ELSEVIER

Physica A 263 (1999) 477–487

PHYSICA A

Renormalization and fixed points in finance, since 1962

Benoit B. Mandelbrot

Yale University, New Haven CT 06520-8283 USA

Abstract

In diverse sciences that lack Hamiltonians, the analysis of complex systems is helped by the powerful tools provided by renormalization, fixed points and scaling. As one example, an intrinsic form of exact renormalizability was long used by the author in economics and related fields, most notably in finance. In 1962-3, its use led to a model of price variation founded on the (Cauchy-Polyà-Lévy) stable distribution, with striking data collapse that accounted for observed large deviations from Gaussianity. In 1965, a different form of exact renormalization led to fractional Brownian motion, which neglected large deviations but accounted for long dependence and the resulting non periodic cyclic behavior. Finally, from a seed planted in 1972, exact renormalizability and scaling led to a model of price variation of which the M1963 and M1965 models are special examples. This broader model, fractional Brownian motion in multifractal time, accounts simultaneously for both large deviations and long dependence. These three steps are in loose parallelism with space, time and joint renormalization in statistical physics. This presentation surveys the old works and many new developments described in the author's 1997 books on fractals and scaling in finance.

Deep interest attaches, at this point in time, to the question of "Has statistical physics already contributed to economics and finance?" Work I did around 1960 answered this question by *Yes*, as documented in [1][2] and sketched in Chapter 37 of [3]. Additional insight is provided by [4] and [5], and developments that followed [1] and [2] are described in [6]. This paper proposes to amplify the above *Yes* and indicate how [1] and [2] might be of assistance in core physics.

1 A rare case of historical precedence (without direct influence).

My major contribution to economics was to show how price variation on financial markets can be "accounted for" solely on the basis of renormalization and suitable fixed points. The trigger was the "folklore" that maintains that all records of (say) 500 successive values of a price look alike, independently of the time increment Δt between observations. I interpreted the passage from a given Δt to a multiple thereof in terms of classical thermodynamical "coarse-graining." Three successive steps taken around 1963, 1965 and 1972 centered on "privileged models" that are fixed points under three increasingly general

PII: S0378-4371(98)00520-2

forms of renormalization. The most general yields a parsimonious process called fractional Brownian motion in multifractal trading time. This process and its generalizations add up to a multifractal model of asset returns (MFAR).

The dates, 1963, 1965, and 1972. They suffice to show that this work was performed or at least begun before the flowering of the study of critical phenomena. Like most scientists, I did not hear of the classic papers by Wilson and Wilson and Fisher until they were about to be printed. (The information came from Herbert Callen, who heard me tell my story as the evening speaker in an American Physical Society meeting.) My inspiration came from outside of contemporary mainstream physics, namely, from old statistical thermodynamics, old mathematical esoterica going back to Cauchy, and old empirical esoterica going back to Pareto.

The cultural context. Much of that work was completed or undertaken in time for a visiting professorship of Economics at Harvard (1962-3). Aside from [7], the key papers on the M1963 model appeared in mainstream economics periodicals. Reprints are found in the second half of [1]. Photographic reproductions of advance abstracts ([7,9]) are found in [2]. The follow-up of [9] is found in [5].

My insistence on long-tailedness (and later on long dependence) was discussed very energetically in the economics literature. A basic anthology [10] includes the following words by the editor, Paul H. Cootner (M.I.T. School of Management): "Mandelbrot ... has forced us to face up in a substantive way to those uncomfortable empirical observations that there is little doubt most of us have had to sweep under the carpet up till now. He has marshalled evidence of a more complicated and much more disturbing view of the economic world than economists have hitherto endorsed. There can be no doubt that his hypotheses are the most revolutionary development in the theory of speculative prices since Bachelier's work in 1900."

Bachelier and Brownian motion. Louis Bachelier's name is mysterious to physicists, but his Ph.D. thesis in mathematics was both the first work in quantitative finance and the first treatment of Brownian motion. This last term came up only after this process was rediscovered in statistical physics, and explored mathematically by Wiener — hence the term Wiener Brownian motion (as contrasted to fractional Brownian motion). But the concept first was first fully articulated as an "ideal market hypothesis" relative to the bond market.

