cond-mat/9810071 v2 19 Oct 1998

Crashes as Critical Points

Anders Johansen[1], Olivier Ledoit[2] and Didier Sornette[1,3,4]
[1] Institute of Geophysics and Planetary Physics
University of California, Los Angeles, California 90095
[2] Anderson Graduate School of Management
University of California, Box 90095-1481, 110 Westwood Plaza
Los Angeles CA 90095-1481
[3] Department of Earth and Space Science
University of California, Los Angeles, California 90095
[4] Laboratoire de Physique de la Matière Condensée
CNRS UMR6622 and Université de Nice-Sophia Antipolis
B.P. 71, Parc Valrose, 06108 Nice Cedex 2, France

October 19, 1998

Abstract

We study a rational expectation model of bubbles and crashes. The model has two components: (1) our key assumption is that a crash may be caused by *local* self-reinforcing imitation between noise traders. If the tendency for noise traders to imitate their nearest neighbors increases up to a certain point called the "critical" point, all noise traders may place the same order (sell) at the same time, thus causing a crash. The interplay between the progressive strengthening of imitation and the ubiquity of noise is characterized by the hazard rate, i.e. the probability per unit time that the crash will happen in the next instant if it has not happened yet. (2) Since the crash is not a certain deterministic outcome of the bubble, it remains rational for traders to remain invested provided they are compensated by a higher rate of growth of the bubble for taking the risk of a crash. Our model distinguishes between the end of the bubble and the time of the crash: the rational expectation constraint has the specific implication that the date of the crash must be random. The theoretical death of the bubble is not the time of the crash because the crash could happen at any time before, even though this is not very likely. The death of the bubble is the most probable time for the crash. There also exists a finite probability of attaining the end of the bubble without crash. Our model has specific predictions about the presence of certain critical log-periodic patterns in pre-crash prices, associated with the deterministic components of the bubble mechanism. We provide empirical evidence showing that these patterns were indeed present before the crashes of 1929, 1962 and 1987 on Wall Street and the 1997 crash on the Hong Kong Stock Exchange. These results are compared with statistical tests on synthetic data.

1 Introduction

Stock market crashes are momentous financial events that are fascinating to academics and practitioners alike. Within the efficient markets literature, only the revelation of a dramatic piece of information can cause a crash, yet in reality even the most thorough *post-mortem* analyses are typically inconclusive as to what this piece of information might have been. For traders, the fear of a crash is a perpetual source of stress, and the onset of the event itself always ruins the lives of some of them.

Our hypothesis is that stock market crashes are caused by the slow buildup of long-range correlations leading to a collapse of the stock market in one critical instant. The use of the word "critical" is not purely literary here: in mathematical terms, complex dynamical systems can go through so-called "critical" points, defined as the explosion to infinity of a normally well-behaved quantity. As a matter of fact, as far as nonlinear dynamical systems go, the existence of critical points is the rule rather than the exception. Given the puzzling and violent nature of stock market crashes, it is worth investigating whether there could possibly be a link between stock market crashes and critical points. The theory of critical phenomena has been recently applied to other economic models [Aoki, 1996].

On what one might refer to as the "microscopic" level, the stock market has characteristics with strong analogies to the field of statistical physics and complex systems. The individual trader has only 3 possible actions (or "states"): selling, buying or waiting. The transformation from one of these states to another is furthermore a discontinuous process due to a threshold being exceeded, usually the price of the stock. The transition involves another trader and the process is irreversible, since a trader cannot sell the stock he bought back to the same trader at the same price. In general, the individual traders only have information on the action of a limited number of other traders and generally only see the cooperative response of the market as a whole in terms of an increase or decrease in the value of the market. It is thus natural to think of the stock market in terms of complex systems with analogies to dynamically driven out-of-equilibrium systems such as earthquakes, avalanches, crack propagation *etc.* A non-negligible difference however involves the reflectivity mechanism : the "microscopic" building blocks, the traders, are conscious of their action. This has been masterly captured by Keynes under the parable of the beauty contest.

In this paper we develop a model based on the interplay between economic theory and its rational expectation postulate on one hand and statistical physics on the other hand. We find the following major points. First, it is entirely possible to build a dynamic model of the stock market exhibiting well-defined critical points that lies within the strict confines of rational expectations and is also intuitively appealing. Furthermore, the mathematical properties of such a critical point are largely independent of the specific model posited, much more so in fact than the "regular" (non-critical) behavior, therefore our key predictions should be relatively robust to model misspecification. In this spirit, we explore several versions of the model. Second, a statistical analysis of the stock market index in this century shows that the largest crashes may very well be outliers. This is supported by an extensive analysis of an artificial index. Third, the predictions made by our critical model are strongly borne out in several stock market crashes in this century, most notably the U.S. stock market crashes of 1929 and 1987 and the 1997 Hong-Kong stock market crash. Indeed, it is possible to identify clear signatures of near-critical behavior many years before these crash and use them to "predict" (out of sample) the date where the system will go critical, which happens to coincide very closely with the realized crash date. Again, we compare our results with an extensive analysis of an artificial index.

Lest this sound like voodoo science, let us reassure the reader that the ability to predict the critical date is perfectly consistent with the behavior of the rational agents in our model: they all know this date, the crash may happen anyway, and they are unable to make any abnormal risk-ajusted profits by using this information.

The first proposal for a connection between crashes and critical point was made by Sornette et al. (1996) which identified precursory patterns as well as characteristic oscillations of relaxation and aftershock signatures on the Oct. 1987 crash. These results were later confirmed independently (Feigenbaum and Freund, 1996) for the 1929 and 1987 crashes and it was in addition pointed out that the log-frequency of the observed log-periodic oscillations seems to decrease as the time of crash is approached. This was substantiated by the nonlinear renormalization group analysis of Sornette and Johansen (1997) on the 1929 and 1987 crashes. Other groups have also taken up the idea (Vandewalle et al., 1998a) and applied it to the analysis of the US stock market in 1997 (Feigenbaum and Freund, 1998; Gluzman and Yukalov, 1998; Laloux et al., 1998; Vandewalle et al., 1998b). These papers present some of the empirical results that we provide below. These papers were mostly empirical and did not present any detailled theoretical model and were based on qualitative arguments concerning the analogies between the stock market and critical phenomena in statistical physics discussed above. By contrast, we provide a rational economic model explaining how such patterns might arise. This has some fundamental implications, for example we see that the date of the crash must be random, therefore it needs not coincide with the date where the critical point is attained. This was not apparent in the purely "physical" approach taken by the earlier papers. Furthermore, the empirical study conducted in our paper is signicantly more detailed than the earlier ones and includes the most up-to date data as well as extensive back-testing to assess the extent to which our empirical evidence for stock market crashes as critical points is credible. Sornette and Johansen (1998) explore a different modelling route using a dynamical rupture analogy on a hierarchical network of trader connections, with the aim of identifying conditions for critical behavior to appear.

The next section presents and solves a rational model of bubbles and crashes. The third section argues that the mathematical properties so obtained are general signatures common to all critical phenomena. The fourth section discusses the empirical evidence for the model. The last section concludes.

2 Model

2.1 Price Dynamics

We consider a purely speculative asset that pays no dividends. For simplicity, we ignore the interest rate, risk aversion, information asymmetry, and the market-clearing condition. In this dramatically stylised framework, rational expectations are simply equivalent to the familiar martingale hypothesis:

$$\forall t' > t \qquad \mathrm{E}_t[p(t')] = p(t) \tag{1}$$

where $p(t)$ denotes the price of the asset at time t and $\mathrm{E}_t[\cdot]$ denotes the expectation conditional on information revealed up to time t. Since we will uphold Equation (1) throughout the paper, we place ourselves firmly on the side of the efficient markets hypothesis (Fama; 1970, 1991). If we do not allow the asset price to fluctuate under the impact of noise, the solution to Equation (1) is a constant: $p(t) = p(t_0)$, where t_0 denotes some initial time. Since the asset pays no dividend, its fundamental value is $p(t) = 0$. Thus a positive value of $p(t)$ constitutes a speculative bubble. More generally, $p(t)$ can be interpreted as the price in excess of the fundamental value of the asset.

