

Scale invariance and universality in economic phenomena

H E Stanley^{1,3}, L A N Amaral¹, P Gopikrishnan¹, V Plerou¹ and M A Salinger²

¹ Center for Polymer Studies and Department of Physics, Boston University, Boston, MA 02215, USA

² Department of Finance and Economics, School of Management, Boston University, Boston, MA 02215, USA

E-mail: hes@bu.edu.

Received 8 January 2002

Published 22 February 2002

Online at stacks.iop.org/JPhysCM/14/2121

Abstract

This paper discusses some of the *similarities* between work being done by economists and by computational physicists seeking to contribute to economics. We also mention some of the *differences* in the approaches taken and seek to justify these different approaches by developing the argument that by approaching the same problem from different points of view, new results might emerge. In particular, we review two such new results. Specifically, we discuss the two newly discovered scaling results that appear to be ‘universal’, in the sense that they hold for widely different economies as well as for different time periods: (i) the fluctuation of price changes of any stock market is characterized by a probability density function, which is a simple power law with exponent -4 extending over 10^2 standard deviations (a factor of 10^8 on the y -axis); this result is analogous to the Gutenberg–Richter power law describing the histogram of earthquakes of a given strength; (ii) for a wide range of economic organizations, the histogram that shows how size of organization is inversely correlated to fluctuations in size with an exponent ≈ 0.2 . Neither of these two new empirical laws has a firm theoretical foundation. We also discuss results that are reminiscent of phase transitions in spin systems, where the divergent behaviour of the response function at the critical point (zero magnetic field) leads to large fluctuations. We discuss a curious ‘symmetry breaking’ for values of Σ above a certain threshold value Σ_c ; here Σ is defined to be the local first moment of the probability distribution of demand Ω —the difference between the number of shares traded in buyer-initiated and seller-initiated trades. This feature is qualitatively identical to the behaviour of the probability density of the magnetization for fixed values of the inverse temperature.

³ Author to whom any correspondence should be addressed.

1. Introduction

One prevalent paradigm in economics is to marry finance with mathematics, with the fruit of this marriage the development of models. In physics, we also develop and make use of models or, as they are sometimes called, ‘artificial worlds’, but many physicists are fundamentally empirical in their approach to science—indeed, some physicists never make reference to models at all (other than in classroom teaching situations). This empirical approach has led to advances when theory has grown out of experiment. One such example is the understanding of phase transitions and critical phenomena [1]. Might this ‘empirics first’ physics paradigm influence the way physicists approach economics? Our group’s approach to economic questions has been to follow the paradigm of critical phenomena, which also studies complex systems comprised of many interacting subunits, i.e., to first examine the empirical facts as thoroughly as possible before we begin to construct models.

2. Fluctuations in finance

As do economists, physicists view the economy as a collection of interacting units. This collection is complex; everything depends on everything else. The interesting problem is: *how* does everything depend on everything else? Physicists are looking for robust empirical laws that will describe—and theories that will help one understand—this complex interaction [2–8].

To a physicist, the most interesting thing about an economic time series—e.g., the S&P 500 stock average index—is that it is dominated by fluctuations. If we make a curve of the values of the S&P 500 over a 35-year period, we see a fluctuating signal. Statistical physicists are particularly interested in fluctuating signals. The nature of this fluctuation immediately suggests to a physicist a model that was developed 100 years ago by Bachelier: the *biased random walk* [9].

A one-dimensional random walk is a drunk with a coin and a metronome. At each beat of the metronome, the coin is flipped—heads means one step to the right, tails one step to the left. If we look at our S&P 500 plot placed alongside a graph of a one-dimensional *biased* random walk—it is biased because it uses a ‘biased coin’ that has a slight tendency to go up rather than down—we physicists see a reasonable visual similarity. In fact, many economic pricing models—e.g., Black and Scholes—use this biased random walk.

Still there are certain points in the S&P 500 plot—such as 19 October 1987 (‘Black Monday’), or the 15% drop over the week following the events of 11 September 2001—that are not mirrored anywhere in the biased random-walk model. Nowhere do we see a drop anywhere near the 30% drop of Black Monday. This could not occur in a biased random walk—the probability that a walk will move two steps in the same direction is p^2 , three steps is p^3 , and so on—so the probability of many steps in the same direction is exponentially low, and it is virtually impossible.

