

Science for Survival,
Pontificia Academia Scientiarum, MM

Fractal financial fluctuations; do they threaten sustainability?

Benoit B. Mandelbrot

"Were all the fountains of the great deep broken up, and the windows of heaven were opened. And the rain was upon the earth forty days and forty nights" Genesis: 6, 11-12

"There came seven years of great plenty throughout the land of Egypt. And there shall arise after them seven years of famine" Genesis: 41, 29-30.

The very diversity of the speakers at this meeting suffices to show that sustainability is threatened by very diverse obstacles. Several speakers described world-wide trends whose continuation would make sustainability impossible. My presentation is, instead, concerned with instances where the trends themselves might be tolerable but severe deviations from the norm are not, in fact, may well be unacceptable. For good reason, and as witnessed by the quote from the Bible that opens this text and will be commented upon momentarily, a deep concern with those deviations has an extremely long history. Until recently, no scientific tools were available to begin to face those deviations, not even to measure their intensity. This paper will not discuss all the tools now available, only those the speaker introduced very recently.

Based on the two quotes from the Bible, my terminology classifies the threatening deviations from norms in two categories. Because of the story of the Flood and Noah, I denote by the term "Noah Effect" those major changes that occur rapidly (even instantaneously) but have strong and durable consequences. Because of the story of the dream of Pharaoh and Joseph the son of Jacob, I denote by the term "Joseph Effect" those sequences of changes that need not seem threatening when viewed individually but become major in their cumulative effect.

The Noah and the Joseph obstacles to sustainability range from the wholly natural

(earthquakes, volcanoes, non man-made climate changes) to the wholly man-made. Both "effects" have long characterized the phenomena I studied very energetically. Many of those studies were reported or summarized in my 1982 monograph, or included in my *Selecta* books, of which the first three have been published. These are big books and it would be of little use to try to summarize them in these Proceedings. Instead, I shall focus on one topic and make minimal reference to the other aspects of the Noah and Joseph Effects.

That topic is the variation of financial prices, that is, of prices quoted on financial markets that trade securities, commodities, and exchange or interest rates. The wild volatility of those markets has long been known but only recently did it form an intimate link with such issues as sustainable growth of developing or other economically fragile countries.

The details of that intimate link belong to political economy. They will be briefly touched upon at the end of this text, but to keep in focus and manage space, an extensive discussion will be carried out elsewhere. The sole ambition of this paper is to further the knowledge of the underlying facts. Unfortunately and perhaps surprisingly, the existing models of financial price variability for good enough and new research is needed. Such statements tend to be viewed as self-serving - and an excuse for delaying urgent action. But in the present instance skepticism would be unwarranted. New research is keenly needed, because the points that matter most in the discussion of smooth sustainability coincide very precisely with points that previous research had deliberately set aside or disregarded. The text that follows is a draft of Chapter 1 of a forthcoming book of mine.

THIS BOOK TOUCHES MANY TOPICS, BUT ITS MAIN AMBITION is to contribute to a better understanding of price variation. Inevitably, it criticizes previously held views on this topic, particularly, the "coin tossing model," to be described momentarily.

The point of departure is that financial prices, including those of securities, commodities, foreign exchange or interest rates, are largely unpredictable. The best one can do is to evaluate the odds for or against some desired or feared outcomes, the most extreme being "ruin." Those odds will also be used as inputs for decisions

concerning economic policy or changes in institutional arrangements. To handle all those issues, the first step - but far from the last! - is to represent different instances of price variation by suitable random processes.

The word "suitable" and the plural in "processes" will surprise many readers. It is, indeed, widely believed that "random change" is a synonym for "prices that move up a bit or down a bit following the toss of a coin." The technical term is "simple random walk." It was made popular by a book title that won the high distinction of becoming a cliché., namely, "random walk down the Street."

The belief that there is no alternative is strengthened by the fact that coin tossing is indeed, the oldest and by far the most widely used model of price variation. The (unsaid) point of departure of this book and this Survey is that the term, random, has a far broader meaning, allowing the coin tossing model to be replaced by alternatives. Many of the alternatives are "unsuitable," but it will be argued that the alternative I put forward in this book, based on "multifractals," is very suitable indeed.