In conventional economics, markets are assumed to be efficient if all available information is reflected in current market prices. This constitutes a first-order approximation, from which to gauge and quantify any observed departures from market efficiency. This approach is based on the theory of rational expectation, which was an answer to the previous dominent theory, the so-called Adaptive Expectation theory. According to the latter, expectations of future inflation, for example, were forecast as being an average of past inflation rates. This effectively meant that changes in expectations would come slowly with the changing data. In contrast, rational expectation theory argues that the aggregate effect of people's forecasts of any variable future behavior is more than just a clever extrapolation but is close to an optimal forecast from using all available data. This is true notwithstanding the fact that individual people forecasts may be relatively limited : the near optimal result of the aggregation process reflects a subtle cooperative adaptative organization of the market (Potters et al., 1998). Recently, the assumption of rational expectations has been called into question by many economists. The idea of heterogeneous expectations has become of increasing interest to specialists. Shiller (1989), for example, argues that most participants in the stock market are not "smart investors" (following the rational expectation model) but rather follow trends and fashions. Here, we take the view that the forecast of a given trader may be sub-optimal but the aggregate effect of all forecasts on the market price leads to its rational expected determination. The no-arbitrage condition resulting from the rational expectation theory is more than a useful idealization, it describes a self-adaptive dynamical state of the market resulting from the incessant actions of the traders (arbitragers). It is not the out-of-fashion equilibrium approximation sometimes described but rather embodies a very subtle cooperative organization of the market. Our distinction between individual sub-optimal traders and near-optimal market is at the basis of our model constructed in two steps.

Now let us introduce an exogenous probability of crash. Formally, let j denote a jump process whose value is zero before the crash and one afterwards. The cumulative distribution function (cdf) of the time of the crash is called $Q(t)$, the probability density function (pdf) is $q(t) = dQ/dt$ and the hazard rate is $h(t) = q(t)/[1 - Q(t)]$. The hazard rate is the probability per unit of time that the crash will happen in the next instant if it has not happened yet. Assume for simplicity that, in case of a crash, the price drops by a fixed percentage $\kappa \in (0, 1)$. Then the dynamics of the asset price before the crash are given by:

$$dp = \mu(t)\, p(t)\, dt - \kappa p(t) dj \tag{2}$$

where the time-dependent drift $\mu(t)$ is chosen so that the price process satisfies the martingale condition, i.e. $\mathrm{E}_t[dp] = \mu(t)p(t)dt - \kappa p(t)h(t)dt = 0$. This yields: $\mu(t) = \kappa h(t)$. Plugging it into Equation (2), we obtain a ordinary differential equation whose solution is:

$$\log \left[\frac{p(t)}{p(t_0)}\right] = \kappa \int_{t_0}^{t} h(t')dt' \qquad \text{before the crash.} \tag{3}$$

The higher the probability of a crash, the faster the price must increase (conditional on having no crash) in order to satisfy the martingale condition. Intuitively, investors must be compensated by the chance of a higher return in order to be induced to hold an asset that might crash. This is the only effect that we wish to capture in this part of the model. This effect is fairly standard, and it was pointed out earlier in a closely related model of bubbles and crashes under rational expectations by Blanchard (1979, top of p.389). It may go against the naive preconception that price is adversely affected by the probability of the crash, but our result is the only one consistent with rational expectations. Notice that price is driven by the hazard rate of crash $h(t)$, which is so far totally unrestricted.

A few additional points deserve careful attention. First, the crash is modelled as an exogenous event: nobody knows exactly when it could happen, which is why the rational traders cannot earn abnormal risk-ajusted profits by anticipating it. Second, the probability of a crash itself is an exogenous variable that must come from outside this first part of the model. There may or may not be feedback loop whereby prices would in turn affect either the arrival or the probability of a crash. This may not sound totally satisfactory, but it is hard to see how else to obtain crashes in a rational expectations model: if rational agents could somehow trigger the arrival of a crash they would choose never to do so, and if they could control the probability of a crash they would always choose it to be zero. In our model, the crash is a random event whose probability is driven by external forces, and *once this probability is given* it is rationally reflected into prices.

2.2 The Crash

The goal of this section is to explain the macro-level probability of a crash in terms of micro-level agent behavior, so that we can derive specific implications pinning down the hazard rate of crash $h(t)$.

We start by a discussion in naive terms. A crash happens when a large group of agents place sell orders simultaneously. This group of agents must create enough of an imbalance in the order book for market makers to be unable to absorb the other side without lowering prices substantially. One curious fact is that the agents in this group typically do not know each other. They did not convene a meeting and decide to provoke a crash. Nor do they take orders from a leader. In fact, most of the time, these agents disagree with one another, and submit roughly as many buy orders as sell orders (these are all the times when a crash *does not* happen). The key question is: by what mechanism did they suddenly manage to organise a coordinated sell-off?

We propose the following answer: all the traders in the world are organised into a network (of family, friends, colleagues, etc) and they influence each other *locally* through this network. Specifically, if I am directly connected with k nearest neighbors, then there are only two forces that influence my opinion: (a) the opinions of these k people; and (b) an idiosyncratic signal that I alone receive. Our working assumption here is that agents tend to *imitate* the opinions of their nearest neighbors, not contradict them. It is easy to see that force (a) will tend to create order, while force (b) will tend to create disorder. The main story that we are telling in this paper is the fight between order and disorder. As far as asset prices are concerned, a crash happens when order wins (everybody has the same opinion: selling), and normal times are when disorder wins (buyers and sellers disagree with each other and roughly balance each other out). We must stress that this is exactly the opposite of the popular characterisation of crashes as times of chaos.

Our answer has the advantage that it does not require an overarching coordination mechanism: we will show that macro-level coordination can arise from micro-level imitation. Furthermore, it relies on a somewhat realistic model of how agents form opinions. It also makes it easier to accept that crashes can happen for no rational reason. If selling was a decision that everybody reached independently from one another just by reading the newspaper, either we would be able to identify unequivocally the triggering news after the fact (and for the crashes of 1929 and 1987 this was not the case), or we would have to assume that everybody becomes irrational in exactly the same way at exactly the same time (which is distasteful). By contrast, our reductionist model puts the blame for the crash simply on the tendency for agents to imitate their nearest neighbors. We do not ask why agents are influenced by their neighbors within a network: since it is a well-documented fact (see e.g. Boissevain and Mitchell, 1973), we take it as a primitive assumption rather than as the conclusion of some other model of behavior.[1] Note, however, that there is no information in our model, therefore what determines the state of an agent is pure noise (see Black, 1986, for a more general discussion of noise traders).

We first present a simple "mean field" approach to the imitation problem and then turn to a more microscopic description.

2.2.1 Macroscopic modelling

In the spirit of "mean field" theory of collective systems (see for instance (Goldenfeld, 1992)), the simplest way to describe an imitation process is to assume that the hazard rate

[1]Presumably some justification for these imitative tendencies can be found in evolutionary psychology. See Cosmides and Tooby (1994) on the relationship between evolutionary psychology and Economics.

$h(t)$ evolves according to the following equation:

$$\frac{dh}{dt} = C\ h^{\delta}\ , \qquad \text{with } \delta > 1\ , \tag{4}$$

where C is a positive constant. Mean field theory amounts to embody the diversity of trader actions by a single effective representative behavior determined from an average interaction between the traders. In this sense, $h(t)$ is the collective result of the interactions between trader. The term h^{δ} in the r.h.s. of (4) models that fact that the hazard rate will increase or decrease due to the presence of *interactions* between the traders. The exponent $\delta > 1$ quantifies the effective number equal to $\delta - 1$ of interactions felt by a typical trader. The condition $\delta > 1$ is crucial to model interactions and is, as we now show, essential to obtain a singularity (critical point) in finite time. Indeed, integrating (4), we get

$$h(t) = (\frac{h_0}{t_c - t})^{\alpha}\ , \qquad \text{with } \alpha \equiv \frac{1}{\delta - 1}\ . \tag{5}$$

The critical time t_c is determined by the initial conditions at some origin of time. The exponent α must lie between zero and one for an economic reason: otherwise, the price would go to infinity when approaching t_c (if the bubble has not crashed in the mean time), as seen from (3). This condition translates into $2 < \delta < +\infty$: for this theory to make sense, this means that a typical trader must be connected to more than one other trader.

It is possible to incorporate a feedback loop whereby prices affect the probability of a crash. The higher the price, the higher the hasard rate or the increase rate of the crash probability. This process reflects the phenomenon of a self-fulfilling crisis, a concept that has recently gained wide attention, in particular with respect to the crises that occurred during the past four years in seven countries – Mexico, Argentina, Thailand, South Korea, Indonesia, Malaysia, and Hong Kong (Krugman, 1998). They all have experienced severe economic recessions, worse than anything the United States had seen since the thirties. It is believed that this is due to the feedback process associated with the gain and loss of confidence from market investors. Playing the "confidence game" forced these countries into macroeconomic policies that exacerbated slumps instead of relieving them (Krugman, 1998). For instance, when the Asian crisis struck, countries were told to raise interest rates, not cut them, in order to persuade some foreign investors to keep their money in place and thereby limit the exchange-rate plunge. In effect, countries were told to forget about macroeconomic policy; instead of trying to prevent or even alleviate the looming slumps in their economies, they were told to follow policies that would actually deepen those slumps, all this for fear of speculators. Thus, it is possible that a loss of confidence in a country can produce an economic crisis that justifies that loss of confidence– countries may be vulnerable to what economists call "self-fulfilling speculative attacks." If investors believe that a crisis may occur in absence of certain actions, they are surely right, because they themselves will generate that crisis. Van Norden and Schaller (1994) have proposed a Markov regime switching model of speculative behavior whose key feature is similar to ours, namely that the overvaluation over the fundamental price increases the probability and expected size of a stock market crash.

