD data are more consistent with the MMAR.

6.1 GARCH

Weak memory in GARCH has recently become interesting to researchers. Taylor (1986), Dacorogna et al. (1993), Baillie (1996), and others note that many financial time series display hyperbolically declining sample autocorrelation functions. In contrast, the GARCH decay rate is exponential. This finding is not a technicality or a matter of moment fitting, but has important qualitative implications. For example, under certain regularity conditions rescaled weak memory processes weakly converge to Brownian Motion.³⁴ This can be depicted visually.

Four GARCH(1,1) simulations, with parameters from Baillie and Bollerslev (1989), appear in Figure 17. The parameter values cover a typical range for GARCH(1,1) estimated on daily and weekly data. The plots contain 40,000 observations in order to show low frequency effects clearly. Convergence to Gaussian is apparent in 17*a-c*. The occasional spikes do not cluster except at high frequencies, and thus their contribution is washed out over long intervals.

Figure 17*d* is different. The parameter values used in this simulation are $\{\alpha_1, \beta_1\} = \{.072, .923\}$, which very nearly sum to one. Thus, squared returns are nearly integrated, or IGARCH, explaining noticeable low frequency effects. Still, the variations in volatility appear very unnatural, lacking significant high frequency variation.³⁵ Because weak-memory, even near integration, in volatility appears inconsistent with the data, we predict that GARCH will not replicate scaling patterns found in the data.

Expectations for the GARCH partition function are not easily calculated, so simulation is our main tool.³⁶ Partition functions of the previous simulations appear in Figure 18. None is

³⁴This is a generalization of Donsker's Theorem, which concerns the weak convergence of random walks. Intuitively, the claim is also related to Central Limit Theorems for weakly dependent sequences.

³⁵In this paper, "high frequency" should be interpreted relative to the scale of the graph or data.

³⁶For general GARCH(q, p) processes, $\mathbb{E}[S_q(\Delta t)]$ is a complicated expression in combinatorial sums. As $\Delta t \rightarrow \infty$, we expect GARCH to scale like Brownian Motion, since very large increments are approximately Gaussian with negligible dependence.

particularly straight, and repeated simulation (not shown) shows a tendency towards the downward bend observed in three examples. Moreover, the apparent slopes of the straighter plots are predicted lines for Brownian Motion, rather than the slopes of the data.

The GARCH trend component is omitted from the simulations in Figure 18. Figure 18b repeats the exercise, but includes the estimated trend. Clearly, the combination of trend and GARCH component is nonscaling.

GARCH(1,1) fails to capture important aspects of financial data. Temporal dependence becomes negligible too quickly as the sampling interval increases. Also, neither simulated GARCH nor sample partition functions of simulated GARCH replicate scaling features of the DM/USD data.

6.2 FIGARCH

FIGARCH improves on previous processes by parsimoniously obtaining long memory in volatility. It generates behavior in-between weak memory, typified by 17a-c, and unit root integration in volatility, approximated by 17d.

Baillie, Bollerslev, and Mikkelsen (1996), which also studies DM/USD exchange rates, prefers the FIGARCH(1, d , 0) specification. We use parameter estimates from this paper to generate the simulations in Figure 18. These improve upon Figure 17 in volatility correlations. On the other hand, the simulations are more symmetric than the data, and too systematic in volatility at high frequencies. Excessive regularity at high frequencies appears as thick clumps where increments of like size group together. In contrast, the DM/USD data contains clustering, but within clusters, subgroups of intervals cluster again into more or less variable groups. In turn, each subinterval can be broken down again, ad infinitum. This characteristic typifies multifractality, and is not captured by FIGARCH.