Physics (statistical or not) has influenced economics on innumerable occasions and very deeply. Therefore, little surprise would be generated by a claim that statistical physics can find new applications in finance. In fact, Bachelier clearly predated physics, though he had no influence on it whatsoever. A second instance of the same kind is provided by my work sketched in this paper. It remains true that scaling and renormalization went much farther in physics, but this added sophistication does not appear to be immediately transferable to economics.

2 The distinction between mild and wild variability, a new concept from economics/finance that may help in core physics.

Inverting the roles of two fields and changing from past to future, a converse question arises, "Has the study of finance based on scaling (etc.) contributed notions of its own that might in the future be transferred to core physics?" This question can be answered by *Yes* on several accounts, of which one will be examined.

Chapter E5 of [1] and Part II of [2] introduce the concept of "states of variability and randomness." To quote from p. 120 of [1]. "While a unique theory of physical interactions applies to every form of matter, the detailed consequences of those unique general laws differ sharply [and] physics has to distinguish between several *states of matter*. [I] argue that a similar distinction should be useful in probability theory. Nearly every scientist engaged in statistical modeling ... used to deal with a special form of randomness, which will be characterized as *mild*. It will also be argued that entirely different states of randomness must be distinguished and faced, namely, *wild* and *slow*. Mildness is characterized by an *absence of structure* [recalling] a gas. Wildness [is] characterized by the presence of structure [recalling] a solid [and slow randomness recalls liquids.]"

The ideal market is the prototype of mild variability, and there are other mild models. But all the other models considered in this paper are wildly variable. The distinction, neither philosophical nor simply metaphorical, concerns the fundamental fact that risk is reduced by averaging. In mild randomness, the central limit theorem says that this reduction proceeds as $1/\sqrt{N}$, N being the number of items being averaged. In wild randomness, the typical reduction is smaller or nonexistent. Moreover, the variability around the typical value is greater and can be very large, accounting for discontinuity and concentration in economics. There is no space to dwell further on this effect, which, once again, is the topic of [1] and also of [4] and [5].

I view wild variability/randomness as one of the frontiers of scientific knowledge.

3 From financial folklore to the ideal market hypothesis on to a challenge.

Description of some characteristics of the charts that illustrate price records. In the past, scientific models and reality could be compared via short lists of numerical quantities. Today the actual data and simulated samples of the models can also be displayed side by side and compared visually. To be sure, it is even easier to lie by pictures than by words, numbers and statistics, but Figures 1 and 3 cannot lie. Incidentally, the eye discriminates better between price *increments* than between price records themselves. The "pen width" in Figure 1 is of the same order as the time lag between observations. The resulting "strip" is an artefact but a very useful one. The top line illustrates white noise (increments of Brownian motion). The lines M3 and M4 illustrate the M1963 and M1965 models to be described momentarily. The remainder is a medley: at least one is a real record and at least one is a forgery, that is, a computer generated sample of the M1972 model. The reader is free to guess which diagrams are real and which ones forgeries. It is hoped that the forgeries will be perceived as surprisingly effective, but only lines BI and DD are real: they refer to the prices of IBM shares and the dollar-deutschmark exchange rate.

Many features observed throughout the medley at the bottom of Figure 1 are familiar even to those who know little about financial markets. Firstly, neglecting the "spikes," a high proportion of relative changes falls in a strip whose width constantly varies as a token of variable "volatility" (independently of the definition chosen for that very vague concept). Secondly, a substantial number of spikes stand out clearly outside the strip: all correspond to unusually large price changes and many correspond to instantaneous