Mathematically, in the spirit of the mean field approach, the simplest way to model this effect is to assume

$$\frac{dh}{dt} = D\ p^{\mu}\ , \qquad \text{with } \mu > 0\ . \tag{6}$$

D is a positive constant. This equation, together with (3), captures the self-fulfilling feature of speculators: their lack of confidence, quantified by $h(t)$, deteriorates as the market price departs increasingly from the fundamental value. As a consequence, the price has to increase further to remunerate the investors for their increasing risk. The fact that it is the rate of change of the hazard rate which is a function of the price may account for panic psychology, namely that the more one lives in a risky situation, the more one becomes accutely sensitive to this danger until one becomes prone to exageration.

Putting (6) in (3) gives

$$\frac{d^2x}{dt^2} = \kappa D\ e^{\mu x}\ , \tag{7}$$

where $x \equiv \ln p$. Its solution is, for large x,

$$x = \frac{2}{\mu}\ \ln\left(\frac{\sqrt{\mu\kappa/2}}{t_c - t}\right)\ . \tag{8}$$

The log-price exhibits a weak logarithmic singularity of the type proposed by Vandewalle et al. (1998). This solution (8) is the formal analytic continuation of the general solution (14) below for $\beta \to 0$.

2.2.2 Microscopic modelling

Consider a network of agents: each one is indexed by an integer $i = 1, \ldots, I$, and $N(i)$ denotes the set of the agents who are directly connected to agent i according to some graph. For simplicity, we assume that agent i can be in only one of two possible states: $s_i \in \{-1, +1\}$. We could interpret these states as "buy" and "sell", "bullish" and "bearish", "calm" and "nervous", etc, but we prefer to keep the discussion about imitation at a general level for now. We postulate that the state of trader i is determined by:

$$s_i = \text{sign}\left(K \sum_{j \in N(i)} s_j + \sigma \varepsilon_i\right) \tag{9}$$

where the sign$(\cdot)$ function is equal to $+1$ (to -1) for positive (negative) numbers, K is a positive constant, and ε_i is independently distributed according to the standard normal distribution. This equation belongs to the class of stochastic dynamical models of interacting particles (Liggett, 1985, 1997), which have been much studied mathematically in the context of physics and biology.

In this model (9), the tendency towards imitation is governed by K, which is called the coupling strength; the tendency towards idiosyncratic behavior is governed by σ. Thus the value of K relative to σ determines the outcome of the battle between order and disorder, and eventually the probability of a crash. More generally, the coupling strength K could be heterogeneous across pairs of neighbors, and it would not substantially affect the properties of the model. Some of the K_{ij}'s could even be negative, as long as the average of all K_{ij}'s was strictly positive.

Note that Equation (9) only describes the state of an agent at a given point in time. In the next instant, new ε_i's are drawn, new influences propagate themselves to neighbors, and agents can change states. Thus, the best we can do is give a statistical description of the states. Many quantities can be of interest. In our view, the one that best describes the chance that a large group of agent finds itself suddenly in agreement is called the *susceptibility* of the system. To define it formally, assume that a global influence term G is added to Equation (9):

$$s_i = \text{sign}\left(K \sum_{j \in N(i)} s_j + \sigma\varepsilon_i + G\right). \tag{10}$$

This global influence term will tend to favour state $+1$ (state -1) if $G > 0$ (if $G < 0$). Equation (9) simply corresponds to the special case $G = 0$: no global influence. Define the average state as $M = (1/I)\sum_{i=1}^{I} s_i$. In the absence of global influence, it is easy to show by symmetry that $\mathrm{E}[M] = 0$: agents are evenly split between the two states. In the presence of a positive (negative) global influence, agents in the positive (negative) state will outnumber the others: $\mathrm{E}[M] \times G \geq 0$. With this notation, the susceptibility of the system is defined as:

$$\chi = \left.\frac{d(E[M])}{dG}\right|_{G=0} \tag{11}$$

In words, the susceptibility measures the sensitivity of the average state to a small global influence. The susceptibility has a second interpretation as (a constant times) the variance of the average state M around its expectation of zero caused by the random idiosyncratic shocks ε_i. Another related interpretation is that, if you consider two agents and you force the first one to be in a certain state, the impact that your intervention will have on the second agent will be proportional to χ. For these reasons, we believe that the susceptibility correctly measures the ability of the system of agents to agree on an opinion. If we interpret the two states in a manner relevant to asset pricing, it is precisely the emergence of this global synchronisation from local imitation that can cause a crash. Thus, we will characterise the behavior of the susceptibility, and we will posit that the hazard rate of crash follows a similar process. We do not want to assume a one-to-one mapping between hazard rate and susceptibility because there are many other quantities that provide a measure of the degree of coordination of the overall system, such as the correlation length (i.e. the distance at which imitation propagates) and the other moments of the fluctuations of the average opinion. As we will show in the next section, all these quantities have the same generic behavior.

2.3 Interaction Networks

It turns out that, in the imitation model defined by Equation (9), the structure of the network affects the susceptibility. We propose two alternative network structures for which the behavior of susceptibility is well understood. In the next section, we will show how the results we get for these particular networks are in fact common to a much larger class of models.

2.3.1 Two-Dimensional Grid

As the simplest possible network, let us assume that agents are placed on a two-dimensional grid in the Euclidean plane. Each agent has four nearest neighbors: one to the North, one to the South, the East and the West. The relevant parameter is K/σ. It measures the tendency towards imitation relative to the tendency towards idiosyncratic behavior. In the context of the alignment of atomic spins to create magnetisation, this model is related to the so-called two-dimensional Ising model which has been solved explicitly by Onsager (1944). There exists a critical point K_c that determines the properties of the system. When $K < K_c$, disorder reigns: the sensitivity to a small global influence is small, the clusters of agents who are in agreement remain of small size, and imitation only propagates between close neighbors. Formally, in this case, the susceptibility χ of the system is finite. When K increases and gets close to K_c, order starts to appear: the system becomes extremely sensitive to a small global perturbation, agents who agree with each other form large clusters, and imitation propagates over long distances. In the Natural Sciences, these are the characteristics of so-called *critical* phenomena. Formally, in this case the susceptibility χ of the system goes to infinity. The hallmark of criticality is the *power law*, and indeed the susceptibility goes to infinity according to a power law:

$$\chi \approx A(K_c - K)^{-\gamma}. \tag{12}$$

where A is a positive constant and $\gamma > 0$ is called the *critical exponent* of the susceptibility (equal to 7/4 for the 2-d Ising model).

We do not know the dynamics that drive the key parameter of the system K. At this stage of the enquiry, we would like to just assume that it evolves smoothly, so that we can use a first-order Taylor expansion around the critical point. K need not even be deterministic: it could very well be a stochastic process, as long as it moves slowly enough. Let us call t_c the first time such that $K(t_c) = K_c$. Then prior to the critical date t_c we have the approximation: $K_c - K(t) \approx \text{constant} \times (t_c - t)$. Using this approximation, we posit that the hazard rate of crash behaves in the same way as the susceptibility (and all the other measures of coordination between noise traders) in the neighborhood of the critical point. This yields the following expression:

$$h(t) \approx B \times (t_c - t)^{-\alpha} \tag{13}$$

where B is a positive constant. The exponent α must lie between zero and one for an economic reason: otherwise, the price would go to infinity when approaching t_c (if the bubble has not crashed yet). The probability per unit of time of having a crash in the next instant conditional on not having had a crash yet becomes unbounded near the critical date t_c.