The FIGARCH(p, d, q) equations are

$$\begin{aligned}\varepsilon_t &= \sigma_t u_t \\ \phi(L)(1-L)^d \varepsilon_t^2 &= \omega + [1 - \beta(L)] v_t,\end{aligned}\tag{12}$$

where $\phi(L) \equiv [1 - \alpha(L) - \beta(L)](1-L)^{-1}$, $\alpha(L) \equiv \sum_{k=0}^q \alpha_k L^k$, $\beta(L) \equiv \sum_{k=0}^p \beta_k L^k$, $v_t \equiv \varepsilon_t^2 - \sigma_t^2$. In estimation, $\{u_t\}$ are typically assumed identically distributed, but not independent. Also,

their distribution is usually specified Student t rather than Gaussian. In simulation, we assume independence for lack of an explicit alternative, and use the Gaussian for u_t . These assumptions should not qualitatively affect our conclusions.

For FIGARCH(1, d , 0), (12) can be written more intuitively. Starting with

$$(1 - L)^d \varepsilon_t^2 = \omega + [1 - \beta_1 L] v_t,$$

rewrite

$$\begin{aligned} \sigma_t^2 &= \omega - \beta_1 v_{t-1} + \left[1 - (1 - L)^d\right] \varepsilon_t^2 \\ &= \omega - \beta_1 v_{t-1} - \sum_{k=1}^{\infty} \psi_k \varepsilon_{t-k}^2. \end{aligned}$$

where

$$\psi_k = \frac{\Gamma(k - d)}{\Gamma(k + 1)\Gamma(-d)}$$

The ARMA representation in volatility clarifies how simulation proceeds, and how long memory arises. (Plotting ψ_k will help readers gain further intuition for the process.) Because the Gamma function is smooth, the path of volatility is also smooth. This is the property that leads to clumping of like-sized returns.

Like ARFIMA, FIGARCH is a hybrid from the scaling point of view. For ARFIMA, intuition is clear since it is asymptotically equivalent in increments to Fractional Gaussian Noise. Short memory components are non-scaling, but become negligible over long sampling increments. This intuition disappears for FIGARCH since squared increments are asymptotically scaling, not increments themselves.

An alternative intuition is the literature on diffusion limits of ARCH sequences. The continuous time limit of discrete GARCH sequences is an Itô diffusion or jump-diffusion in certain cases. (Nelson, 1990; Drost and Nijman, 1996). Intuitively, GARCH is not scaling since Itô diffusions are not scaling.³⁷ Continuous time limits of FIGARCH, on the other hand, are not known.³⁸

Twenty simulated partition functions, appearing in Figure 19, are used to investigate the scaling properties of FIGARCH. Each simulation consists of 100,000 observations, allowing like comparison

³⁷A reference is not available, but intuitively, it results from smoothness of volatility.

³⁸Nelson (1990) notes that his method of demonstrating convergence to a continuous time diffusion does not apply to long memory processes.

with the MMAR (Figure 15). Moments up to $q = 5$ appear in 19*a*. Moments around $q = 2$ appear in 19*b*.

The high moments are irregular, and many show a tendency towards concavity. The partition functions for the MMAR, in contrast, are mostly straight. Some partition functions for FIGARCH approximately follow the best fit line of the data, although not as convincingly as MMAR partition functions. A reasonable degree of ambiguity suggests that developing stronger inference methods will be important in the future.

A more clear distinction exists in low moments. The MMAR partition functions around $q = 2$ follow predicted values very closely, while almost half of FIGARCH partition functions are quite irregular. Of the remaining partition functions, most are closer to the predictions for Brownian Motion than best fit lines of the DM/USD data.

This section investigated scaling properties of simulated GARCH and FIGARCH processes. Neither process appears as likely as the MMAR to scale like the DM/USD data. Barring theoretical results that allow distributional inference, more systematic simulation evidence is needed.

7 Robustness Tests

7.1 Alternative Reporting Methods

This section tests sensitivity of results to alternative reporting methods for the data, including price approximation by bid or ask or average, and source of data. We obtained daily data from the Federal Reserve Bank of New York covering approximately the same sample span as the Olsen daily data. The Fed data is recorded at noon in New York, and is a buying rate for foreign currencies. Because of missing data, two separate data sets were constructed. FED_1 supplements missing values with Olsen data. FED_2 eliminates missing observations. Appendix 9.3 provides more information on construction.