Then how do we quantify these S&P 500 fluctuations? We begin by graphing the values of the fluctuations as a function of time. We place the plot of the empirical data next to the predictions of Bachelier’s model. The fluctuations in the model are normalized by one standard deviation. Note that the biased random walk has a probability density function (PDF) that is a Gaussian, so the probability of having more than five standard deviations is essentially zero—you can see that a line drawn at five standard deviations is outside the range of the fluctuations.

If we normalize the empirical data we see a difference. A line drawn at five standard deviations is *not* outside the range of the fluctuations—there are many ‘shocks’ that exceed five standard deviations. A bar placed on the positive side at five standard deviations also

has 30 or 40 hits—fluctuations that equal or exceed five standard deviations in the positive direction. Some, such as Black Monday, are more than 34 standard deviations. The exponential of $(-1/2)(34)^2$ is approximately $10^{-267/2}$.

Because big economic shocks affect the economy around the world ('everything depends on everything else'), the possibility of an economic 'earthquake' is one that we must take seriously. Big changes in stocks affect not only people with large amounts, but also those on the margins of society. One person's portfolio collapse is another's physical starvation; e.g., literal starvation in some areas was one result of the recent Indonesian currency collapse.

Another example is the recent Merriwether LTCM (long-term capital management) collapse, caused in part by the use of models that do not take into account these catastrophic rare events. Thus there are many reasons that we physicists might be interested in understanding economic fluctuations.

3. One possible conceptual framework

We shall see that our analysis of empirical data shows that those catastrophic rare events are a part of the overall picture: they are not simply inexplicable disasters beyond any possible understanding. Although this sounds as though we physicists think that we can contribute to economics, it is possible that the converse may be even more true. If we join economists in studying economics, we may stumble onto some ideas that will help us back in our more traditional research areas of physics. An example is turbulence. If one stirs a bucket of water, energy is added to the system on a big scale. This energy then dissipates over progressively smaller scales. This is an unsolved physics problem; many empirical facts can be stated, but little can be said about understanding it [10–12].

One could hypothesize that the economy is perhaps analogous to this example of turbulence. One can add *information* on a big scale to an economic system—e.g., the news of who wins a presidential election—and that information is dissipated on smaller and smaller scales. The way that you handle the 'turbulence' associated with this dissipation of information in a financial market may help us understand how to approach turbulence in our physics research. As attractive as this analogy might appear intuitively, it is not so accurate quantitatively, since the actual laws of turbulence are not at all the same as the empirical laws of economics [11, 12], despite early claims to the contrary [10].

Much of physics comes down to solving a differential equation, i.e., most functions in physics have some kind of characteristic scale. Once you have determined the scale, you can express the function in an exponential form—which has the wonderful property that the derivative of the function is also an exponential. In particular, the parameter r sets the scale of the problem. If r is positive, the function grows—and tells you the doubling time for the quantity of interest. Solutions to this look like Gaussians, and Gaussians are wonderful, tractable functions.

Such systems with scales describe almost everything in nature, including disordered things. Even raindrops on a sidewalk almost always have a characteristic scale. (If one 'zooms in' or 'zoom out', one can find the scale.) But there is a set of systems in nature that lack a scale. Systems in this set are described by power laws.

The framework for our approach to systems with many interacting subunits is something that is usually called 'scale invariance'. These systems vary greatly from systems that do have scales [1, 13].

We are all familiar with algebraic equations such as

$$x^2 = 4, \tag{1a}$$

and we know the solution is a number, ± 2 . Most of us are also familiar with functional equations, which are statements not about relations between numbers, but about the functional form of a function $f(x)$. Algebraic equations have solutions that are numbers, but functional equations have solutions that are functional forms. Power-law functions are the solutions of certain functional equations of the form

$$f(\lambda x) = \lambda^p f(x). \quad (1b)$$

In a functional equation, the converse also holds, i.e., every function that is of this power-law form also obeys this functional equation. This applies to a large number of contexts—in particular, to physical systems that have been tuned to be near critical points. An example is a binary mixture of two fluids in which the temperature has been tuned to be a special value called the critical temperature. At that temperature, there occur fluctuations in density in the binary mixture that extend over all length scales up to and including the wavelength of light. If you shine a flashlight on a tube of the binary mixture, you see an eerie glow—because the density fluctuations are so big in spatial extent, they become comparable to the wavelength of the light that is interacting with them. When that occurs, you see something that is visible—‘critical opalescence’. The same conceptual framework as describes this system appears to be able to describe economic systems [14].