We stress that t_c is not *the* time of the crash because the crash could happen at any time before t_c, even though this is not very likely. t_c is the mode of the distribution of the time of the crash, i.e. the most probable value. In addition, it is easy to check that there exists a residual probability $1-Q(t_c)>0$ of attaining the critical date without crash. This residual probability is crucial for the rational expectations hypothesis, because otherwise the whole model would unravel because rational agents would anticipate the crash. Other than saying that there is some chance of getting there, our model does not describe what happens at and after the critical time. Finally, plugging Equation (13) into Equation (3) gives the following law for price:

$$\log[p(t)] \approx \log[p_c] - \frac{\kappa B}{\beta} \times (t_c - t)^{\beta} \qquad \text{before the crash.} \tag{14}$$

where $\beta = 1-\alpha \in (0,1)$ and p_c is the price at the critical time (conditioned on no crash having been triggered). We see that the logarithm of the price before the crash also follows a power law. It has a finite upper bound $\log[p_c]$. The slope of the logarithm of price, which is the expected return per unit of time, becomes unbounded as we approach the critical date. This is to compensate for an unbounded probability of crash in the next instant.

2.3.2 Hierarchical Diamond Lattice

The stock market constitute an ensemble of inter-actors which differs in size by many orders of magnitudes ranging from individuals to gigantic professional investors, such as pension funds. Furthermore, structures at even higher levels, such as currency influence spheres (US$, DM, YEN ...), exist and with the current globalization and de-regulation of the market one may argue that structures on the largest possible scale, *i.e.*, the world economy, are beginning to form. This means that the structure of the financial markets have features, which resembles that of hierarchical systems and with "traders" on all levels of the market. Of course, this does not imply that any strict hierarchical structure of the stock market exists, but there are numerous examples of qualitatively hierarchical structures in society. In fact, one may say that horizontal organisations of individuals are rather rare. This means that the plane network used in the previous section may very well represent a gross over-simplification.

Another network structure for which our local imitation model has been solved is the following. Start with a pair of traders who are linked to each other. Replace this link by a diamond where the two original traders occupy two diametrically opposed vertices, and where the two other vertices are occupied by two new traders. This diamond contains four links. For each one of these four links, replace it by a diamond in exactly the same way, and iterate the operation. The result is a diamond lattice (see Figure 1). After p iterations, we have and $\frac{2}{3}(2+4^p)$ traders and 4^p links between them. Most traders have only two neighbors, a few traders (the original ones) have 2^p neighbors, and the others are in between. Note that the least-connected agents have 2^{p-1} times fewer neighbors than the most-connected ones, who themselves have approximately 2^p fewer neighbors than there are agents in total. This may be a more realistic model of the complicated network of communications between financial agents than the grid in the Euclidean plane of Section 2.3.1.

A version of this model was solved by Derrida et al. (1983). The basic properties are similar to the ones described above: there exists a critical point K_c; for $K < K_c$ the susceptibility is finite; and it goes to infinity as K increases towards K_c. The only difference – but it is an important one – is that the critical exponent can be a complex number. The general solution for the susceptibility is a sum of terms like the one in Equation (12) with complex exponents. The first order expansion of the general solution is:

$$\begin{aligned} \chi &\approx \mathsf{Re}[A_0(K_c - K)^{-\gamma} + A_1(K_c - K)^{-\gamma+i\omega} + \ldots] && (15) \\ &\approx A_0'(K_c - K)^{\gamma} + A_1'(K_c - K)^{\gamma}\cos[\omega \log(K_c - K) + \psi] + \ldots && (16) \end{aligned}$$

where A_0', A_1', ω and ψ are real numbers, and $\mathsf{Re}[\cdot]$ denotes the real part of a complex number. We see that the power law is now corrected by oscillations whose frequency explodes as we reach the critical time. These accelerating oscillations are called "log-periodic", and $\frac{\omega}{2\pi}$ is called their "log-frequency". There are many physical phenomena where they decorate the power law (see Sornette (1998) for a review). Following the same steps as in Section 2.3.1, we can back up the hazard rate of a crash:

$$h(t) \approx B_0(t_c - t)^{-\alpha} + B_1(t_c - t)^{-\alpha}\cos[\omega \log(t_c - t) + \psi']. \tag{17}$$

Once again, the hazard rate of crash explodes near the critical date, except that now it displays log-periodic oscillations. Finally, the evolution of the price before the crash and before the critical date is given by:

$$\log[p(t)] \approx \log[p_c] - \frac{\kappa}{\beta}\left\{B_0(t_c - t)^{\beta} + B_1(t_c - t)^{\beta}\cos[\omega \log(t_c - t) + \phi]\right\} \tag{18}$$

where ϕ is another phase constant. The key feature is that oscillations appear in the price of the asset just before the critical date. The local maxima of the function are separated by time intervals that tend to zero at the critical date, and do so in geometric progression, i.e. the ratio of consecutive time intervals is a constant

$$\lambda \equiv e^{\frac{2\pi}{\omega}} . \tag{19}$$

This is very useful from an empirical point of view because such oscillations are much more strikingly visible in actual data than a simple power law: a fit can "lock in" on the oscillations which contain information about the critical date t_c. If they are present, they can be used to predict the critical time t_c simply by extrapolating frequency acceleration. Since the probability of the crash is highest near the critical time, this can be an interesting forecasting exercise. Note that, for rational traders in our model, such forecasting is useless because they already know the hazard rate of crash $h(t)$ at every point in time (including at t_c), and they have already reflected this information in prices through Equation (3).

3 Generalisation

Even though the predictions of the previous section are quite detailed, we will try to argue in this section that they are very robust to model misspecification. We claim that models of crash that combine the following features:

1. A system of noise traders who are influenced by their neighbors;
2. Local imitation propagating spontaneously into global cooperation;
3. Global cooperation among noise traders causing crash;
4. Prices related to the properties of this system;
5. System parameters evolving slowly through time;

would display the same characteristics as ours, namely prices following a power law in the neighborhood of some critical date, with either a real or complex critical exponent. What all models in this class would have in common is that the crash is most likely when the locally imitative system goes through a *critical* point.

In Physics, critical points are widely considered to be one of the most interesting properties of complex systems. A system goes critical when local influences propagate over long distances and the average state of the system becomes exquisitely sensitive to a small perturbation, *i.e.* different parts of the system becomes highly correlated. Another characteristic is that critical systems are self-similar across scales: in our example, at the critical point, an ocean of traders who are mostly bearish may have within it several islands of traders who are mostly bullish, each of which in turns surrounds lakes of bearish traders with islets of bullish traders; the progression continues all the way down to the smallest possible scale: a single trader (Wilson, 1979). *Intuitively speaking, critical self-similarity is why local imitation cascades through the scales into global coordination.*

Because of scale invariance (Dubrulle et al., 1997), the behavior of a system near its critical point must be represented by a power law (with real or complex critical exponent): it is the only family of functions that are homogeneous, i.e. they remain unchanged (up to scalar multiplication) when their argument gets rescaled by a constant. Mathematically, this means that scale invariance and critical behavior is intimately associated with the following equation:

$$F(\gamma x) = \delta F(x) \tag{20}$$

where F is a function of interest (in our example, the susceptibility), x is an appropriate parameter, and δ is a positive constant that describes how the properties of the system change when we re-scale the whole system by the factor γ. It is easy to verify that the general solution to Equation (20) is:

$$F(x) = x^{\log(\delta)/\log(\gamma)}\, \pi\left[\frac{\log(x)}{\log(\gamma)}\right] \tag{21}$$

where π is a periodic function of period one. Equation (16) is nothing but the terms of order 0 and 1 in the Fourier expansion of the periodic function π.

In general, physicists study critical points by forming equations such as (20) to describe the behavior of the system across different scales, and by analysing the mathematical properties of these equations. This is known as *renormalisation group theory* (Wilson, 1979), and its introduction was the major breakthrough in the understanding of critical points ("renormalisation" refers to the process of rescaling, and "group" refers to the fact that iterating Equation (21) generates a similar equation with the rescaling factor γ^2).[2] Before renormalisation group theory was invented, the fact that a system's critical behavior had to be correctly described at all scales simultaneously prevented standard approximation methods from giving satisfactory results. But renormalisation group theory turned this liability into an asset by building its solution precisely on the self-similarity of the system across scales. Let us add that, in spite of its conceptual elegance, this method is nonetheless quite challenging mathematically.

For our purposes, however, it is sufficient to keep in mind that the key idea of the paper is the following: the massive and unpredictable sell-off occuring during stock market crashes comes from local imitation cascading through the scales into global cooperation when a complex system approaches its critical point. Regardless of the particular way in which this idea is implemented, it will generate the same universal implications, which are the ones in Equations (14) and (18).