Figure 22 shows partition functions for FED_1 and FED_2 . Figure 23 compares the multifractal spectra of these series with the previously estimated spectrum of Olsen data. Results are not substantially affected by the use of these alternative data sets. Thus, reasonable variations, including treatment of missing values, use of bid or ask or an average, and use of price quotes or government reported rates, do not seem to substantially alter results.³⁹

³⁹Although the two daily spectra are quite similar, the Fed spectra are narrower. This might be expected based

7.2 Consistency in Subsamples

We split the twenty-four year sample span of FED_1 into two separate twelve year samples, 1973 – 1985, and 1985 – 1997. Figure 22 shows partition functions and Figure 23 shows estimated multi-fractal spectra. The results are very similar to full sample results, suggesting near stationarity in the model. This is an important result, and probably the most substantial test of robustness in DM/USD scaling.

Also of interest is the small, but noticeable, difference between the two subsamples. The break point occurs in 1985, a major turning point in international monetary policy. In February of 1985, a concerted intervention by major central banks reversed appreciation of the dollar, and marked a shift from “near complete flexibility to a managed float” (Catte, et al., 1994).⁴⁰ Consistent with this institutional shift, we notice a longer tail in the spectrum of the earlier, less managed, period.

7.3 Seasonal Adjustment Filters

This section tests robustness of scaling in the high frequency data to different deseasonalizing methods. We use four additional seasonal adjustment filters. Briefly, $SEAS_1$ adjusts by number of quotes received per fifteen minutes. Like the previously used $SEAS_2$, this method is *market-activity based*. The other three methods simply remove slow (on average) trading times from the data, expanding remaining time so as to keep total time constant. These are *crude* deseasonalizing methods. Details of construction are in Appendix 9.2.

Figure 20 shows graphs of the partition function around $q = 2$ for each new seasonal filter. Figure 21 shows partition functions for higher moments. For comparison with Figures 6 and 7, the format and aspect ratio are identical, and the high frequency plots are shown next to the same daily data as previously.

All methods, except to some degree $SEAS_4$, result in strongly linear partition functions. Crossover occurs at roughly the same point, again with the possible exception of $SEAS_4$, which appears to have a later crossover. The slopes of all partition functions appear slightly lower than in Figure 7. The lower slope is pronounced for $SEAS_4$. Lower slopes result in wider estimated

upon lower variability in the ask series than in the bid-ask midpoint. It would be interesting to see whether the same effect is present in other currencies.

⁴⁰This shift is clearest when viewed from the perspective of international agreements. Individual countries varied with respect to their exchange rate policies from 1973-1985.

multifractal spectra. Figure 22 plots envelopes of the estimated multifractal spectra for all five methods. The narrowest stands out among the rest, and was generated using $SEAS_2$.

This is probably the result of over-deseasonalization. $SEAS_2$ adjusts on the basis of absolute price changes, which also form the basis of our empirical tests. Due to a limited sample, the seasonality index is constructed with the same data that is later transformed by the index. Ideally, the seasonality index would be constructed with out of sample data, since large price changes are partially filtered away after they contribute to a high index of weekly activity for their time period.⁴¹

On the other hand, the three crude methods under-deseasonalize. Crude filters eliminate slow (on average) time periods, but do nothing to smooth remaining parts of the week. This leaves more irregularity in the data, leading to higher estimated fractal dimensions for low values of α , and thus a wider spectrum. The ordering of the three crude filters confirms this hypothesis: $SEAS_4$, which removes the fewest time periods, is the widest. $SEAS_0$, which removes the most time, is narrowest. Informally, we can think of $SEAS_0$ as an upper bound estimate of the spectrum (for low values of α), and $SEAS_2$ as a lower bound estimate.

Despite non-optimality of the seasonal adjustment filters, the range of variation appears tolerable. The multifractal scaling law in the DM/USD series is thus fairly robust to choice of deseasonalizing method. Estimated multifractal spectra are predictably affected by the seasonal adjustment procedure, and permit upper and lower bound estimates of an “optimally” deseasonalized multifractal spectrum.