The multifractal model does not belong to esoteric mathematics and it must not be allowed to remain part of pure science. Its practical consequences are many and very serious. The first but not last is in the spirit of the Hippocratic Oath, "do no harm," which deserves to be generalized to finance and is best expressed in nautical term. When a ship was built to navigate placid lakes by fair weather, to send it across the ocean in typhoon season is to do serious harm. Similarly, the "coin tossing" model of financial prices may well be beloved by mathematicians, but it denies the existence of hurricanes; therefore it is dangerous.

The preceding nautical analogy will be heavily used throughout this text, because it resides at the very center of the present study. Not only alternatives to the coin tossing model are available, but the multifractal alternative differs in "qualitative" ways that have immediate consequences for finance and economic policies.

The coin tossing model exemplifies a form of randomness (a "state of randomness," as I shall argue) that can be called "mild." Had the evidence agreed with this model - but it does not at all - variability in finance would be as easily controllable as is variability in physics.

However, the coin tossing model must not be criticized too hard. It is always best to start with the simplest possible model and hold to it until it has begun to bring more harm than value. In its time, it played a fundamental and positive role in creating awareness of the difficulty of even the simplest forms of randomness. One can also argue that for the "man in the street" coin tossing is an adequate description of the facts. But policy makers and the professionals in finance are (or should be) far more demanding. It matters very much for them that, as will be seen, coin tossing is very far from accounting for some essential facts. Once again, the history of price variation is filled with "financial hurricanes" while we shall see that coin tossing claims that they practically never happen. Shipbuilders and ship owners cannot predict the dates and destructiveness of the hurricanes their vessel will encounter over its lifetime. But the knowledge that hurricanes will happen - and realistic evaluation of the corresponding odds - permeates ship-building, ship ownership and navigation.

This book argues that tools needed to acquire some mastery of the intensities of financial hurricanes are already available. They are those of fractal and multifractal geometry, a discipline better known as describing the shapes of coastlines and clouds and the distribution of galaxies and as having led to the discovery of the Mandelbrot set. My claim is that it also describes the growth and collapse of financial prices.

PICK THE FAKES

This book includes a multitude of words or numbers and also of formulas and dry diagrams. Without mastering them, my claims and contributions cannot be fully understood and appreciated. But to make the central point, as this Survey proposes, words, formulas and diagrams are not really necessary.

To explain my central point, the best and quickest way is to encourage the reader to participate in a test concerning Figures 1 and 2. No one is asked to accept pictures as the sole or final arbiter in scientific discourse, only as a useful additional tool. Pictures are often used to delude, but in this instance they deserve to be described as uncovering a widespread delusion and assisting in the selection of an improvement.

Among the many graphs in Figure 1, some are "real." They follow the practice of

financial journals and trace the sequence of daily closing of some price series such as security, commodity, foreign exchange or interest rate. The other graphs in Figure 1 will be called "forgeries." They correspond to imitations of financial reality provided by mathematical models that are fully specified in quantitative fashion, therefore can be sampled and illustrated without resorting to unreported stretching and reducing or other such manipulations.

Figure 2 plots price differences from one day to next.

Now the "find-the-fakes test" can be described: you are asked to separate reality and forgery as completely as you can. For a perfect score, you must rank the diagrams from "most obviously a forgery" to "apparently real."

When the test relates to Figure 1, it is very difficult to separate the real and forged records. Indeed, all such records tend to look alike. This




Figure 1 A collection of Diagrams, illustrating – in no particular order – the behavior in time of some actual financial prices and of some mathematical models of this behavior. It would be very difficult to pick the fakes.

impression is confirmed by looking through the financial press and the books on the mathematics of finance. The optimist will rush to conclude that coin tossing, which is represented by one of the graphs in Figure 1, is perfectly acceptable. In fact, Figure 1. A collection of diagrams, illustrating - in no particular order - the behavior in time of some actual financial prices and of some mathematical models of this behavior. It would be very difficult to pick the fakes.

we shall see momentarily that this optimism would be seriously misplaced. The resemblance between those curves is due to the fact that graphs of prices themselves do not reveal or enhance important differences, but hide them. In other words, plots of prices are a very inefficient way of presenting information. This is well-known to students of the psychology of vision: position is seen less accurately than change.