Strictly speaking, these equations are approximations valid only in the neighborhood of the critical point. Sornette and Johansen (1997) propose a more general formula with additional degrees of freedom to better capture behavior away from the critical point. The specific way in which these degrees of freedom are introduced is based on a finer analysis of the renormalisation group theory that is equivalent to including the next term in a systematic expansion around the critical point. It is given by:

$$\log\left[\frac{p_c}{p(t)}\right] \approx \frac{(t_c-t)^\beta}{\sqrt{1+\left(\frac{t_c-t}{\Delta_t}\right)^{2\beta}}}\left\{B_0 + B_1\cos\left[\omega\log(t_c-t) + \frac{\Delta_\omega}{2\beta}\log\left(1+\left(\frac{t_c-t}{\Delta_t}\right)^{2\beta}\right)+\phi\right]\right\} \tag{22}$$

where we have introduced two new parameters: Δ_t and Δ_ω. Two new effects are included in this equation: (i) far from the critical point the power law tapers off; and (ii) the log-frequency shifts from $\frac{\omega+\Delta_\omega}{2\pi}$ to $\frac{\omega}{2\pi}$ to as we approach the critical point. Both transitions take place over the same time interval of length Δ_t. Thus, in addition to the critical regime defined by Equation (18), we allow for a pre-critical regime where prices oscillate around a constant level with a different log-periodicity (see Sornette and Johansen, 1997, for the formal derivation). This generalisation describe prices over a longer pre-crash period than the original formula, but it still captures exactly the same underlying phenomenon: critical self-organisation. Note that for small $\frac{t_c-t}{\Delta_t}$ Equation (22) boils down to Equation (18).

[2]Rigorously speaking, this only constitutes a *semi*-group because the inverse operation is not defined.

Before we continue with the empirical results obtained from analysis of stock market data, we would like to stress that replacing equation (18) with equation (22) *does not* correspond to simply increasing the number of free variables in the function describing the time evolution of the stock market index. Instead, we are including the *next order term* in the analytical expansion of the solution (21) to equation (20), which is something quite different. As we shall see in section 4.2, this has some rather restrictive implications.

4 Empirical Results

4.1 Large Crashes are Outliers

In figure 2, we see the distribution of draw downs (continuous decreases) in the closing value of the Dow Jones Average larger than 1% in the period 1900-94. The distribution resembles very much that of an exponential distribution while three events are standing out. The derivation of the exponential distribution of draw downs is given in appendix A. If we fit the distribution of draw downs DD larger than 1% by an exponential law, we find

$$N(\mathrm{DD}) = N_0 \; e^{-|\mathrm{DD}|/\mathrm{DD}_c}, \tag{23}$$

where the total number of draw downs is $N_0 = 2789$ and $DD_c \approx 1.8$ %, see figure 2.

The important point here is the presence of these three events that should not have occurred at this high rate. This provides an empirical clue that large draw downs and thus crashes might result from a different mechanism, that we attribute to the collective destabilizing imitation process described above. Ranked, the three largest crashes are the crash of 1987, the crash following the outbreak of World War I[3] and the crash of 1929.

To quantify how much the three events deviates from equation (23), we can calculate what accordingly would be the typical return time of a draw down of an amplitude equal to or larger than the third largest crash of 23.6%. Equation (23) gives the number of draw downs equal to or larger than DD and predicts the number of drawn downs equal to or larger than 23.6% per century to be ≈ 0.0056. The typical return time of a draw down equal to or larger than 23.6% would then be the number of centuries n such that $0.0056 \cdot n \sim 1$, which yields $n \sim 180$ centuries. In contrast, Wall Street has sustained 3 such events in less than a century. If we take the exponential distribution of draw downs larger than 1% suggested by figure 2 at face value, it suggest that the "normal day-to-day behavior" of the stock market index is to a large extent governed by random processes but the largest crashes are signatures of cooperative behavior with long-term build-up of correlations between traders.

As an additional test of this hypothesis, we have applied a more sophisticated null-hypothesis than that of an exponential distribution of draw downs and used a GARCH(1,1)

[3]Why the out-break of World War I should affect the Wall Street much more than the Japanese bombing of Pearl Harbour or the out-break of the Korean War seems very odd. World War I was largely an internal European affair and in the beginning something the populations of the warring countries were very enthusiastic about.

model estimated from the true Dow Jones Average, see appendix B. The background for this model is the following (Bollerslev et al., 1992). First, the stock market index itself cannot follow a simple random (Brownian) walk, since prices cannot become negative. In fact, one of the basic notions of investment is that the gain should be proportional to the sum invested. Using the logarithm of the stock market index instead of the index itself takes care of both objections and allows us to keep the random walk assumption for day-to-day variations in the stock return. However, the logarithm of stock returns is not normally distributed, but follows a fat-tailed distribution, such as the Student-t. Furthermore, the variance around some average level should change as a function of time in a correlated way, since large price movements tend to be followed by large price movements and small price movements tend to follow small price movements.

¿From this GARCH(1,1) model, we have generated 10.000 independent data sets corresponding in total to approximately *one million* years of "garch-trading" with a reset every century. Among these 10.000 surrogate data sets only two had 3 draw downs above 22% and none had 4[4]. However, each of these data sets contained crashes with a rather abnormal behavior in the sense that they were preceded by a draw up of comparable size as the draw down. This means that in a million years of garch-trading *not once* did 3 asymmetric crashes occur in a single century.

We note that the often symmetric behavior of artificial indices, *i.e.*, that large/small price movements in either direction are followed by large/small price movements in either directions is not quite compatible with what is seen in the Dow Jones Average [Brock et al., 1992]. In general, we see that the build-up of index is relatively slow and crashes quite rapid. In other words, "bubbles" are slow and crashes are fast. This means that not only must a long-term model of the stock market prior to large crashes be highly non-linear but it must also have a "time-direction" in the sense that the crash is "attractive" prior to the crash and "repulsive" after the crash. This breaking of time-symmetry, in other word the fact that the time series does not look statistically similar when viewed under reversal of the time arrow, has been demonstrated by Arnéodo et al. (1998). Using wavelets to decompose the volatility (standard deviation) of intraday (S&P500) return data across scales, Arnéodo et al. (1998) showed that the two-point correlation functions of the volatility logarithms across different time scales reveal the existence of an asymmetric causal information cascade from large scales to fine scales and *not* the reverse. Cont (1998) has recently confirmed this finding by using a three-point correlation function suggested by Pomeau (1982), constructed in such a way that it becomes non-zero under time asymmetry.

The analysis presented in this section suggests that it is *very* likely that the largest crashes of this century have a different origin than the smaller draw downs. We have suggested that these large crashes may be viewed as critical points and the purpose of the next section is to quantify the characteristic signatures associated with such critical points, specifically that of a critical divergence and log-periodic oscillations.

[4]The first attempt of "only" 1000 data sets did not contain any "GARCH(1,1) centuries" with 3 crashes larger than 22%.

4.2 Fitting Stock Market Indices

The simple power law in Equation (14) is very difficult to distinguish from a non-critical exponential growth over a few years when data are noisy. Furthermore, systematic deviations from a monotonous rise in the index is clearly visible to the naked eye. This is why all our empirical efforts were focused on the log-periodic formulas.

Fitting the stock market with a complex formula like equation (18) and especially equation (22) involves a number of considerations, the most obvious being to secure that the best possible fit is obtained. Fitting a function to some data is nothing but a minimisation algorithm of some cost-function of the data and the fitting-function. However, with noisy data and a fitting-function with a large number of degrees of freedom, many local minima of the cost-function exist where the minimisation algorithm can get trapped. Specifically, the method used here was a downhill simplex minimisation (Press *et al.* , 1992) of the variance. In order to reduce the number of free parameters of the fit, the 3 linear variables have been "slaved". This was done by requiring that the cost-function has zero derivative with respect to A, B, C in a minimum. If we rewrite equations (18) and (22) as $\log(p(t)) \approx A + Bf(t) + Cg(t)$ then we get 3 linear equations in A, B, C

$$\sum_i^N \begin{pmatrix} \log(p(t_i)) \\ \log(p(t_i))\, f(t_i) \\ \log(p(t_i))\, g(t_i) \end{pmatrix} = \sum_i^N \begin{pmatrix} N & f(t_i) & g(t_i) \\ f(t_i) & f(t_i)^2 & f(t_i)\, g(t_i) \\ g(t_i) & f(t_i)\, g(t_i) & g(t_i)^2 \end{pmatrix} \cdot \begin{pmatrix} A \\ B \\ C \end{pmatrix} \tag{24}$$

to solve for each choice of values for the non-linear parameters. Equations (24) was solved using the LU decomposition algorithm in (Press *et al.* , 1992) thus expressing A, B and C as functions of the remaining non-linear variables. Of these, the phase ϕ is just a time unit and has no physical meaning[5]. By changing the time unit of the data (from f.ex. days to months or years) it can be shown that β and ω as well as the timing of the crash represented by t_c is independent of ϕ. This leaves 3 (equation (18)) respectively 5 (equation (22)) physical parameters controlling the fit, which at least in the last case, is quite a lot. Hence, the fitting was preceded by a so-called Taboo-search (Cvijović D.and J. Klinowski J. 1995) where only β and ϕ was fitted for fixed values of the other parameters. All such scans, which converged with $0 < \beta < 1$ was then re-fitted with all non-linear parameters free. The rationale behind this restriction is, as explained in section 2.3.1, economical since $\beta < 0$ would imply that the stock market index could go to infinity.