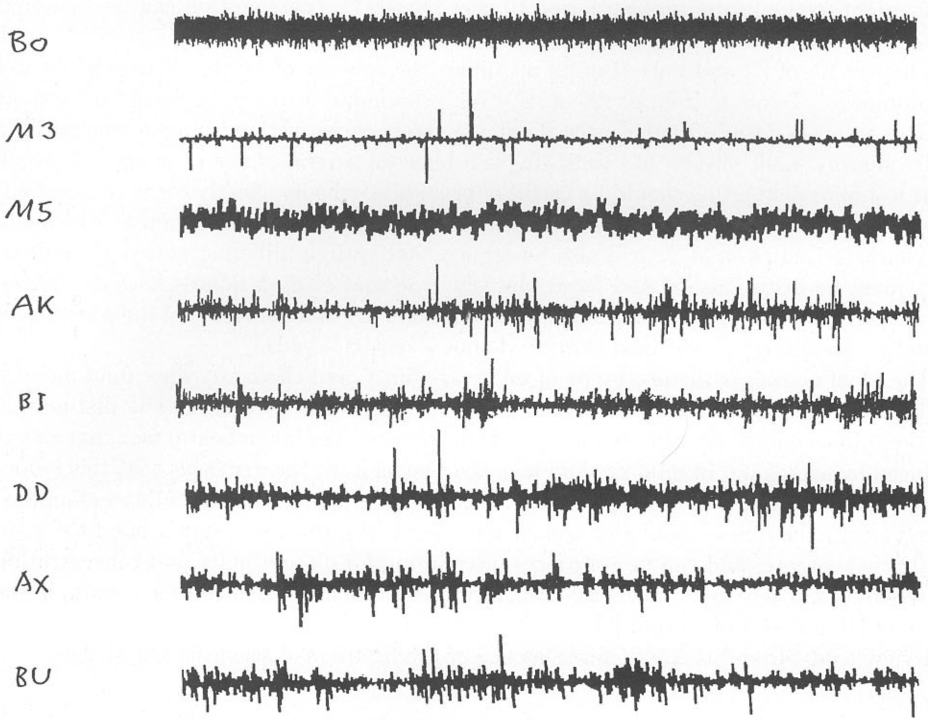

Fig. 1. Stack of graphs representing price increments under several old and new models, and real data. Sec. 3 solves the mystery created by the absence of labels.

discontinuities. Thirdly, the spikes tend to cluster and to occur during periods when the strip is broad, hence the volatility is large. An ideal market model that is correct 95% of the time is not sufficient if one agrees that most major events concern the 5% remainder. That is, the study of finance cannot be blind to extreme price changes.

Those extremes deserve to be documented. IBM saw its stock fall instantaneously by 10% early in 1996, and later in that year rise instantaneously by 13.2%. Concentration with or without discontinuity is striking even in the extensively averaged portfolio based on the Standard & Poor 500 index. Of this portfolio's returns over the 1980s, fully 40% was earned during ten days (0.5% of the number of trading days in a decade.)

The ideal market hypothesis: Bachelier's Brownian motion model. The preceding raw numbers are best understood against the fact that of the many characteristics of real markets, Bachelier's ideal market incorporates none, hence represents the observed price series $P(t)$ very poorly.

For example, the white noise strip on the top of Figure 1 is of constant width and no spike stands out. Suppose that both an ideal and a real chart have a width of the order of 2σ. Real charts often show spikes of the size 10σ. In an ideal market where σ is defined as the standard deviation, those spikes would have a probability of 10^{-23}: roughly the inverse of Avogadro's number. The ideal market disregards them, but a realistic model cannot. As to the concentration of the growth of the Standard & Poor 500 index in 0.5% of the number of trading days, it is contradicted by a fundamental theorem about the ideal market, namely that even the most active day makes a negligible contribution.

A further criticism of the ideal market hypothesis is qualitative but deep. Financial dailies, weeklies, and monthlies can thrive because one day in the market is unlike any other day, one week, month or year unlike any other. In an ideal market, to the contrary, daily chapters of the history books might vary from one another, but all monthly and yearly chapters would seem effectively alike.

A challenge. The limitations of the ideal market hypothesis are generally acknowledged but the ideal model using $B(t)$ became the basis of a very sophisticated "modern portfolio theory" and of a "calculus of risk" highlighted by the Black-Sholes-Merton theory. This unavoidable development sensibly followed an example set by the exploration of matter, which began by inventing the concept of perfect gas. Today, the time to move on to realistic models has come.