In sharp contrast, the lines in the stack of this Figure 2 show striking differences between




Figure 2. A stack of diagrams, illustrating the successive "daily" differences in some actual financial prices and some mathematical models. Pick the fakes!

one another. The meaning of those differences will be refined through this Survey, ending with the solution of the test.

LARGE STOCK MARKET MOVEMENTS AND THEIR ODDS

Individual investors and professional stock and currency traders know better than ever that prices quoted in any financial market often change with heart-stopping swiftness. Fortunes are made and lost in sudden bursts of activity when the market seems to speed up and the volatility soars. In September 1998, for instance, the stock for Alcatel, a French telecommunications equipment manufacturer, dropped about 40 percent one day and another 6 percent over the next few days.

Figure 3. Shapes of the Gaussian distribution and of two "stable" distributions. The latter provide a far better fit for many financial data, but the multifractal model is even more satisfactory.

In a reversal, the stock shot up 10 percent on the fourth day. On a longer time scale, most real price changes behave like those in the lower portion of Figure 2. However, not all lines at the bottom of Figure 2 are real. (That is, I am not giving away the test the reader is taking!)

The coin tossing model, which served as foundation for the theory of finance used most widely in this century, is represented by the top line of Figure 2 (now, I am giving away part of the test). We shall see in a moment that precipitous events like the Alcatel debacle are given totally negligible odds in that theory. Certainly, they should never happen in the lifetime of this generation and the next few. A cornerstone of finance is modern portfolio theory, which tries to maximize returns for a given level of risk. The mathematics underlying portfolio ignore the possibility of a typhoon?

This term, coin-tossing, is actually an oversimplification, but the risk-reducing formulas behind portfolio theory rely on a number of demanding premises that are mathematically attractive but rely on hope rather than reality.

First, they suggest that price changes are statistically independent of one another: for example, that today's price has no influence on the changes between the current price and tomorrow's. This is the "efficient market" hypothesis to ignore the possibility of a typhoon?

This term, coin-tossing, is actually an oversimplification, but the risk-reducing formulas

behind portfolio theory rely on a number of demanding premises that are mathematically attractive but rely on hope rather than reality.

First, they suggest that price changes are statistically independent of one another: for example, that today's price has no influence on the changes between the current price and tomorrow's. This is the "efficient market" hypothesis of Louis Bachelier. As a result, prediction of future market movements is never possible.

The second assumption is that all price changes are distributed in a pattern that conforms to the standard bell curve. Of the three curves in Figure 3, the bell curve is the flattest in the center. The width of its bell (as measured by its "sigma," or standard deviation) depicts how far price changes diverge from the mean. In this perspective, 95% of all cases fall into the narrow range between minus two sigmas and plus two sigmas. As was already mentioned and will be elaborated momentarily, the bell curve declares extreme events to be extremely rare. Typhoons are, in effect, defined out of existence.

Do financial data neatly conform to such assumptions? Of course, they never do! This is shown by a more attentive inspection of the bottom portion of Figure 2. It is true that charts of stock or currency changes over time reveal a constant background of small up and down price movements - though not as uniform as one would expect of price changes that fit the bell curve. Invariably, however, these patterns constitute only one aspect of the graph. A substantial number of sudden large changes - spikes on the chart that shoot up and down as with the Alcatel




Figure 3. Shapes of the Gaussian distribution and of two "stable" distributions. The latter provide a far better fit for many financial data, but the multifractal model is even more satisfactory.

stock - stand out from the background of more moderate perturbations. Moreover, it is typical of the magnitude of price movements - both large and small - to remain roughly constant for an extended period and then suddenly unpredictably increase for another extended period. Big price jumps become more common as the turbulence of the market grows -they cluster on the chart, expressing an obvious amount of dependence.

According to the coin tossing model, these large fluctuations often exceed ten sigmas, meaning ten standard deviations. This value is so huge that standard textbook tables of the Gaussian fail to include it. But a good calculator should yield their probability a few millionths of a millionth of a millionth, that is, approximately one day out of ten million million years. If risks were so tiny, they would be truly negligible, not worth even a passing thought. But this tiny value grossly contradicts the evidence. The real world of finance produces "ten sigma" spikes on a regular basis - as often as every month, to give an idea - and their probability should be expected to be a few hundredths.