7.4 The Japanese Yen/ US Dollar Series

Our final investigation involves the Japanese Yen/ US Dollar (“JPY/ USD”) series, also obtained from Olsen and Associates. We conducted a similar, but less thorough set of tests, and found this price record less consistent with the MMAR than is DM/USD. Figure 24 shows JPY/ USD, its increments and log increments. Figure 25 shows patterns in seasonality. Figure 26 shows the estimated partition function using $SEAS_2$ seasonal adjustment for the high frequency data and daily data from the Federal Reserve Bank of New York.⁴²

⁴¹Smoothing by moving averages would reduce this problem, but the exact method is a significant empirical problem on its own, and deserves a separate treatment.

⁴²The daily data is described in more detail in Appendix 9.3.

This last figure is a vignette, reproduced at a scale that will not overemphasize its importance. The irregularities in the partition function should not be viewed simply as failures of the model, but as opportunities to learn about different types of behavior and their causes. Comparing the JPY/USD series in levels with the DM/USD series, JPY/USD appears much smoother, particularly lacking in sharp local peaks. Institutional differences may matter.

On this point, a recent study of concerted interventions by G-3 Central banks finds almost twice the dollar value of intervention in the JPY/USD market as in the significantly larger DM/USD market.⁴³ While the connection between activism in exchange rate policies and scaling behavior has not been empirically established, it seems reasonable that the institutional rules or customs that govern a market affect the appearance of its price series. In turn, this may give a distinctive signature to partition functions or multifractal spectra. Comprehensive study of other data sets is the first step in studying this relationship.

More fundamentally, the MMAR does not exhaust the set of possible scaling models. While the JPY/USD series does not appear entirely consistent with the MMAR, no completely satisfactory alternative is apparent.⁴⁴ Scaling models may be modified in the future to describe a greater variety of behavior.

8 Conclusion

This paper has conducted the first empirical investigation of a multifractal model in economics. Consistent with the Multifractal Model of Asset Returns, we find evidence of a multifractal scaling law in an important financial data set, Deutschemark / US Dollar returns. The scaling properties of the data set are robust to variations in data reporting, including use of bids, asks, or averages as approximations to price; use of market price quotes or government reported exchange rates; and treatment of missing values. The DM/USD scaling properties are also robust to seasonal adjustment methods. Further, within subsamples of the data, the scaling laws are quite similar. We thus propose that a single, stationary model can provide a good approximation to the generating process of the data, covering not only a long and highly variable sample span (1973-1997), but also

⁴³Catte, et al. (1994) studies a sample of nineteen major concerted interventions from 1985-1991. Dollar interventions were about \$100 billion for JPY/USD vs. about \$54 billion for DM/USD.

⁴⁴The JPY/USD first differences are quite different from the FIGARCH simulations, the most obvious difference being the lack of variation in volatility at high frequencies in FIGARCH.

a broad range of sampling frequencies (two hours to six months).

In modelling financial prices, multifractality generates long memory in squared returns, long tails, and allows substantial flexibility in modelling either martingale or long-memory behavior in returns. Repeated simulations show that, among existing models of financial prices, the MMAR is most like typical financial data in its appearance. Further, the MMAR is best able to replicate the scaling law found in DM/USD returns.

The conclusion we draw is that multifractality is a promising new tool in economics. In the near future, it should be possible to use the MMAR in the evaluation of risk. With assumptions on a pricing model, we also envision pricing derivative assets through simulation methods. Naturally, this will require further work on estimation methods for multifractal processes, including convergence properties, simulation methods, and inference procedures. Another component of this research program is empirical study of other data sets. Given that all of these areas are relatively unexplored, there is considerable potential for new discovery.