4. Quantifying finance fluctuations

One topic we physicists are interested in is symmetry. An example of traditional symmetry is sodium chloride. One can displace the lattice an amount equal to exactly two lattice constants and the configuration will remain the same. One can rotate it 90° , or invert it, and the configuration will remain the same. Not only are these properties fascinating to mathematicians, they are also very relevant to solid-state physics. This simple symmetry and the mathematics and physics that are built on it have led to extremely useful inventions, e.g., the transistor.

The scale-invariance symmetry involved here is just as much a symmetry as the translational invariance symmetry in sodium chloride. We do not know how useful this scale-invariance symmetry will ultimately prove to be. Over the past 30 years physicists have used the theme of scale-invariance symmetry to understand systems near their critical points. Previous to this period of time, this class of problems was one nobody could solve: there were many, many length scales, not just one. The length scales could run from one nearest-neighbour spacing out to approximately 5000 (approximately the wavelength of light). The elements that make up this system are molecules that interact only over a short range—almost entirely with nearest neighbours. But this nearest-neighbour interaction propagates a small amount of torque through the system of nearest-neighbour interactions, so the entire system is somewhat affected.

This is beginning to sound like economics, in which ‘everything affects everything else’, and in economics, the first thing a physicist would do is look for the correlations. If we look at a graph of the autocorrelation function, we see a measure of the quantity G , which is a price change over some time horizon Δt . If we look at how G is now correlated with G at a time τ later, we measure that quantity as a function of τ , and as the size of τ increases, the correlation decreases. It is remarkable that this decrease happens in a regular fashion. How do we interpret this decrease? If we put the autocorrelation function in logarithmic units and the time lag in linear units, we see that the data fall on an approximate straight line. This means that the function is decaying exponentially, which means it does indeed have a characteristic scale [15–17]. So the autocorrelation function is not scale invariant. This differs from systems near their critical points in which the autocorrelation functions are scale invariant.

5. Statistical features of price fluctuations

The decay time in this economic example is short (4 min), so one cannot easily ‘make money’ on these correlations [15, 16]. A little less well known is the measure of the volatility [16, 17]. One way to quantify volatility is to replace G (the price change) with the absolute value of G . The data now are not at all linear on log–linear paper, but they are linear on log–log paper. And, of course, a power law $y = x^p$ is linear on log–log paper, because $\log y = p \log x$. The slope of the log–log plot p is the value of the exponent. These exponents turn out to be fundamental quantities. In this case, $p = -0.3$. The data are straight from about 200 min out to about 10^5 min—a range of almost 1000. With the data graphed, one can see the approximate region in which the data are straight—the data are not straight over all regions. Qualitatively, we have known for a long time that there are long-range correlations in the volatility, e.g., volatility ‘clustering’ and ‘persistence’, but this graph helps quantify this known empirical fact.

If we cannot find an ordinary correlation near a critical point, we must try something else. For example, we might simply dump all of our data ‘on the floor’. After we do that, the data no longer have time ordering nor do they have long- or short-range power-law correlations in the volatility of the autocorrelation function itself. Now we pick the data points up off the floor and make a histogram. Mandelbrot did this in 1963 with 1000 data points—a tiny number by today’s standards—for cotton-price fluctuations [14]. He concluded that those data were consistent with a Lévy distribution, i.e., a power-law distribution in that histogram—a so-called ‘fat tail’.

In 1995, Mantegna and Stanley decided to test this result using data with Δt shorter than for the daily data available in 1963 [15]. We used approximately one million data points: three orders of magnitude greater than Mandelbrot’s data set. Instead of Mandelbrot’s daily returns on cotton prices, we had returns approximately every 15 s on the S&P 500 index. We found that on a log–linear plot (i) the histogram of the G data points for the S&P 500 clearly is not a Bachelier/Black–Scholes Gaussian, and (ii) although the *centre* of the histogram agrees fairly well with Mandelbrot’s Lévy distribution, it begins to disagree after a few standard deviations. This disagreement led us to develop a class of mathematical processes called *truncated* Lévy distributions—which has attracted the attention of a number of mathematicians, who have carried this concept far further than we could [18–23].