The tiny probability mentioned a few lines above illustrates that the Gaussian practically vanishes near the left and right ends of the graph. Had the horizontal axis been part of Figure 3 (which it is not), it would have hidden those insignificant tails.

Reality is far better represented by the other two curves in Figure 3, both with more peaked heads and fatter tails. These curves belong to the "M 1963 model" produced by my first attack on financial data, in 1963. Having revealed this fact, it is best to narrow down the test the reader is supposed to be taking. Price changes according to the M 1963 model are the source of the second line in Figure 2. This is clearly better than the top line, to be sure, but far from being the last word.

COIN TOSSING NORMALITY VERSUS THE FINANCIAL REALITY

The bell curve is often described as illustrating the "normal" distribution. But does it follow that financial markets should be described as "abnormal or anomalous"? Of course not: they are what they are, and it is the coin tossing model, therefore portfolio theory, that are flawed.

Modern portfolio theory poses a danger to those who believe in it too strongly and is a

powerful challenge for the theoretician. Though sometimes acknowledging faults in the present body of thinking, the extreme bearish answer is that there is no alternative: large market swings are anomalies, individual "acts of God" that present no conceivable regularity. Other adherents suggest that coin tossing must be maintained, "faute de mieux," because no other premises can be handled through tough mathematical modeling leading to a rigorous quantitative description of at least some features of major financial upheavals.

An increasingly wide agreement is being




Figure 4. Constructing a "pseudo-Brownian cartoon" of the idealized coin-tossing model that underlies modern portfolio theory. The construction starts with a linear trend ("the initiator") and breaks it repeatedly by following a prescribed "generator." The interpolated generator is inverted for each descending piece. The pattern that emerges increasingly resembles market price oscillations.

reached, that the extreme bearish view is untenable and that the coin tossing model must be replaced by one that allows (near-) instantaneous price changes and substantial temporal dependence. This agreement marks a change of mood in the "mainstream," bringing it toward the views I have been campaigning for since 1963 and 1965, respectively. From this point on, however, two general approaches are in conflict, leading to what I shall call "micromanaged" and "macromanaged" models.

Micromanaged models agree with my diagnosis but not my follow-up. They proceed through a series of "fixes." Each fix "patches" a perceived defect of coin tossing, independently of its other defects. The outcome is that this approach accumulates a large number of parameters and no property is present that was not knowingly incorporated in the construction. In the nautical analog, the fixes consist in lengthening a small boat's keel, lengthening its mast, reinforcing its engine, etc..., one by one.

It is clearly preferable to design a large boat from scratch. Similarly, my experience of successful modeling in other fields has fostered deep a priori doubts about the chances of micromanaged modeling in finance. But personal prejudices would not have mattered if a posteriori modeling had been effective. I think it is not.

My own work - carried out over many years - takes a very different and decidedly bullish position. I claim that a financial model can be redesigned following an approach that is macromanaged by being guided by a principle of fractal invariance to which we shall come soon. The outcome, as I propose to show, is that the variation of financial prices can be accounted for by a model derived from my work in fractal geometry. Once again, fractals - or their later elaboration, called multifractals - do not claim the ability to predict the future with certainty. But they do create a more realistic picture of market risks. Given the recent troubles confronting the large investment pools called hedge funds, it would be foolhardy not to investigate models providing more accurate estimates of risk.

FRACTALS, MULTIFRACTALS AND THE MARKET

An extensive mathematical basis already exists for fractals and multifractals. Fractal patterns do not only appear in the price changes of securities but

also in the distribution of galaxies throughout the cosmos, in the shape of coastlines and in the decorative designs generated by innumerable computer programs.

A fractal is a geometric shape that can be separated into parts, each of which is a reduced-scale version of the whole. In finance, this concept is not a rootless abstraction but a theoretical reformulation of a down-to-earth bit of market folklore, namely, the notion that movements of a stock or currency all look alike when a market chart is enlarged or reduced so that it fits some prescribed time and price scales. This implies that an observer cannot tell which of the data concern prices that change from week to week, day to day, or hour to hour. This quality defines the charts as fractal curves and many powerful tools of mathematical and computer analysis become available.