This effort, like every feature of financial markets, pits bulls against bears. To the question, "can large events be handled quantitatively?", the bears answer that this is impossible, large events being individual "acts-of-God" or "anomalies" presenting no conceivable regularity. My work, to the contrary, takes the decidedly bullish position that the variation of financial prices can be accounted for.

4 Self-affinity, self-affine interpolation and cartoons.

A classical invariance property of the Wiener ("ordinary") Brownian motion $B(t)$. It is well-known that, if $B(0) = 0$, the function $|\mu|^{-H}B(\mu t)$ has the same distribution for all $\mu \neq 0$, in particular for $\mu = 1$, which yields $B(t)$ itself. In the 1960s, this special form of scaling being standard but nameless, I put forward the term, *self-affinity*, which was generally accepted. The Brownian self-affinity exponent is $H = 1/2$.

Self-affinity As is widely known, a fractal is a geometric shape that can be separated into parts such that each part is a reduced-scale version of the whole. To implement this characterization, one must define the notion of reduction. Fractals using isotropic reduction are called *self-similar*. They have become well-known, but prices call for the more novel concept of self-affine fractality.

Self-affinity expresses mathematically the claim that all market charts look alike. The "whole" chart is usually wider than it is high, but its small parts are higher than they are wide. That is, in order to move from the whole to a part, one must reduce the time scale far more than the price scale. Self-affine fractality makes available many powerful tools of analysis. Some are very new; others are described in [3,1,4,5].

"Privileged models." From mathematical esoterica that I knew it followed that, if one ceases to impose either $H = 1/2$ or continuity, self-affinity holds more broadly than $B(t)$. I interpreted Cauchy's broader solution, and these additional ones that had to be developed, as being "privileged models," observed that all involve diverse forms of scaling, and compared them with the empirical evidence. "Privileged models" succeeded beyond expectation, improving through three successive forms of exact renormalizability. The M1963 model assumed price changes to be independent, and focussed on the non-

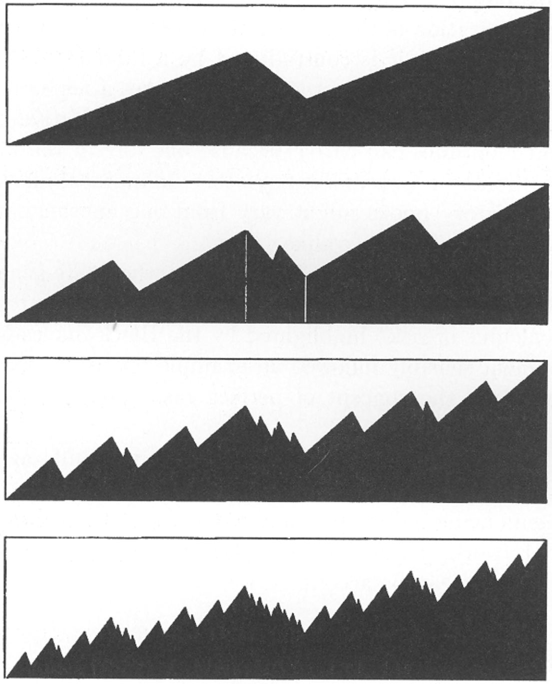

Fig. 2. Recursive construction of a non-random grid-bound cartoon of Wiener Brownian motion (see text for explanation).

Gaussian long-tailedness. The M1965 model assumed price changes to be Gaussian, and focussed on their very long dependence. The M1972 model incorporated the earlier models as special cases, and went beyond by combining long tails and long dependence.

"Privileged models" do not explain, but show fundamentally important observed features of real data to be *necessary consequences* of suitable exact renormalizability. This approach is far preferable to those common would-be improvements of the Brownian model, in which basic aspects of reality s erve as *input*.