What about ‘universality’, the notion in statistical physics that many laws seem to be remarkably independent of details? A physics example is that dramatically different materials behave exactly the same near their respective critical points. Binary mixtures, binary alloys, ferromagnets, even biological systems that involve switching, all behave the same way. An analogue of this universality appears in economics. For example, Skjeltorp [24] did a study that utilized the Mantegna approach. Instead of 1500 000 points from the S&P 500 (15 s intervals spread over six years), Skjeltorp did a parallel study for the Norwegian stock exchange and got almost exactly the same result.

We assumed that the reason we saw the truncated Lévy distribution while Mandelbrot did not was because we had more data—by three orders of magnitude. Gopikrishnan *et al* recently acquired a data set three orders of magnitude larger still (of order 10^9)—one that records *every* transaction for every stock. They found that when their data were graphed on log–log paper, the result was linearity [25–28]. This is the log of the cumulative distribution, the same quantity Mandelbrot plotted for cotton. But where Mandelbrot’s straight line had a slope of about 1.7 (well inside the Lévy regime, which stops at slope 2.0), Gopikrishnan’s straight line has a slope of ≈ 3.0 (far outside the limit for a Lévy distribution). The fact that these data are approximately linear over two orders of magnitude means that fluctuations that are as much as 100 standard deviations are still conforming to the same law as describes

the smaller fluctuations. This is reminiscent of the Gutenberg–Richter law that describes earthquakes [29–31]. Thus it would seem that these very rare events, which are conventionally treated as totally unexpected and unexplainable, have a precise probability describable by the same law as describes much more common events. These rare events occur with a frequency eight orders of magnitude less than the common, everyday event.

This means that Mandelbrot's results for cotton (10^3 points) are at total odds with Gopikrishnan's results for the S&P 500 (10^9 points). Why this difference? Is it simply because Mandelbrot did not have enough data to draw reliable conclusions? Or do commodities intrinsically have fatter tails? In recent work with data from British Petroleum, it appears that commodity data may have a slightly smaller slope—consistent with the possibility that perhaps there is not one universal behaviour for everything, but at least two separate universal behaviours—one for commodities and one for equities [32]. This smaller slope is still above 2, so the commodity data are not in the Lévy regime (even current data on cotton display a slope above 2).

6. Some similarities with 'diffusion in a tsunami wave'

Over this past year, we and our collaborators have been trying to understand these exponents using procedures similar to those used in critical phenomena, e.g., we relate one exponent to another and call the relation a scaling law, or we derive some microscopic model.

In particular, there appears to be an intriguing analogue not with the classic diffusion process studied in 1900 by Bachelier [9], but rather with a generalization called anomalous diffusion. It is plausible that classical diffusion does not describe all random motion. The Brownian motion seen in the behaviour of a grain of pollen in relatively calm water becomes something quite different if the grain of pollen is in a tsunami wave. The histograms would certainly be perturbed by a tsunami. A tsunami is an apt verbal metaphor for such economic 'earthquakes' as the Merriwether disaster, so why not explore the stock market as an example of anomalous diffusion?

In one-dimensional classic diffusion, a particle moves at constant velocity until it collides with something. One calculates, e.g., the end position of the particle, and (of course) finds a Gaussian. Within a fixed time interval Δt , one might calculate a histogram for the number of collisions $p(N)$, and also find a Gaussian. And if one calculates a histogram of the variance W^2 , one also finds a Gaussian. The fact that these are relatively narrow Gaussians means that there is a characteristic value, i.e., the width of that Gaussian, and that this is the basis for classical diffusion theory.

The quantity in the stock market corresponding to the displacement x is the price. At each transaction there is a probability that the price will change, and after a given time horizon there is a total change G . We have seen the histogram of G -values—the cumulative obeyed an inverse cubic law, and therefore the PDF, by differentiation, obeys an inverse quartic law.

What about these histograms? Apparently nobody had calculated these previously. Plerou *et al* set about using the same data analysed previously for G to calculate the histograms of N and W^2 . They also found power laws—not Gaussians, as in classic diffusion. That means there is no characteristic scale for the anomalous diffusion case (there is a characteristic scale for the classic diffusion case) and for an obvious reason. If you are diffusing around in a medium—such as the 'economic universe' that we live in—in which the medium itself is changing, then the laws of diffusion change and, in particular, they adopt this scale-free form. Further, the exponents that describe $p(N)$ and $p(W^2)$ appear [33, 34] to be the analogues of exponents in critical phenomena in the sense that they seem to be related to one another in interesting ways.