A technical term for this form of close resemblance between the parts and the whole is self-affinity. This property is related to the better-known concept of fractals called self-similarity, in which every feature of a picture is reduced or blown up by the same ratio, a process familiar to anyone who ordered a photographic enlargement or a xerox copy. Financial market charts, however, are far from

being self-similar. If we simply focus on a detail of a graph, the features become increasingly higher than they are wide — as are the individual up-and-down price ticks of a stock. Hence, the transformation from the whole to the parts must shrink the time scale (the horizontal axis) more than the price scale (the vertical axis). This task can be performed by copiers using lasers. The geometric relation of the whole to its parts is said to be one of self-affinity.

Unchanging properties are not given much weight by most statisticians. But they are beloved of physicists and mathematicians like myself, who call them invariances and are happiest with models that present an attractive invariance property. A good idea of what I mean is provided by drawing a simple chart that inserts (interpolates) price changes from time 0 to a later time 1 in successive steps. The intervals themselves are chosen arbitrarily; they may represent a second, an hour, a day or a year.

The process begins with a price represented by a straight trend line called "initiator," shown in the top panel of Figure 4. Next, a broken line called "generator" is used to

create the pattern that corresponds to a slow up-and-down price oscillation. It is obviously essential for the number and positions of the pieces of the generator to be completely specified. If it is not, or if one allows oneself the right to fiddle with the generator during the construction, no prediction could be made. In Figure 4, the generator consists of three pieces that are inserted (interpolated) along the straight trend line. (A generator with fewer than three pieces could not simulate a price that must be able to move up and down.) Then, each of the generator's three pieces is interpolated by three shorter ones. Repeating these steps reproduces the shape of the generator, or price curve, but at increasingly compressed scales. Both the horizontal axis (time scale) and the vertical axis (price scale) are squeezed to fit the horizontal and vertical boundaries of each piece of the generator.

INTERPOLATIONS CONTINUED (NOT QUITE) FOREVER

Only four construction stages are shown in Figure 4, but the same process continues. In theory, it has no end, but in practice, it makes no sense to interpolate down to time intervals shorter than those between trading transactions, which may be of the order of a minute. The fact that each piece ends up with a shape like the whole is not a surprise: this scale invariance is present simply because it was built in. The novelty (and surprise) is that these very simply defined self-affine fractal curves exhibit a wealth of structure. The beauty of fractal geometry is that it does not consist in micromanaged models in which everything of interest has been inputted separately. Fractals involve only macromanaged instructions, yet make possible a model general enough to reproduce the patterns that characterize portfolio theory's placid markets as well as the tumultuous trading conditions of real markets.

Indeed, the construction's outcome, if plotted as in Figure 2, is very sensitive to the exact shape of the generator.

For example, Figure 4 uses a very special generator that - according to a theory I developed - will produce a behavior that is pseudo-Brownian, that is, close to the relatively tranquil, "mildly random," picture of the market ruled by coin tossing. But this level of tranquillity prevails only under extraordinarily special conditions that are satisfied only by equally special generators.

Figure 3 satisfies those conditions because each generator segment's height - namely, $2/3$, $1/3$ or $2/3$ - was made equal to the square root of the corresponding width - namely, $4/9$, $1/9$ or $4/9$. This "square root rule" is a characteristic of a process physicists call "simple diffusion" Adherence to the assumptions behind this oversimplified model is one of the central mistakes of modern portfolio theory. It is much like a theory of sea waves that forbids the swells to exceed six feet.