9 Appendix

9.1 Goal of Seasonal Adjustment Filters

Mathematically, the multifractal trading time $\theta(t)$ is the c.d.f. of a statistically self-similar random measure. Statistically self-similar random measures have invariant conditional distributions of measure under isotropic contractions. In the MMAR, the conceptualization is fairly simple. The random measure, denoted μ_θ , is defined on $I^* = [0, T]$, an interval of the real line. Consider subintervals $I_1, I_2 \subseteq I^*$, and the class of affine transformations on the real line, denoted \mathcal{S} . We require invariance of conditional distribution of measure under affine transformation. (This rewrites Assumption 1 of MFC, Section 3.5.)

$$\frac{\mu_\theta[S(I_1)]}{\mu_\theta[I_1]} \stackrel{d}{=} \frac{\mu_\theta[S(I_2)]}{\mu_\theta[I_2]} \text{ whenever } I_1, I_2, S(I_1), S(I_2) \subseteq I^*$$

Seasonality violates statistical self-similarity. For example, for two one week intervals in clock time containing identical measures of trading time, the distribution of trading time in the first seventh of the interval clearly depends upon whether that particular interval begins on a Sunday or a Wednesday. Thus the goal of the seasonal adjustment filters is to remove weekly patterns in trading time.

9.2 Construction of Seasonal Adjustment Filters

Since trading time is an abstract concept, deseasonalizing filters are based upon proxy measures of market activity. The strategy followed in this paper is to use several reasonable seasonal adjustment methods and thereby test the sensitivity of our results. To limit the number of alternatives, we do not investigate transformations of the raw indices of market activity. The five prefilters used in this paper are summarized in the following table

Table 1: Seasonal Filters

$SEAS_0$ - Crude; Removes (2100 Fri-2300 Sun) and (weekdays 300-430)
$SEAS_1$ - Market activity; Deseasonalizes by average number of quotes
$SEAS_2$ - Market activity; Deseasonalizes by average absolute price change
$SEAS_3$ - Crude; Removes (2300 Fri-1800 Sun) and (weekdays 300-400)
$SEAS_4$ - Crude; Removes (2300 Fri-1800 Sun)

The *market activity based filters* use some measure of market activity as a proxy for trading time. By expanding periods of high market activity and contracting periods of low activity, the goal is to create a new time scale under which, on average, all times of the week have the same measure of market activity per unit of time. $SEAS_1$ and $SEAS_2$ are based upon the market activity measures in Figures 5a and 5b. Respectively, these are the average number of quotes per 15 minute interval, and the average absolute price change over 15 minute intervals. In the graph, the vertical axis gives the trading time/clock time ratio that would be required to (approximately) eliminate weekly seasonality in that particular measure. For example, a value of two on the vertical axis indicates that, on average, the relevant time interval is twice as active per unit of time as the average time interval in the week.

The *crude filters* simply remove sparse periods of market activity from the data set. Remaining time periods are uniformly expanded by an amount which returns the data set to its original time span. “Sparse” time periods were determined for $SEAS_0$ by printing out all time periods during which no quotes arrived for a period of at least 30 minutes. The most frequently occurring times of the week with thirty minute lulls were determined by inspection, and occur from 9 p.m. Friday to 11 p.m. Sunday, plus the period from 3:00-4:30 AM on weekdays, during which activity in Asian markets slows because of the lunch hour, and no other markets are open. Sparse time periods for $SEAS_3$ and $SEAS_4$ were determined after inspecting the measures of market activity in Figure 5. $SEAS_3$ eliminates less of the sparse time periods than $SEAS_0$ and $SEAS_4$ eliminates only weekends, leaving in the period around the Asian lunch hour.

9.3 Daily Data Series

The daily data used in Sections 3 to 5 was provided by Olsen & Associates. The set contains DM/USD rates for every weekday between June 4, 1973 and December 31, 1996 except Christmas and New Year’s Day. The price is recorded as a linear interpolation of the quotes immediately preceding and immediately following 16:00 GMT. The geometric mean of bid and ask proxies for price, and there are 6119 total observations in the data set.

Section 6 uses daily data from the Federal Reserve Bank of New York. The recorded price is the noon buying rate in New York, used for customs purposes. The data set spans June 4, 1973 to June 3, 1997. The following table provides information on missing values and the construction of

the two series FED_1 and FED_2 . The table is broken down into two twelve year periods, reflecting the subsamples used in section 6.2.