Implementation of self-affinity by "cartoons." As shown in Figure 2, a cartoon consists of interpolating a price change from time 0 to time 1 by successive steps. The construction begins with an initiator, that is, a square crossed diagonally by a straight trend line. To avoid clutter and make this illustration easier to read the square is stretched horizontally into a rectangle and the trend is not drawn. Next, this trend is interpolated by a generator, which here is a broken line formed of three intervals, whose lengths $(\Delta_i t, \Delta_i x)$, from left to right, are $(4/9, 2/3)$, $(1/9, 1/3)$, $(4/9, 2/3)$. Each interval is again interpolated by three shorter ones, reproducing the generator after it has been squeezed to fit the same endpoints. The white lines in the second figure from the top emphasize the different squeezings.

Since each interval of the first approximation ends up with a shape just like the whole,

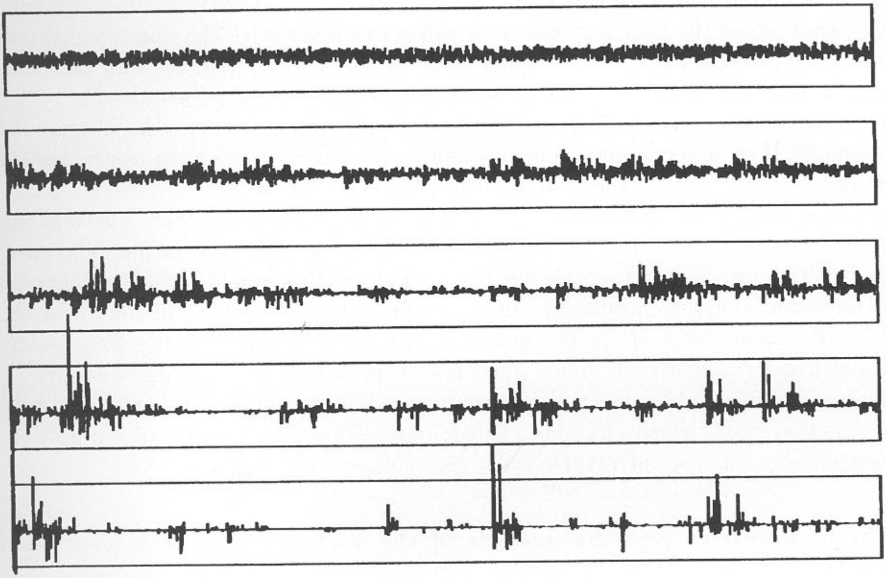

Fig. 3. Stack of graphs representing (from the top down) the increments of a cartoon of Wiener Brownian motion, then those of several cartoons of Wiener Brownian motion occuring in increasingly "wild" multifractal time.

scale invariance is present because it was built in.

Unifractality in grid-bound cartoons. The computer-generated diagrams in Figure 3 are plotted in the form of price increments. The top one corresponds to a "randomized" variant of Figure 2. It changes the generator before each interpolation by scrambling at random the sequence of its three intervals. Figure 3 demonstrates that great wealth of structure can be generated by an extremely simple formula. This unexpected possibility is crucial to fractal geometry and to the theory of chaos.

In Figure 2, $\Delta_i x = \sqrt{\Delta_i t}$ for each generator interval. The same is true of every interval of the "prefractal" broken line constructed at a later stage of recursion and is approximately the case for arbitrary increment or in the fractal limit. This "square root" property also characterizes Brownian "diffusion" and expresses that increments are not correlated.

In a first generalization, $\Delta_i x = (\Delta_i t)^H$, where $H \neq 1/2$ and $0 < H < 1$. The resulting constructions, called "unifractals," are cartoons of processes called "fractional Brownian motions."

Multifractal trading time. The idea of the next generalization is that trading time does not reduce to clock time but flows slowly on some days and very fast on some others. Price, in turn, is a Wiener or fractional Brownian motion in trading time.