A first and very important generalization of Figure 4 yields models that are non Brownian but can be called "unifractal." It consists in continuing to require that the height of every segment of the generator be linked to its width by the same relation in the form of a power H . Previously, we set $H = 1/2$, but a different value of H can be chosen, as long as it is positive and less than 1. Taking $H = .7$ suffices to change the




Figure 5. This open cube illustrates related generators: The "right wall" shows an oscillating generator in trading time. This is the pseudo-Brownian (unifractal) generator of Figure? The "back wall" shows four multifractal oscillating generators in clock time. The "floor" shows the generators that relate the clock time to trading time. Each is an increasing function of the other. Moving a piece of the fractal generator to the left causes the same amount of market activity in a shorter time interval for the first piece of the generator and the same amount in a longer interval for the second piece.

top line of Figure 2 into its third line. On the corresponding graphs in Figure 1, the place of tranquility and mildness is taken by movements that are non-periodic but described by everyone as "cyclic," with many apparent cycle lengths, ranging from very small up to "about three cycles in a sample." (This last rule is a remarkable observation that cannot be elaborated here.) Here, cyclic behavior is present in the output without having been incorporated in the input. This is lovely, but large spikes of variation were lost and must be reinstated.

There is a second and far more drastic generalization of Figure 3. So far, market activity was assumed constant but one can allow it to speed up and slow down. This variability is the essence of volatility, in fact, practical people describe the diverse lines at the bottom of Figure 2 as proceeding at many different speeds at different times. This is why models that allow for variability add the prefix "multi" before the word "fractal."

To define "activity" is beyond our concern and not necessary. The key idea is that the market does not follow the physical time that precedes with the relentless regularity of a clock, but




Figure 6. The underlying pattern is as in Figure 5, but limited to the left-most generator, and the generators are replaced by the curve obtained by repeating it recursively as done in Figure 3. To make the picture clearer, the back and right wall are moved away from the floor.

instead a subjective time that flows slowly during some periods and fast during others.

In this spirit, the theory provides a neat "transmutation" from uni - to multifractal. The key step shown in Figure 5 is to lengthen or shorten the horizontal time axis so that the pieces of the generator are either stretched or squeezed. At the same time, the vertical price axis may remain untouched. As seen on the "back" wall of Figure 5, the first piece of the unifractal generator is progressively shortened, which also provides room to lengthen the second piece. After making these adjustments, the generators become multifractal (M1 to M4). As seen on the "floor," of Figure 6, market activity speeds up in the interval of time represented by the first piece of the generator and slows in the interval that corresponds to the second piece.

When the generators in Figure 5 are used recursively, one obtains the patterns shown in Figure 7. Recall that those patterns do not pretend to exhaust all the possibilities offered by either theory or the facts. Their sole aim is to show the power of the very simplest fractal models. Such an alteration to the generator can produce a full




Figure 7. Randomized multifractal "price increments" that correspond to the five multifractal generators in Figure 5. On top a pseudo-Brownian sequence of "price increments". A gradual displacement of the generator to the left causes market activity to increase gradually, becoming more and more volatile.

simulation of price fluctuations over a given period, using the process of interpolation described earlier. Each time, the first piece of the generator is further shortened. The process of successive interpolation produces a chart that increasingly resembles the characteristics of volatile markets (Figure 7).

Once again, the unifractal (U) chart that prevails before any shortening corresponds to the becalmed markets postulated in the portfolio theorists' model. Proceeding down the stack (M1 to M4), each chart diverges further from that model, exhibiting the sharp, spiky price jumps and the persistently large movements that characterize financial trading.

An important detail has not been mentioned yet. To make these models of volatile markets achieve the necessary realism, the three pieces of each generator were scrambled - a process not shown in the illustrations. It works as follows. Altogether, the three pieces of the generator allow the following six permutations:

1,2,3; 1,3,2; 2,1,3; 2,3,1; 3,1,2 and 3,2,1.

Conveniently enough, a die has six sides; imagine that each bears the image of one of the six permutations. Before each interpolation, the die is thrown and the permutation that comes up is selected.

BACK TO THE GAME OF PICKING THE FAKES

How do simulations of the multifractal model stand up against actual records of changes in financial prices? To respond, let us return to Figure 2, which is a composite of several historical series of price changes with a few outputs of artificial models.

As we have already observed the goal of modeling the patterns of real markets is certainly not fulfilled by the first chart, which is extremely monotonous and reduces to a static background of small price changes, analogous to the static noise from a radio. Volatility stays uniform with no sudden jumps. In a historical record of this kind, daily chapters would vary from one another, but all the monthly chapters would read very much alike.