Table 2: DM/USD Daily Data from NY FED

Years	Obs	Nmiss	Nreplace
1973-1985	3131	126	110
1985-1997	3131	114	93

FED_1 replaces missing values with Olsen data when it is available. The “Nreplace” column gives the number of replacement observations used. When no replacement value is available, we use the previous day’s price. FED_2 simply eliminates all missing values from the data set.

The JPY/ USD data used in Section 6 contains the same number of observations and missing values as the DM/USD data, and also comes from the Federal Reserve Bank of New York. An alternative data source was not used for the JPY/ USD data. The reported results replace missing values with the previous day’s price. We performed the same tests after eliminating all missing values, and found the results not to be substantially different.

References

- [1] Bachelier, L. (1900), Théorie de la Spéculation, *Annales de l'Ecole Normale Supérieure* **3**, Paris: Gauthier Villars. English translation in Cootner (1964)
- [2] Baillie, R.T. (1996), Long Memory Processes and Fractional Integration in Econometrics, *Journal of Econometrics* **73**, 5-60
- [3] Baillie, R. T., Bollerslev, T., and Mikkelsen, H. O. (1996), Fractionally Integrated Generalized Autoregressive Conditional Heteroskedasticity, *Journal of Econometrics* **74**, 3-30
- [4] Baillie, R.T., and Bollerslev, T. (1989), The Message in Daily Exchange Rates: A Conditional Variance Tale, *Journal of Business & Economic Statistics* **7**, 297-305
- [5] Bollerslev, T. (1986), Generalized Autoregressive Conditional Heteroskedasticity, *Journal of Econometrics* **31**, 307-327
- [6] Bollerslev, T., Engle, R. F., and Nelson, D. B. (1994), ARCH Models, ch 49 in: R. Engle and D. McFadden eds., *Handbook of Econometrics*, Vol.4, Amsterdam: North-Holland
- [7] Bollerslev, T., and Mikkelsen, H.O. (1996), Modeling and Pricing Long Memory in Stock Volatility, *Journal of Econometrics* **73**, 151-184
- [8] Calvet, L., Fisher, A., and Mandelbrot, B.B. (1997), A Multifractal Model of Asset Returns, Working Paper, Yale University
- [9] Catte, P., Galle, G., and Rebecchini, S. (1994), Concerted Interventions and the Dollar, in *The International Monetary System*, ed. P. B. Kenen, F. Papadia, and F. Saccomanni, Cambridge: Cambridge University Press
- [10] Clark, P. K. (1973), A Subordinated Stochastic Process Model with Finite Variance for Speculative Prices, *Econometrica* **41**, 135-156
- [11] 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
- [12] Drost, F.C., and Nijman, T.E. (1993), Temporal Aggregation of GARCH Processes, *Econometrica* **61**, 909-927
- [13] Drost, F.C., and Werker, B.J., Closing the GARCH Gap: Continuous Time Garch Modeling (1996), *Journal of Econometrics* **74**, 31-58
- [14] Engle, R. F. (1982), Autoregressive Conditional Heteroscedasticity with Estimates of the Variance of United Kingdom Inflation, *Econometrica* **50**, 987-1007
- [15] Evertsz, C. J. G., and Mandelbrot, B. B. (1992), Multifractal Measures, in: Peitgen, H. O., Jürgens, H., and Saupe, D. (1992), *Chaos and Fractals: New Frontiers of Science*, 921-953, New York: Springer Verlag
- [16] 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