An implementation of multifractal trading time that preserves self-affinity consists in modifying the generator as follows: the intervals' vertical extents are not affected, but their horizontal extents are either lengthened or shortened, while preserving a sum equal to 1. The final effect of such slowing and speeding up is illustrated by three of the bottom curves in Figure 3. Going down the stack, those cartoons diverge increasingly from Brownian motion. To increase the strip's spikiness, it suffices to begin with the above pseudo-ideal three-interval generator and move its first break to the left! By playing with cartoons of this sort, a multiplicity of levels and kinds of volatility can be obtained at will.

A last step consists in allowing the intervals of the cartoon generator to be either vertical or horizontal. If so, a very strong property can be proved: every cartoon can be analyzed into a unifractal cartoon of a multifractal time. (Section 3.5 of [4].)

It remains to describe quantitatively the effect of these manipulations. The velocity of a self-affine fractal curve is not defined at most points, vanishes at some and is infinite at others. It is far more helpful to follow the mathematician Hölder and divide the logarithms of the increments of the coordinates. In the unifractal case, their ratio converges at all points to the same value H. In the multifractal case, this limit α varies from point to point and usually is neither 0 or ∞. Speeding yields $\alpha < 1$ and slowing down, $\alpha 1$. The well-known multifractal formalism leads again to a function $\tau(q)$ and a generalized form of the function $f(\alpha)$, which is a kind of distribution function. Particularly important is an exponent H_T, the root of $\tau(1/H_T) = 1$. (See Chapter E6 of [1].)

5 Three successive grid-free models of the variation of prices on financial markets.

Last section's grid-based cartoons are extraordinarily useful in creating intuition but cannot be viewed as the last word. Furthermore, a full description of a grid-based cartoon involves more than the value of H_T and the function $f(\alpha)$. Therefore, the ideal of parsimony makes it desirable to forsake the cartoons' simplicity and convenience, and prefer processes that are fully determined by H and $f(\alpha)$, with no artificial grid. At least four such models were considered even before the cartoons were introduced.

The ideal market hypothesis: Bachelier's Brownian motion model. This model's counterparts are all the multifractal cartoons with $H = 1/2$.

The M1963 model for the observed long-tailedness. In 1962-3, [6] and [7] proposed a model of price variation founded on the (Lévy) stable distribution, with striking data collapse that accounted for observed large deviations from the Gaussian "norm." As W. H. Taylor and I realized in 1967 (see Chapter 21 of [1]), this process can be represented as Brownian motion in a *fractal* trading time. As a result, this model's cartoon counterparts are those whose generators include both vertical intervals and slanted intervals with $H = 1/2$.

It should be recalled that the symmetric stable distributions are such that, if $\phi(s)$ is the Fourier transform of the density (the "characteristic function"), there is an exponent α, with $0 < \alpha \leq 2$, such that $\log \phi(s) = -\gamma|s|^\alpha - \delta$. Here δ specifies position and γ specifies scale. In the general case, there is a fourth parameter β that specifies skewness.

The classic test of scaling using "collapse" of graphs, found a very early use in Figure 4. It was first presented in [8] (and in preliminary form in [7]), and later reproduced in [10], on Plate 340 of [3], page 391 of [1] and page 24 of [2].

A simulation of the M1963 model is shown on the second line of Figure 1. Consider the large deviations from the norm; in many cases, their size is accounted for correctly, but they are not clustered in time, contrary to the actual data represented in the bottom part