The rather simple second chart is less unrealistic, because it shows many spikes; however, these are isolated against an unchanging

background in which the overall variability of prices remains constant. The third chart has interchanged strengths and failings, because it lacks any precipitous jumps.

The eye tells us that these three diagrams are unrealistically simple. Let us now recall the sources. Chart 1 illustrates price fluctuations in a model introduced in 1900 by French mathematician Louis Bachelier. The changes in prices follow a "random walk" that conforms to the bell curve and illustrates the model that underlies modern portfolio theory. Charts 2 and 3 are partial improvements on Bachelier's work: one is the "M 1963" model I proposed in 1963 (based on Levy stable random processes). The other is the "M 1965" model I published in 1965 (based on fractional Brownian motion). These revisions of coin tossing are inadequate, except under certain special market conditions.

By now, the test around which this Survey is structured has been reduced to a careful inspection of the more important five lower diagrams of the graph. Let me now add a piece of information that was withheld until now: at least one record is real and at least one is a computer-generated sample of my latest multifractal model. The reader is free to sort those five lines into the appropriate categories. I hope the forgeries will be perceived as surprisingly effective. In fact, only two are real graphs of market activity. Chart 5 refers to the changes in price of IBM stock and chart 6 shows price fluctuations for the dollar-deutsche mark exchange rate. The remaining charts (4, 7 and 8) bear a strong resemblance to their two real-world predecessors. But they are completely artificial.

Two technical points must be mentioned before moving on to conclusions. The recursive constructions in the body of the paper were nothing but "cartoons". The artificial charts 4, 7 and 8 were, instead, generated through a refined form of my multifractal model, called "fractional Brownian motion in multifractal trading time." Secondly, this introductory Survey necessarily emphasizes graphics, but - once again - the theory of multifractals is endowed with full numerical tools of analysis.

VERY TENTATIVE CONCLUSIONS: DIVERSIFICATION AND REINSURANCE

What conclusions should be drawn from all this? Does this matter to a corporate treasurer, currency

trader or other market strategists? Does this matter to the central banker and others concerned with overall financial and economic policy? Does this matter to the economist who seeks to explain the workings of the economy and concedes that his task may be helped by an accurate description of part of what is to be explained?

All those questions arise because the discrepancies between coin-tossing and the actual movement of prices have become too obvious to be ignored much longer. Prices do not vary continuously, and they are subjected to wild fluctuations at all time scales. Volatility – far from a static entity to be ignored or easily compensated for – is at the very heart of what goes on in financial markets. In the past, nearly everyone embraced the modern portfolio theory because of the absence of strong alternatives. But one need no longer accept it at face value.

However, the multifractal alternative is very young and very far from being fully developed. It deserves to draw attention (and criticism). By contrast, modern portfolio theory was formulated twenty-four years ago and was energetically developed ever since. Moreover, wild variability is a new notion endowed with little inherited capital. Modern portfolio theory inherited a large accumulated capital of techniques that statisticians designed to deal with mild Gaussian variability. The challenge was to adapt them to the context of financial prices.

Therefore, it is necessary, as we near a conclusion, to separate thoughts concerning the near future from thoughts concerning the longer range. Multifractals can immediately be put to work to "stress-test" portfolios, in particular, from the viewpoint of a quantity called "value at risk," whose definition is unfortunately beyond the scope of this text. stress-testing begins by questioning how a portfolio would have performed if it had been designed a while ago. That is, the simplest stress test merely uses historical data. But the actual market test will not come in the past, rather in the future, and a future that simply repeats the past is only one of many alternatives, and not a very likely one. The goal of every model of price variation (coin-tossing not being an exception) is to use the past to create the same patterns of variability as do the unknown rules that govern actual markets. This attempt should yield a collection of alternative scenarios for the future, and Stress-testing should include tests based on many such alternatives.

According to the coin-tossing model, the differences between those alternatives are comparatively slight. Not so with the multifractal models. They describe the past market fluctuations more realistically and the scenarios they propose for the future include a quota of extreme events that will really stress a portfolio. This is all that can be said on this subject at this point.

Of greatest interest, at least to me, are problems that the multifractal model confronts on broader institutional, temporal and spatial horizons. They are more important than any detail and question the worth of the widespread faith in the power of diversification and other forms of lumping and averaging. Here an enlightening analogy and powerful guidance for the future is provided by a distinction between different levels of insurance that relates to my distinction between mild and wild "states" of chance.