7. Some similarities with critical point phenomena

Stock prices respond to fluctuations in demand, just as the magnetization of an interacting spin system responds to fluctuations in the magnetic field. Periods with large number of market participants buying the stock imply mainly positive changes in price, analogous to a magnetic field causing spins in a magnet to align. Recently, Plerou *et al* [36] addressed the question of how stock prices respond to changes in demand. They quantified the relations between price change G over a time interval Δt and two different measures of demand fluctuations: (a) Φ , defined as the difference between the number of buyer-initiated and seller-initiated trades, and (b) Ω , defined as the difference in number of shares traded in buyer- and seller-initiated trades. They find that the conditional expectations $\langle G \rangle_\Phi$ and $\langle G \rangle_\Omega$ of price change for a given Φ or Ω are both concave. They find that large price fluctuations occur when demand is very small—a fact which is reminiscent of large fluctuations that occur at critical points in spin systems, where the divergent nature of the response function leads to large fluctuations. Their findings are reminiscent of phase transitions in spin systems, where the divergent behaviour of the response function at the critical point (zero magnetic field) leads to large fluctuations [1]. Further, Plerou *et al* [37] find a curious ‘symmetry breaking’ for values of Σ above a certain threshold value Σ_c ; here Σ is defined to be the local first moment of the probability distribution of demand Ω , the difference between the number of shares traded in buyer-initiated and seller-initiated trades. This feature is qualitatively identical to the behaviour of the probability density of the magnetization for fixed values of the inverse temperature.

8. Cross-correlations in price fluctuations of different stocks

We know that a stock price does not vary in isolation from other stock prices, but that stock prices are correlated. This fact is not surprising because we know that ‘in economics everything depends on everything else’. How do we *quantify* these cross-correlations of one stock with another? If we take the G -values of four companies out of the 1000 that we have studied—corresponding to the shrinking or growing of each of these four companies in, say, a 30 min interval. How does the behaviour of these four companies during that half-hour interval affect your response to their price activity? If two of the companies were Pepsi and Coke, there would probably be some correlation in their behaviours.

In order to quantify these cross-correlations, we begin by calculating a cross-correlation matrix. If we have 1000 firms, we have a 1000×1000 matrix \mathbf{C} , each element C_{ij} of which is the correlation of firm i with firm j . This large number of elements (one million) does not frighten a physicist with a computer. Eugene Wigner applied random-matrix theory 50 years ago to interpret the complex spectrum of energy levels in nuclear physics [35,38–47,49]. We do exactly the same thing, and apply random-matrix theory to the matrix \mathbf{C} . We find that certain eigenvalues of that 1000×1000 matrix deviate from the predictions of random-matrix theory, which has not eigenvalues greater than an upper bound of ≈ 2.0 . Furthermore, the contents of the eigenvectors corresponding to those eigenvalues correspond to well-defined business sectors. This allows us to define business sectors without knowing anything about the separate stocks: a Martian who cannot understand stock symbols could nonetheless identify business sectors [47,48].

9. Statistical physics and firm growth

In the economy, each firm depends on every other firm, and the interactions are not short ranged nor are they of uniform sign. For example, Ford Motor Company is in trouble because they have been selling their Explorer vehicle with extremely unsafe tyres—and the price of

their stock is going down. Prospective buyers purchase General Motors cars instead. There is thus a *negative* correlation between the stock prices of the two companies. But now General Motors needs to hire more workers to make a larger number of cars, and the McDonald's near the assembly plant has many more customers at lunch-time—a *positive*. Sometime later the situation may change again. So we can say that the 'spins' all interact with one another, and that these interactions change as a function of time.

Nevertheless, the general idea of a critical phenomenon seems to work. If the firms were spread out in a kind of chain, the correlations among them would decay exponentially. Instead, the firms interact with each other in much the same way as subunits in critical phenomena interact with each other. This fact motivated a study carried out about five years ago by a group of physicists interacting with an economist [51–53]. They calculated the fluctuations in business firms from one year to the next. They found that if they broke the fluctuations into bins by size, a tent-shaped distribution function was produced for each day of trading. The width of the tent was narrower for large firms than the width of the tent for small firms. This is not surprising, since a small firm has a potential to grow or shrink much more rapidly than a larger firm. When the widths of these tent-shaped distribution functions were plotted on log–log paper as a function of histogram size, the decreasing function turns out to be a straight line—corresponding to a power-law behaviour in that function. The exponent in that power law is ≈ 0.2 . The linearity extends over a number of decades, indicating that the data collapse onto a single plot irrespective of scale. This new result, which appears to be quite robust, has caught the attention of John Sutton [54], one of the leading economists at the London School of Economics, and of a number of other economists.