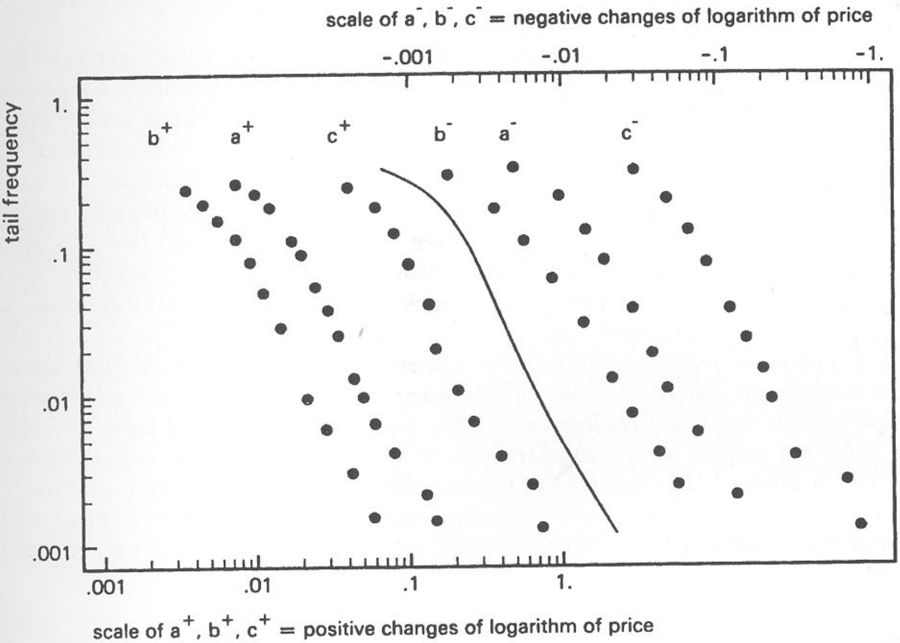

Fig. 4. Collapse of the graphs of daily and monthly price changes for cotton. A full legend is found in [1], page 32.

of Figure 1. For data to which the M1963 model does not apply, there is a temptation to "fix" the marginal distribution (usually by truncating the data beyond a threshold) without tackling the absence of clustering in time. This destroys scaling, instead of generalizing it (as I did in the M1972 model).

The M1965 model for the observed long dependence. In 1965 ([9], see also [1]), a different form of exact renormalization led to fractional Brownian motion (FBM), which neglected large deviations but accounted for long dependence and the resulting non periodic cyclic behavior. The FBM process $B_H(t)$ is fully specified by the property that $B_H(t') - B_H(t)$ is a Gaussian variable having zero mean and satisfying $\langle [B_H(t') - B_H(t)]^2 \rangle = |t - t'|^{2H}$, where $0 < H < 1$ and $H \neq 1/2$. Not only are the increments of B_H definitely correlated, but they possess the essential property of global dependence. In particular, global dependence never dies out, since the same correlation is found between price changes over successive days, months, or years. The bulk of [5] is devoted to FBM.

This model's cartoon counterparts are those whose generators include only slanted intervals with $0 < H < 1$ and $H \neq 1/2$. A simulation of the M1965 model is shown in the second line of Figure 1. Its being Gaussian is a very serious defect.

The M1972 model of price variation, which accounts for both large deviations and long dependence. Finally, from a seed planted in 1972, exact renormalization and scaling led to fractional Brownian motion in multifractal time, which assumes that price is a Brownian motion of a "trading time" that is itself a multifractal function of clock time, namely, the integral of a multifractal measure.

Let us reflect on the relation between, on the one hand, the ideal market and the M1963 and M1965 models and, on the other hand, their multiplicity of acceptable cartoon models. The latter involve grids and a great deal of arbitrariness, and the former involve no grid and (basically) become uniquely determined when they are forced to be Gaussian. Can injecting the Gaussianity similarly eliminate arbitrariness from the multifractal cartoons, more precisely, from the relation of clock to trading time? A statistician may settle for the lognormal distribution, but (see [4] Chapter 14) I showed long ago that this wish is self-contradictory. The multifractal closest to the ideal of Gaussianity is a function for which $f(\alpha)$ is a cap convex parabola. It is technically described as "limit lognormal" and was the first to be investigated.

At this point, a clear-cut empirical question arises: could it be that the facts agree to the sub-ideal that comes just after the Brownian ideal, namely to fractional Brownian motion in limit lognormal trading time? This question was faced and most remarkably (see [1,6]) the answer is to the affirmative. The dollar-deutschmark rate of exchange (line DD of Figure 1) can be fitted using this limit lognormal multifractal measure. The remaining bottom lines on Figure 1 are simulations of the M1972 process fitted to the dollar-deutschmark data.