In the case of equation (18), the described procedure always produced one or a few distinct minima. Unfortunately, in the case of equation (22) things were never that simple and quite different sets of parameter values could produce more or less the same value of the variance. This degeneracy of the variance as a function of $\alpha, t_c, \omega, \Delta_t$ and Δ_ω means that the error of the fit alone is not enough to decide whether a fit is "good or bad". Naturally, a large variance means that the fit is bad, but the converse is not necessarily true. If we return to the derivation of equations (18) and (22), we note three things. First, that α is determining the strength of the singularity. Second, that the frequency

[5] In fact, A, B, C can be regarded as simple units as well.

$\omega \approx 2\pi/\ln\lambda$ is determined by the underlying hierarchical structure quantified by the allowed re-scaling factors γ^n, as previously defined. Last, we recall that Δ_t in equation 22 describes a transition between two regimes of the dynamics, i.e., far from respectively close to the singularity. These considerations have the following consequences. If we believe that large crashes can be described as critical points and hence have the same physical background, then β, ω (actually γ) and Δ_t should have values which are comparable. Furthermore, Δ_t should not be much smaller or much larger than the time-interval fitted, since it describes a transition time. Realizing this, we can use the values of β, ω and τ together with the error of the fit to discriminate between good and bad fits. This provides us with much stronger discriminating statistics than just the error of the fit.

4.3 Estimation of Equations (18)

For the period before the crash of 1987, the optimal fit of Equation (18) is shown in Figure 3, along with two trend lines representing an exponential and a pure power law. The log-periodic oscillations are so strong that they are visible even to the naked eye. The estimation procedure yields a critical exponent of $\beta = 0.57$. The position of the critical time t_c is within a few days of the actual crash date. This result was reported earlier by Sornette et al. (1996) and Feigenbaum and Freund (1996). The latter authors also show a similar fit for the crash of 1929. See also Leland (1988) and Rubinstein (1988) for discussions of the crash of 1987 along different lines.

The turmoil on the financial US market in Oct. 1997 fits into the framework presented above. Detection of log-periodic structures and prediction of the stock market turmoil at the end of october 1997 has been made, based on an unpublished extension of the theory. This prediction has been formally issued ex-ante on september 17, 1997, to the French office for the protection of proprietary softwares and inventions under number registration 94781. Vandewalle et al. have also issued an alarm, published in the Belgian newspaper Tendances, 18. September 1997, page 26, by H. Dupuis, entitled "Un krach avant novembre" (a krach before november). It turned out that the crash did not really occur : only the largest dayly loss since Oct. 1987 was observed with ensuing large volatilities. This has been argued by Laloux et al. (1998) as implying the "death" of the theory. In fact, this is fully consistent with our rational expectation model of a crash : the bubble expands, the market believes that the crash is more and more probable, the prices develop characteristic structures but the critical time passes out without much happening. This corresponds to the finite probability $1 - Q(t_c)$, mentionned above, that no crash occurs over the whole time including the critical time of the end of the bubble. Feigenbaum and Freund have also analyzed, after the fact, the log-periodic oscillations in the S&P 500 and the NYSE in relation to the October 27'th "correction" seen on Wall Street.

In figure 4 we see the Hang Seng index fitted with Equation (18) approximately 2.5 years prior to the recent crash in October 1997. Rather remarkably the values obtained for the exponent α and the frequency ω differs less 5 % from the values reported for the 1987 crash on Wall Street by Sornette *et al.* (1996). Note also the similar behavior of the two indices over the last few months prior to the crashes.

4.4 Estimation of Equation (22)

The hope here is that we will be able to capture the evolution of the stock price over a longer pre-crash period, say 7-8 years. For estimation purposes, Equation (18) was rewritten as:

$$\log\left[p(t)\right] = A + B\frac{(t_c - t)^{\beta}}{\sqrt{1+\left(\frac{t_c-t}{\Delta_t}\right)^{2\beta}}}\left\{1 + C\cos\left[\omega\log(t_c - t) + \frac{\Delta_\omega}{2\beta}\log\left(1+\left(\frac{t_c-t}{\Delta_t}\right)^{2\beta}\right)\right]\right\} \tag{25}$$

We use the same least-squares method as above. We again concentrate away the linear variables A, B and C in order to form an objective function depending only on t_c, β, ω and ϕ as before, and the two additional parameters Δ_t and Δ_ω. We can not put any bounds on the second frequency Δ_ω, but since Δ_t is, as mentioned, a transition time between two regimes we require it to full-fill $3 < \tau < 16$ years. As before, the nonlinearity of the objective function creates multiple local minima, and we use the preliminary grid search to find starting points for the optimiser.

We fit the logarithm of the Dow Jones Industrial Average index prior to the 1929 crash and the $S\&P500$ index prior to the 1987 crash, both starting approximately 8 years prior to the crash thus extending the fitted region from two years to almost eight years, see figures 5a and 6a.

We see that the behavior of the logarithm of the index is well-captured by equation (22) over almost 8 years with a general error below 10% for both crashes, see figures 5b and 6b. Furthermore, the value for t_c is within two weeks of the true date for both crashes. The reason for fitting the logarithm of the index and not the index itself can be seen as an attempt to "de-trend" the index from the underlying exponential growth on long time scales.

As we shall further elaborate in the next section, the small values obtained for the $r.m.s.$ is of course encouraging but not decisive, since we are fitting a function with considerable "elasticity". Hence, it's quite reassuring that the parameter values for the "physical variables" α, ω and Δ_t obtained for the two crashes are rather similar

- $\alpha_{29} \approx 0.63$ compared to $\alpha_{87} \approx 0.68$
- $\omega_{29} \approx 5.0$ compared to $\omega_{87} \approx 8.9$ or $\lambda_{29} \approx 3.5$ compared to $\lambda_{87} \approx 2.0$.
- $\Delta_{t_{29}} \approx 14$ years compared to $\Delta_{t_{87}} \approx 11$ years

4.5 Fitting the Model to Other Periods

In order to investigate the significance of the results obtained for the 1929 and 1987 crashes, we picked at random 50[6] 400-week intervals in the period 1910 to 1996 of the Dow Jones average and launched the fitting procedure described above on these surrogate data sets.

[6]Naturally, we would have liked to use a much larger number, but these 50 intervals already correspond to 4 centuries.

The approximate end-date of the 50 data sets was

1951, 1964, 1950, 1975, 1979, 1963, 1934, 1960, 1936, 1958, 1985, 1884, 1967, 1943, 1991, 1982, 1972, 1928, 1932, 1946, 1934, 1963, 1979, 1993, 1960, 1935, 1974, 1950, 1970, 1980, 1940, 1986, 1923, 1963, 1964, 1968, 1975, 1929, 1984, 1944, 1994, 1967, 1924, 1974, 1954, 1956, 1959, 1926, 1947 and 1965.

The results were very encouraging. Of the 11 fits with a *r.m.s.* comparable with the *r.m.s.* of the two crashes, only 6 data sets produced values for α, ω and Δ_t which were in the same range as the values obtained for the 2 crashes, specifically $0.45 < \alpha < 0.85$, $4 < \omega < 14$ $(1.6 < \lambda < 4.8)$ and $3 < \Delta_t < 16$. All 6 fits belonged to the periods prior to the crashes of 1929, 1962 and 1987. The existence of a crash in 1962 was before these results unknown to us and the identification of this crash naturally strengthens our case. In figure 7 we see the best fit obtained for this crash and except for the "postdicted" time of the crash $t_c \approx 1963.0$, which is rather off as seen in figure 7, the values obtained for the parameters $\alpha_{62} \approx 0.83, \omega_{62} \approx 13, \Delta_{t62} \approx 14$ years is again close to those of the 1929 and 1987 crashes. A comment on the deviation between the obtained value for t_c and the true date of the crash is appropriate. As seen in figure 7, the crash of 1962 was anomalous in the sense that it was "slow". The stock market declined approximately 30% in 3 months and not in less than one week as for the other two crashes. One may speculate on the reason(s) for this and in terms of the model presented here some external shock may have provoked this slow crash before the stock market was "ripe". In fact within our rational expectation model, a bubble that starts to "inflate" with some theoretical critical time t_c can be perturbed and not go to its culmination due to the influence of external shocks. This does not prevent the log-periodic structures to develop up to the time when the course of the bubble evolution is modified by these external shocks.