- [17] Ghysels, E., On the Economics and Econometrics of Seasonality (1994), *Advances in Econometrics, Sixth World Congress, Vol. 2*, ed. C. A. Sims, Cambridge: Cambridge University Press
- [18] Ghysels, E., Gouriéroux, C., and Jasiak, J. (1995), Market Time and Asset Price Movements Theory and Estimation, Working Paper
- [19] Ghysels, E., Gouriéroux, C., and Jasiak, J. (1996), Trading Patterns, Time Deformation and Stochastic Volatility in Foreign Exchange Markets, *CREST Working Paper n° 9655*
- [20] Granger, C.W.J.(1980), Long Memory Relationships and the Aggregation of Dynamic Models, *Journal of Econometrics* 14, 227-238
- [21] 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, 15-31
- [22] Grether, D.M., and Nerlove, M. (1970), Some Properties of ‘Optimal’ Seasonal Adjustment, *Econometrica* 38, 682-703
- [23] 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
- [24] Hamilton, J.D. (1994), Time Series Analysis, Princeton: Princeton University Press
- [25] Hosking, J.R.M. (1981), Fractional Differencing, *Biometrika* 68, 165-176
- [26] Mandelbrot, B. B. (1963), The Variation of Certain Speculative Prices, *Journal of Business* 36, 394-419
- [27] Mandelbrot, B. B. and Taylor, H. W. (1967b), On the Distribution of Stock Price Differences, *Operations Research* 15, 1057-1062
- [28] Mandelbrot, B. B., and Ness, J. W. van (1968), Fractional Brownian Motion, Fractional Noises and Application, *SIAM Review* 10, 422-437
- [29] Mandelbrot, B.B., and Wallis, J.R. (1969), Computer Experiments with Fractional Gaussian Noises, *Water Resources Research* 5, 228-267
- [30] Mandelbrot, B.B. (1971), A Fast Fractional Gaussian Noise Generator, *Water Resources Research* 7, 543-553
- [31] 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
- [32] Mandelbrot, B. B. (1974), Intermittent Turbulence in Self Similar Cascades; Divergence of High Moments and Dimension of the Carrier, *Journal of Fluid Mechanics* 62
- [33] Mandelbrot, B.B. (1976), Note on the Definition and Stationarity of Fractional Gaussian Noise, *Journal of Hydrology* 30, 407-409
- [34] Mandelbrot, B. B. (1982), *The Fractal Geometry of Nature*, New York: Freeman

- [35] Mandelbrot, B.B. (1988), A Class of Multinomial Multifractal Measures with Negative (Latent) Values for the “Dimension” $f(\alpha)$, *Fractals’ Physical Origins and Properties*, Proceedings of the Erice Meeting, ed. L. Pietronero, New York: Plenum Press, 3-29
- [36] Mandelbrot, B. B. (1989a), Multifractal Measures, Especially for the Geophysicist, *Pure and Applied Geophysics* **131**, 5-42
- [37] Mandelbrot, B. B. (1990), Limit Lognormal Multifractal Measures, in: E. A. Gotsman et al. eds., *Frontiers of Physics: Landau Memorial Conference*, 309-340, New York: Pergamon
- [38] Mandelbrot, B.B. (1995), Negative Dimensions and Hölders, Multifractals and their Hölder Spectra, and the Role of Lateral Preasymptotics in Science, *Journal of Fourier Analysis* Special J.P. Kahane Issue, 409-432
- [39] Mandelbrot, B. B. (1997), *Fractals and Scaling in Finance: Discontinuity, Concentration, Risk*, New York: Springer Verlag
- [40] Mandelbrot, B. B., Calvet, L., and Fisher, A. (1997), Large Deviation Theory and the Multifractal Spectrum of Financial Prices, Working Paper, Yale University
- [41] Melino, A. (1994), Estimation of Continuous-Time Models in Finance, in *Advances in Econometrics, Sixth World Congress, Vol. 2*, ed C.A. Sims, Cambridge: Cambridge University Press
- [42] Mills, T. C. (1993), *The Econometric Modelling of Financial Time Series*, Cambridge: Cambridge University Press
- [43] 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
- [44] Nelson, D.B. (1990), ARCH Models as Diffusion Approximations, *Journal of Econometrics* **45**, 7-38
- [45] Nelson, D. B. (1991), Conditional Heteroskedasticity in Asset Returns: A New Approach, *Econometrica* **59**, 347-370
- [46] Nelson, D.B., and Foster, D.P. (1994), Asymptotic Filtering Theory for Univariate ARCH Models, *Econometrica* **62**, 1-41
- [47] Rossi, P. E. (1996), *Modelling Stock Market Volatility, Bridging the Gap to Continuous Time*, New York: Academic Press