10. Universality in company growth

Takayasu and Okuyama [55] have demonstrated that the above results are universal by moving outside the realm of US economies and studying firm behaviour in other parts of the world.

Buldyrev *et al* [53] have shown that organizations (such as business firms) that are organized like trees will fluctuate in size. The hierarchical structure is set up such that instructions from the top of the hierarchy propagate down to the branching lower levels of the structure. Within that structure is a disobedience factor—those lower down do not always obey the directives handed down from those above them. This factor is, of course, crucial to the survival of the system. If employees always did only and exactly what they were told, any small mistake put into the system by a manager would grow and do an increasing amount of damage as it propagated through the expanding tree structure of the organization. On the other hand, the probability of an instruction being disobeyed cannot be one—or chaos would result. The propensity to disobey can be neither infinitesimal nor unity. The 'obeying probability' needs to settle at a point at which the organization can maintain both its integrity and self-corrective flexibility. And the behaviour of the exponent describing this probability is very similar to the behaviour of critical exponents.

This result is fairly robust, not only as far as business firm fluctuations are concerned, but also in the size of countries. Lee *et al* [56] extend the same analysis as was used for business firms to countries—and with the same exponent. Data can therefore be graphed on the same curve both for firms and for countries—where country size is measured by GDP.

We can see a similar pattern in the funding of university-based research. We researchers compete for research money the same way as business firms compete for customers. Plerou *et al* [57] analysed the funding of research groups over a 17-year period in the same way as fluctuations in firm size were analysed. The results were very similar, with the data collapsing onto the same curve.

As a final example, we mention the case of fluctuating bird populations in North America. In this case the exponent is 0.35 instead of 0.2. But, nevertheless, there seems to be some kind of property of organizations that we do not understand well [58].

11. 'Take-home message'

So—what have we learned? First, that the approach that we have emphasized is an empirical approach where one first seeks to uncover features of the complex economy that are challenges to understand. We find that there are two new universal scaling models in economics: (i) the fluctuation of price changes of any stock market is characterized by a PDF which is a simple power law with exponent $\alpha + 1 = 4$ that extends over 10^2 standard deviations (a factor of 10^8 on the y -axis); (ii) for a wide range of economic organizations, the histogram that shows how size of organization is inversely correlated with fluctuations in size with an exponent $\beta \approx 0.2$.

Neither of these two new laws has a firm theoretical foundation. This situation parallels the situation in the 1960s when the new field of critical phenomena also did not have a firm theoretical foundation for its new laws, but was awaiting the renormalization group. It is my hope that some of you will rise to the challenge and try to find a firm theoretical foundation for the structure of the empirical laws that appear to be describing (i) finance fluctuations, and (ii) economic organizations.