The results from fitting the surrogate data sets generated from the real stock market index show that fits, which in terms of the fitting parameters corresponds to the 3 crashes mentioned above, are not likely to occur "accidentally". Encouraged by this, we decided on a more elaborate statistical test in order to firmly establish the significance of the results presented above.

4.6 Fitting a GARCH(1,1) Model

We generated 1000 surrogate date sets now of length 400 weeks using the same GARCH(1,1) model used in section 4.1. On each of these 1000 data sets we launched the same fitting routine as for the real crashes. Using the same parameter range for the α, ω and Δ_t as for the 50 random intervals above showed that 66 of the best minima had values in that range. All of these fits, except two, did not resemble the true stock market index prior to the 1929 and 1987 crashes on Wall Street very much, the reason primarily being that they only had one or two oscillations. However, two fits looked rather much like that of the 1929 and 1987 crashes, see the two fit in figure 7. Note that one of the two fits shown in figure 7 have parameter values right on the border of the allowed intervals. This result correspond approximately to the usual 95 % confidence interval for *one* event. Opposed to this, we have here provided several examples of log-periodic signatures before large stock market crashes.

4.7 Fitting Truncated Stock Market Data

An obvious question concerns how long time prior to the crash can one identify the log-periodic signatures described in sections 4.3 and 4.4. There are several reasons for this. Not only because one would like to predict future crashes, but also to further test how robust our results are. Obviously, if the log-period structure of the data is purely accidental, then the parameter values obtained should depend heavily on the size of the time interval used in the fitting. We have hence applied the following procedure. For each of the two rapid crashes fitted above, we have truncated the time interval used in the fitting by removing points and re-launching the fitting procedure described above. Specifically, the logarithm of the S& P500 shown in figure 5 was truncated down to an end-date of $\approx$ 1985 and fitted using the procedure described above. Then $\approx$ 0.16 years was added consecutively and the fitting was relaunched until the full time interval was recovered. In table 1, we see the number of minima obtained for the different time intervals. This number is to some extent rather arbitrary since it naturally depends on the number of points used in the preliminary scan as well as the size of the time interval used for t_c. Specifically, 40.000 points were used and t_c was chosen 0.1 years from the last data point used and 3 years forward. What is more interesting is the number of "physical minima" as defined above and especially the values of $t_c, \alpha, \omega, \tau$ of these fits. The general picture to be extracted from this table, is that a year or more before the crash, the data is not sufficient to give any conclusive results at all. This point correspond to the end of the 4'th oscillation. Approximately a year before the crash, the fit begins to lock-in on the date of the crash with increasing precision. In fact, in four of the last five time intervals, we can find a fit with a t_c, which differs from the true date of the crash by only a few weeks. In order to better investigate this, we show in table 2 the corresponding parameter values for the other 3 physical variables β, ω, τ. The scenario resembles that for t_c and we can conclude that our procedure is rather robust up to approximately 1 year prior to the crash. However, if one wants to actually predict the time of the crash, a major obstacle is the fact that our fitting procedure produces several possible dates for the date of the crash even for the last data set. As a naive solution to this problem we show in table 3 the average of the different minima for $t_c, \alpha, \omega, \tau$. We see that the values for β, ω, τ are within 20% of those for the best prediction, *but* the prediction for t_c has not improved significantly. The reason for this is that the fit in general "over-shoot" the true day of the crash.

The same procedure was used on the logarithm of the Dow Jones index prior to the crash of 1929 and the results are shown in tables 4, 5 and 6. We see that we have to wait until approximately 4 month before the crash before the fit locks in on the date of the crash, but from that point the picture is the same as for the crash in 1987. The reason for the fact that the fit "locks-in" at a later time for the 1929 is obviously the difference in the transition time Δ_t for the two crashes which means that the index prior to the crash of 1929 exhibits fewer distinct oscillations.

A general trend for the two crashes is that the t_c's of the physical minima tends to over-shoot the time of the crash. We have also tried to truncate the data sets with respect to the starting point with 2 years and the effect is quite similar. Even though β, ω, τ stay close to the values obtained using the full interval, the fit starts to overshoot and t_c moves up around 88.2 for the 1987 crash and 30.0 for the 1929 crash. The reason for this is rather obvious. The difference between a pure power law behavior (equation 12) and the discrete versions (equations 18 and 22) is that the pure power law is less constraining with respect to determining t_c. This since t_c is now determined not only by the over-all acceleration quantified by β but also by the frequencies ω and Δ_ω in the cosine. This means that by truncating the data we are in fact removing parts of the oscillations and hence decreasing the accuracy in which we can determine ω and Δ_ω and as a consequence t_c.

To conclude this part of the analysis, we have seen that the three physical variables β, ω and τ are reasonably robust with respect to truncating the data up to approximately a year. In contrast, this is not the case for the timing of the crash t_c, which systematically over-shoot the actual date of the crash. We stress that these results are fully consistent with our rational expectation model of a crash. Indeed, recall that t_c in the formulas (18,22) is *not* the time of the crash but the most probable value of the (skewed to the left) distribution of the possible times of the crash. The occurrence of the crash is a random phenomenon which occurs with a probability that increases as time approaches t_c. Thus, we expect that fits will give values of t_c which are in general close to but *systematically* later than the real time of the crash. The phenomenon of "overshot" that has been clearly documented above is thus fully consistent with the theory. This is one of the major improvement over previous works brought by our combining the rational expectation theory with the imitation model.

5 Conclusion

This paper has drawn a link between crashes in the stock market and critical behavior of complex systems. We have shown how patterns that are typical of critical behavior can arise in prices even when markets are rational. Furthermore, we have provided some empirical evidence suggesting that these characteristic patterns are indeed present, at certain times, in U.S. and Hong Kong stock market indices. This supports our key hypothesis, which is that local imitation between noise traders might cascade through the scales into large- scale coordination and cause crashes. To sum up, the evidence we have presented signifying that large financial crashes are outliers to the distribution of draw downs and as a consequence have their origin in cooperative phenomena are the following.

- The recurrence time for a sudden stock market drop equal or larger than 23.6% was estimated to be 180 centuries. The stock market has sustained three such events in this century.

- In a million years of GARCH(1,1)-trading with student-t noise with four degrees of freedom with a reset of every century, only twice did three draw downs larger than

22% occur and never four. However, three of these "crashes" were symmetric and none preceded by the log-periodic behavior seen in real crashes.

- From 50 randomly sampled 400 weeks intervals of the Dow Jones Average in the period 1910-94, all fits that were found with parameters in the range of those obtained for the 1929 and 1987 crashes could be related to those two crashes *or* to the slow crash of 1962. Furthermore, out of a 1000 surrogate data sets generated from a GARCH(1,1) model, only 66 fits had parameter values comparable with that obtained for the crashes of 1929 and 1987. This correspond approximately to the usual 95 % confidence interval for *one* event. In contrast, we have here provided several examples of log-periodic signatures before large stock market crashes.

- We have tested the prediction ability equation (22) by truncating the stock market indices with respect to the end-point of the data set. In general, our procedure tends to over-shoot the actual date of the crash and hence its prediction ability is not very impressive. This is fully consistent with our rational expectation theory according to which the critical time t_c is only the most probable value of the distribution of times of the crash, which occurs (if it occurs) close to but before t_c. However, these tests show that our fitting procedure is robust with respect to *all* variables in equation (22).

- Three of the largest crashes of this century, the 1929 and 1987 crashes on Wall Street and the 1997 crash in Hong-Kong, have been preceded by log-periodic behavior well-described by equations (18) and (22). In addition, prior to the slow crash of 1962 and the turmoil on Wall Street in late October 1997 log-periodic signatures have been observed with parameter values consistent with the three other crashes. The turbulence of Oct. 1997 in the US market was also preceded by a log-periodic structure that led to an estimation of t_c close to the end of the year, with the other parameters consistent with the values found for the other crashes. No crash occurred. This may be interpreted as a realization, compatible with our rational expectation theory of a crash, illustrating that there is a non-vanishing probability for no crash to occur even under inflating bubble conditions.

Acknowledgements: We thank M. Brennan, B. Chowdhry, W.I. Newman, H. Saleur and P. Santa-Clara for useful discussions. We also thank participants to a seminar at UCLA for their feedback. Errors are ours.