6 Miscellaneous Remarks.

Multiplicative growth models of scaling. Before statistical physics moved to problems where power law distributions are common, examples of these distributions concentrated in economics, beginning with the Pareto law for the distribution of personal income. Needless to say, a multitude of authors attempted to explain them but every model that actually yield a power-law distribution included a variant of the following key ingredient. Observation: when U satisfies $Pr\{U > u\} = u^{-\alpha}$, $\log U = V$ satisfies $Pr\{V > v\} = \exp(-\alpha v)$. Physics knows a plethora of different looking arguments that ultimately agree in yielding an exponential behavior for V. Hence the power law behavior for $Pr\{U > u\}$. A skeptic's opinion is described in Chapter 10 of [1].

Implications of multifractality. The good fit of MFAR raises an endless string of hard conceptual issues. Which is the cause of the regularity implicit in MFAR? Multifractals are found throughout physics, but to assume universal regularities in the economic fundamentals would be far-fetched. The next thought, that the complex interactions in financial markets end up by creating some, seems far more likely.

Enough mathematical properties of the multifractal model are already known and the statistical procedures are sufficiently developed, to allow the investigation of price records to proceed well beyond pictures, and showing that price series like DD are indeed multifractal. Without waiting for more mathematics, a first task is to use Monte Carlo calculations to help assess portfolio risks. The ideal market hypothesis has a beautiful mathematical theory, but predicts unrealistically low risks. My multifractal model is not the last word, but looks very promising.

A corporate treasurer, currency trader or other market strategist need not be told that prices do not vary continuously and oscillate wildly, that volatility is at the very heart of what goes on in financial markets. He would not dream of taking current calculus of risk at face value. The techniques I propose come closer, not to forecasting a price drop or rise on a specific day or time, but to estimating the probability of what the market might do and getting ready for it. In other words, I hope to have thrown a light of order in the seemingly impenetrable thicket of the stock market.

References

[1] B. B. Mandelbrot, *Fractals and Scaling in Finance: Discontinuity, Concentration, Risk* (*Selecta*, Volume E) (Springer-Verlag, New York, 1997).

[2] B. B. Mandelbrot, *Fractals, hasard et finance*, (*Selecta*, Volume FE) Flammarion (Collection *Champs*), Paris, 1997.

[3] B. B. Mandelbrot, *The Fractal Geometry of Nature* (W. H. Freeman and Company, New York, 1982).

[4] B. B. Mandelbrot, *Multifractals and 1/f Noise: Wild Self-Affinity in Physics*, (*Selecta*, Volume N) Springer-Verlag, New York, 1998.

[5] B. B. Mandelbrot, *Gaussian Fractals and Beyond* (*Selecta*, Volume H) (Springer-Verlag, New York, Expected in 1998 or 1999).

[6] B. B. Mandelbrot, L. Calvet & A. Fisher, *The Multifractal Model of Asset Returns. Large Deviations and the Distribution of Price Changes. The Multifractality of the Deutschmark/US Dollar Exchange Rate.* Three Discussion Papers. (Cowles Foundation for Economics, Yale University, New Haven, CT, September 1997.) Main Internet address: http://papers.ssrn.com/sol3/paper.taf?ABSTRACT. Additional data on the first paper: ABSTRACT_ID=78588. Second paper: ABSTRACT_ID=78608. Third paper: ABSTRACT_ID=78628

[7] B. B. Mandelbrot, Sur certains prix spéculatifs: faits empiriques et modèle basé sur les processus stables additifs de Paul Lévy. *Comptes Rendus Académie des Sciences* (Paris): **254** (1962) 3968–3970. Photographic reproduction in [2].

[8] B. B. Mandelbrot, The variation of certain speculative prices. *The Journal of Business of the University of Chicago* **36** (1963) 394–419. Reprinted in [1] and photographic reproduction in [10].

[9] B.B. Mandelbrot, Une classe de processus stochastiques homothétiques à soi: Application à la loi climatologique de H. E. Hurst. *Comptes Rendus Académie des Sciences* (Paris): **260** (1965) 3274-3277.

[10] P. H. Cootner(ed), *The Random Character of Stock Market Prices* (MIT Press, Cambridge, MA,1964).