Most life, automobile, or homeowner risks are mild. Very much like the coin-tossing model of price change, they fall within a narrow range and are mutually independent. Even when a portfolio of insurance contracts is small, the wonders of diversification (due to the law of large numbers and related mathematics) can be trusted to create a risk of ruin that is sufficiently small to be profitable even for an insurance company with limited reserves. To play safe and to insure the occasional higher risk, the insurer of mild risks will seek reinsurance – which will seldom be needed, therefore will not be expensive. When a tornado defeats diversification of homeowner policies, the reinsurer is likely to be an entity that collected no premiums, namely an agency of a state.

However, many other risks seeking to be insured are wild, very much like in my multifractal model of price change. They involve the equivalents of the notorious "ten sigma" price changes that were discussed earlier in this Survey. Ordinary diversification would be defeated by such risks, even if the number of cases had sufficed for a law of large numbers to apply. More precisely, the odds of those wild risks, if included in the usual calculations, would imply reserves that are clearly unreasonable. However, such risks become insurable if they are immediately shared with reinsurers (or almost directly with competitors, as is apparently the case in the shipping industry).

The key fact is that insurers cannot survive by only considering the "fair weather" 95% of the claims, which would have easily been diversified. Not only the 5% of large "foul weather" claims cannot be ignored, but their odds are non-negligible and are an essential input of planned and carefully priced reinsurance.

Once again, theories based on coin-tossing legislate this "foul weather" out of existence, but it is evident that many features of the real world are best understood as designed to tackle comparatively rare but potentially disastrous situations. It is indeed filled with state or private institutions and informal or ad-hoc arrangements whose purpose can be viewed as that of reinsurance. A central part of my thinking in finance is that those arrangements may have worked in the past but cannot be relied upon in the future. As to institutions, their role deserves a fresh examination.

As a result of globalization, the relevance of the preceding comments on insurance is bound to increase. Under the coin-tossing model, the effects of globalization are limited. But the actual behavior of financial prices confirms what intuition suggests: the larger the markets, the greater the attention demanded by the potentially disastrous effects of financial storms.

To conclude no overall mathematical technique comes close to forecasting a price drop or rise on a specific day on the basis of past records. Multifractals do not do any better. But they provide estimates of the probability of what the market might do in the future and allow one to prepare for inevitable sea changes. The new modeling techniques are designed to cast a light of order into the seemingly impenetrable thicket of the financial markets. They also recognize the mariner's warning that, as recent events demonstrate, deserves to be heeded: On even the calmest sea, a gale may be just over the horizon.

BIOGRAPHY

Benoit Mandelbrot is unusual in having left a permanent mark on numerous fields of science and art. He is presently the Abraham Robinson Professor of Mathematical Sciences at Yale University. In 1958 he had joined IBM's Thomas J. Watson Research Center, where he served as IBM Fellow in Physics, 1974-1993. He belongs to the National Academy of Sciences of the USA, the American Academy of Arts and Sciences, and the

Norwegian Academy of Science and Letters. Having attended Ecole Polytechnique and earned the Ph.D. from the University of Paris, he also holds numerous honorary doctorates from institutions throughout the world. He received the 1993 Wolf Prize for Physics and also (among others) the Barnard, Franklin and Steinmetz Medals and the Harvey, Humboldt, Nevada, Honda and Arts and Science Prizes.

BIBLIOGRAPHY

B. MANDELBROT, *Fractals and Scaling in Finance: Discontinuity, Concentration, Risk*, Springer-Verlag, 1997.

B. MANDELBROT, *Multifractals and 1/f Noise: Wild Self-Affinity in Physics*, Springer-Verlag, 1999.

B. MANDELBROT, L. CALVET & A. FISHER, *The Multifractal Model of Asset Returns*, three reports of the Cowles Foundation for Economic Research of Yale University. Web address of the first: http://papers.ssrn.com/sol3/paper.taf?/ABSTRACT_ID=78588.

Addresses of the other two: same beginning and the endings `_ID=78608` and `_ID=78628`.