References

- [1] Stanley H E 1971 *Introduction to Phase Transitions and Critical Phenomena* (Oxford: Oxford University Press)
- [2] Mantegna R N and Stanley H E 1999 *An Introduction to Econophysics: Correlations and Complexity in Finance* (Cambridge: Cambridge University Press)
- [3] Bouchaud J P and Potters M 2000 *Theory of Financial Risk* (Cambridge: Cambridge University Press)
- [4] Levy H, Levy M and Solomon S 2000 *Microscopic Simulation of Financial Markets* (New York: Academic)
- [5] Roehner B M 2001 *Hidden Collective Factors in Speculative Trading* (Berlin: Springer)
- [6] Paul W and Baschnagel A 1999 *Stochastic Processes from Physics to Finance* (Berlin: Springer)
- [7] Voit J 2001 *The Statistical Mechanics of Financial Markets* (Berlin: Springer)
- [8] Takayasu H (ed) 2002 *Empirical Science of Financial Fluctuations: The Advent of Econophysics* (Berlin: Springer)
- [9] Bachelier L 1900 Théorie de la spéculation *PhD Thesis in Mathematics; Ann. Sci. Ecole Norm. Supér.* III **17** 21
- [10] Ghashgaie S, Breymann W, Peinke J, Talkner P and Dodge Y 1996 Turbulent cascades in foreign exchange markets *Nature* **381** 767–70
- [11] Mantegna R N and Stanley H E 1996 Turbulence and exchange markets *Nature* **383** 587–8
- [12] Mantegna R N and Stanley H E 1997 Stock market dynamics and turbulence: parallel analysis of fluctuation phenomena *Proc. Int. Conf. on Pattern Formation in Fluids and Materials; Physica A* **239** 255–66
- [13] Stanley H E 1999 Scaling, universality, and renormalization: three pillars of modern critical phenomena *Rev. Mod. Phys.* **71** S358–66
- [14] Mandelbrot B B 1963 The variation of certain speculative prices *J. Bus.* **36** 394–419
- [15] Mantegna R N and Stanley H E 1995 Scaling behaviour in the dynamics of an economic index *Nature* **376** 46–9
- [16] Liu Y, Gopikrishnan P, Cizeau P, Meyer M, Peng C-K and Stanley H E 1999 The statistical properties of the volatility of price fluctuations *Phys. Rev. E* **60** 1390–400
- [17] Ding Z, Granger C W J and Engle R F 1993 A long memory property of stock market returns and a new model *J. Empirical Finance* **1** 83–105
- [18] Mantegna R N and Stanley H E 1994 Stochastic process with ultraslow convergence to a Gaussian: the truncated Lévy flight *Phys. Rev. Lett.* **73** 2946–9
- [19] Podobnik B, Ivanov P Ch, Lee Y, Chessa A and Stanley H E 2000 Truncated Lévy flights in the presence of variance correlations *Europhys. Lett.* **50** 711–7
- [20] Mantegna R N and Stanley H E 1995 Ultra-slow convergence to a Gaussian: the truncated Lévy flight *Lévy Flights and Related Topics in Physics; Proc. 1994 Int. Conf. on Lévy Flights* ed M F Shlesinger, G M Zaslavsky, and U Frisch (Berlin: Springer) pp 300–12
- [21] Mantegna R N and Stanley H E 1998 Modeling of financial data: comparison of the truncated Lévy flight and the ARCH(1) and GARCH(1,1) processes *Proc. Int. IUPAP Conf. on Statistical Physics (Taipei); Physica A* **254** 77–84

- [22] Podobnik B, Ivanov P Ch, Lee Y and Stanley H E 2000 Scale-invariant truncated Lévy process *Europhys. Lett.* **52** 491–7
- [23] Ivanov P Ch, Podobnik B, Lee Y and Stanley H E 2001 Truncated Lévy process with scale-invariant behaviour *Proc. NATO Advanced Research Workshop on Application of Physics in Economic Modeling (Prague, Feb. 2001)*; *Physica A* **299** 154–60
- [24] Skjeltorp J A 2001 Scaling in the Norwegian stock market *Physica* **283** 486–528
- [25] Lux T 1996 The stable Paretian hypothesis and the frequency of large returns: an examination of major German stocks *Appl. Financ. Econ.* **6** 463–75
- [26] Gopikrishnan P, Meyer M, Amaral L A N and Stanley H E 1998 Inverse cubic law for the distribution of stock price variations *Eur. Phys. J. B* **3** 139–40
- [27] Plerou V, Gopikrishnan P, Amaral L A N, Meyer M and Stanley H E 1999 Scaling of the distribution of price fluctuations of individual companies *Phys. Rev. E* **60** 6519–29
- [28] Gopikrishnan P, Plerou V, Amaral L A N, Meyer M and Stanley H E 1999 Scaling of the distributions of fluctuations of financial market indices *Phys. Rev. E* **60** 5305–16
- [29] Gutenberg B and Richter C F 1954 *Seismicity of the Earth and Associated Phenomena* 2nd edn (Princeton, NJ: Princeton University Press)
- [30] Turcotte D L 1992 *Fractals and Chaos in Geology and Geophysics* (Cambridge: Cambridge University Press)
- [31] Rundle J B, Turcotte D L and Klein W 1996 *Reduction and Predictability of Natural Disasters* (Reading, MA: Addison-Wesley)
- [32] Matia K, Gopikrishnan P, Plerou V, Amaral L A N, Goodwin S and Stanley H E Scaling for the distribution of the commodity market variations 2002 *Phys. Rev. E* submitted (cond-mat/0202028)
- [33] Plerou V, Gopikrishnan P, Amaral L A N, Gabaix X and Stanley H E 2000 Diffusion and economic fluctuations *Phys. Rev. E* **62** 3023–6
- [34] Gopikrishnan P, Plerou V, Gabaix X and Stanley H E 2000 Statistical properties of share volume traded in financial markets *Phys. Rev. E* **62** 4493–6
- [35] Mehta M L 1991 *Random Matrices* (Boston, MA: Academic)
- [36] Plerou V, Gopikrishnan P, Gabaix X and Stanley H E 2001 Quantifying stock price response to demand fluctuations *Preprint* cond-mat/0106657
- [37] Plerou V, Gopikrishnan P and Stanley H E 2002 Critical threshold and symmetry breaking in stock demand *Preprint* cond-mat/0111349
- [38] Guhr T, Müller-Groeling A and Weidenmüller H A 1998 Random-matrix theories in quantum physics: common concepts *Phys. Rep.* **299** 189–425
- [39] E P Wigner 1951 On a class of analytic function from the quantum theory of collisions *Ann. Math.* **53** 36–65
- [40] E P Wigner 1951 On the statistical distribution of the widths and spacings of nuclear resonance levels *Proc. Camb. Phil. Soc.* **47** 790–8
- [41] Wigner E P 1956 Results and theory of resonance absorption *Conf. on Neutron Physics by Time-of-Flight (Gatlinburg, TN)* pp 59–70
- [42] Mehta M L and Dyson F J 1963 Statistical theory of the energy levels of complex systems V *J. Math. Phys.* **4** 713–19
- [43] Dyson F J 1971 The distribution of singular values of certain random matrices *Rev. Mex. Fís.* **20** 231
- [44] Sengupta A M and Mitra P P 1999 Distributions of singular values for some random matrices *Phys. Rev. E* **60** 3389–92
- [45] Laloux L, Cizeau P, Bouchaud J-P and Potters M 1999 Noise dressing of financial correlation matrices *Phys. Rev. Lett.* **83** 1469–82
- [46] Plerou V, Gopikrishnan P, Rosenow B, Amaral L A N and Stanley H E 1999 Universal and nonuniversal properties of financial cross-correlation matrices *Phys. Rev. Lett.* **83** 1471–5
- [47] Gopikrishnan P, Rosenow B, Plerou V and Stanley H E 2001 Quantifying and interpreting collective behaviour in financial markets *Phys. Rev. E* **64** 035106-1–035106-4
- [48] Rosenow B, Plerou V, Gopikrishnan P and Stanley H E 2002 Random matrix theory and portfolio optimization *Europhys. Lett.* submitted (cond-mat/0111537)
- [49] Plerou V, Gopikrishnan P, Rosenow B, Amaral L A N, Guhr T and Stanley H E 2002 A random matrix approach to financial cross-correlations *Phys. Rev. E* submitted (cond-mat/0108023)
- [50] Mantegna R N and Stanley H E 1997 Physics investigation of financial markets *Proc. Int. ‘Enrico Fermi’ School of Physics, Course CXXXIV* ed F Mallamace and H E Stanley (Amsterdam: IOS Press)
- [51] Stanley M H R, Amaral L A N, Buldyrev S V, Havlin S, Leschhorn H, Maass P, Salinger M A and Stanley H E 1996 Scaling behaviour in the growth of companies *Nature* **379** 804–6
- [52] Amaral L A N, Buldyrev S V, Havlin S, Leschhorn H, Maass P, Salinger M A, Stanley H E and Stanley M H R 1997 Scaling behaviour in economics: I. Empirical results for company growth *J. Physique I* **7** 621–33

-
- [53] Buldyrev S V, Amaral L A N, Havlin S, Leschhorn H, Maass P, Salinger M A, Stanley H E and Stanley M H R 1997 Scaling behaviour in economics: II. Modeling of company growth *J. Physique I* **7** 635–50
 - [54] Sutton J 2000 *The Variance of Firm Growth Rates: the 'Scaling' Puzzle* London School of Economics working paper
 - [55] Takayasu H and Okuyama K 1998 Country dependence on company size distributions and a numerical model based on competition and cooperation *Fractals* **6** 67–79
 - [56] Lee Y, Amaral L A N, Canning D, Meyer M and Stanley H E 1998 Universal features in the growth dynamics of complex organizations *Phys. Rev. Lett.* **81** 3275–8
 - [57] Plerou V, Amaral L A N, Gopikrishnan P, Meyer M and Stanley H E 1999 Similarities between the growth dynamics of university research and of competitive economic activities *Nature* **400** 433–7
 - [58] Keitt T and Stanley H E 1998 Scaling in the dynamics of North American breeding-bird populations *Nature* **393** 257–60