A Derivation of the exponential distribution (23) of draw downs

We use the dayly time scale as the unit of time. Call p_+ (resp. $p_- = 1 - p_+$) the probability for the stock market to go up (resp. to go down) over a day. If $P_1(x)$ is the probability density function (pdf) of *negative* dayly price increments, the distribution $P(DD)$ of draw downs (continuous decreases) is given by

$$P(DD) = \sum_{n=1}^{\infty} p_+^2 \; p_-^n \; P_1^{\otimes n}(DD) \;, \tag{26}$$

where $P_1^{\otimes n}$ denotes the distribution obtained by n convolutions of P_1 with itself, corresponding to the fact that DD is the sum of n random variables. The factor $p_+^2 \; p_-^n$ weights the probability that a draw down results from consecutive n losses starting after and finishing before a gain. Since $P_1(x)$ is defined for $-\infty < x \leq 0$, the relevant tool is the Laplace transform with argument k: the Laplace transform of $P_1^{\otimes n}(DD)$ is the nth power of the Laplace transform $\hat{P}_1(k)$ of $P_1(x)$. When applied to (26), it gives after summation

$$\hat{P}(k) = p_+^2 \; \frac{p_- \hat{P}_1 k}{1 - p_- \hat{P}_1 k} \;. \tag{27}$$

The dayly price distribution is closely approximated by an exponential distribution $P_1(x) = \frac{c}{2} e^{-c|xl}$ (Bouchaud and Potters, 1997; Laherrère and Sornette, 1998). Then, $\hat{P}_1(k) = \frac{c/2}{c+k}$. Put in (27), this gives $\hat{P}(k) = \frac{c}{2} \; \frac{p_+^2 \; p_-}{c(1-p_-)+k}$. By inverse Laplace transform, we get

$$P(DD) \propto e^{-cp_+|x|} \;. \tag{28}$$

We see that the typical amplitude DD_c of "normal" draw downs is $1/p_+$ times the typical amplitude $1/c$ of the negative price variations. The price distribution of the Dow Jones is very close to being symmetric. Thus, p_+ is close to $1/2$ and thus DD_C is about twice the typical dayly price variation.

B GARCH(1,1) model of the stock market

Estimating the five parameters of a GARCH(1,1) (Generalised Autoregressive with Conditional Heteroskedasticity[7]) model of some process is a simple optimisation problem.

B.1 Parameters

The 5 parameters are related to the data x_t as follows.

$$x_t = \mu + h_t \epsilon_t \tag{29}$$

$$h_t = \alpha + \beta \epsilon_{t-1}^2 + \gamma h_{t-1} \tag{30}$$

where $\epsilon(t)$ belongs to a Student-t distribution with κ degrees of freedom and mean 0 and variance $h(t)$. The interpretation of the parameters are as follows:

μ Mean of the process

κ Number of degrees of freedom ($i.e$, "fat-tailedness") of the Student t-distribution

α Controls the average level of volatility

β Controls the impact of short-term shocks to volatility

γ Controls the long-term memory of volatility

Furthermore, the range of the parameters are

μ Any real number.

κ Any number strictly above 2, so that the variance is finite; the lower the κ the fatter the tails; this number needs not be an integer.

α Between 0 and 1.

β Between 0 and 1.

γ Between 0 and 1.

In addition there is the constraint $\alpha + \beta \leq 1$.

[7] When the errors of different points have different variances but are uncorrelated with each other, then the errors are said to exhibit heteroskedasticity. If the variances are correlated, the heteroskedasticity is said to be conditional.

B.2 Initial values of the parameters

These are the values recommended in order to start the optimisation.

μ Take the sample mean $m = \sum_{t=1}^{T} x_t/t$.

κ Take 4.

α Take 0.05 times the sample variance: $0.05 \times v$, where the sample variance is defined by $v = \sum_{t=1}^{T}(x_t - m)^2/(t-1)$.

β Take 0.05.

γ Take 0.90.

The factor 0.05 for α comes from the fact that the average level of the variance is equal to $\alpha/(1-\beta-\gamma)$, as can be seen by replacing h_t with its average level in equation (30). For $\beta = 0.05$ and $\gamma = 0.90$, this gives the factor 0.05.

B.3 Objective function

The five parameters are obtained by maximising the likelihood function L over the parameter range defined above. This is a function of the 5 parameters and the data (x_t for $t = 1, \ldots, T$), which is computed in three steps.

1. Do a loop over $t = 1, \ldots, T$ to compute the residuals $\epsilon_t = x_t - \mu$.

2. Do a loop to compute recursively the conditional variance h_t at time t. Initialise the recursion by: $h_1 = v$ (the sample variance computed above). Iterate for $t = 2, \ldots, T$ with $h_t = \alpha + \beta\epsilon_{t-1}^2 + \gamma h_{t-1}$.

3. Compute the likelihood function as:

$$L = \log\left[\Gamma\left(\frac{\kappa+1}{2}\right)\right] - \log\left[\Gamma\left(\frac{\kappa}{2}\right)\right] - \frac{1}{2}\log(\kappa) \tag{31}$$

$$- \frac{1}{T}\sum_{t=1}^{T}\left[\frac{\log(h_t)}{2} + \frac{\kappa+1}{2}\log\left(1 + \frac{\epsilon_t^2}{\kappa h_t}\right)\right] \tag{32}$$

where Γ denotes the gamma function.

This procedure yields the value of the logarithm of the likelihood, given the data and the five parameters. Feed this likelihood function into an optimiser and it will give the five GARCH(1,1) parameters.

Specifically for the GARCH(1,1) that generated the surrogate data used in subsection 4.6, we used $\mu = 4.38\ 10^{-4}$, $\gamma = 0.922$, $\alpha = 2.19 \cdot 10^{-5}$, $\beta = 0.044$ and the Student-t distribution had $\kappa = 4$ degrees of freedom and was generated by a NAG-library routine.

Furthermore, the first 5000 data points was always discarded in order to remove any sensitivity on the initial conditions. The logarithm of the index $\ln I$ was then calculated as

$$\ln I_{t+1} = \ln I_t + x_{t+1}, \tag{33}$$

using equations (29) and (30) and the arbitrary initial condition $\ln I_0 = 2$.

References

[Aoki, 1996] Aoki, M., 1996. New approaches to macroeconomic modeling – Evolutionary stochastic dynamics, multiple equilibria, and externalities as field effects. *Cambridge University Press*, Cambridge, UK.

[Anderson et al., 1988] Anderson, P., Arrow, K., and Pines, D. (1988). *The Economy as a Complex Evolving System.* Addison-Wesley.

[Arnéodo et al., 1998] Arnéodo, A., J.-F. Muzy and D. Sornette, 1998. "Direct" causal cascade in the stock market. *European Physical Journal B* 2, 277–282.

[Arthur, 1990] Arthur, B. (1990). *Scientific American*, 261:80.

[Black, 1986] Black, F. (1986). Noise. *Journal of Finance*, 41:529–543.

[Blanchard, 1979] Blanchard, O. J. (1979). Speculative bubbles, crashes and rational expectations. *Economics Letters*, 3:387–389.

[Boissevain and Mitchell, 1973] Boissevain, J. and Mitchell, J. (1973). *Network analysis: Studies in human interaction.* Mouton.

[Bollerslev et al., 1992] Bollerslev, T., Chou, R.Y. and Kroner, K.F., 1992. ARCH Modeling in Finance: A Review of the Theory and Empirical Evidence. *Journal of Econometrics* 51, 5–59.

[Bouchaud and Potters, 1997] Bouchaud, J.P. and M. Potters, 1997. Théorie des risques financiers. Aléa Saclay, CEA.

[Brock et al., 1992] Brock, W., J. Lakonishok and B. LeBaron, 1992. Simple technical trading rules and the stochastic properties of stock returns. *J. Finance* 47, 1731-1764.

[Cont, 1998] Cont, R., 1998. PhD Thesis, Univ. of Jussieu.

[Cvijović and J. Klinowski, 1995] Cvijović D. and Klinowski J. (1995) Taboo search – An approach to the multiple minima problem. *Science* 267(5), 664-666.

[Cosmides and Tooby, 1994] Cosmides, L. and Tooby, J., 1994. Better than rational – evolutionary psychology and the invisible hand. *American Economic Review* 84(2), 327–332.

[DeBondt and Thaler, 1985] DeBondt, W. and Thaler, R., 1985. Does the stock market overreact? *Journal of Finance* 40, 793–805.

[Derrida et al., 1983] Derrida, B., De Seze, L., and Itzykson, C., 1983. Fractal structure of zeros in hierarchical models. *Journal of Statistical Physics* 33, 559.

[Dubrulle et al., 1997] Dubrulle, D., F. Graner and D. Sornette, eds., 1997. Scale invariance and beyond. EDP Sciences and Springer, Berlin.

[Fama, 1970] Fama, E. (1970). Efficient markets: A review of theory and empirical work. *Journal of Finance* 25,383–417.

[Fama, 1991] Fama, E. (1991). Efficient capital markets: II. *Journal of Finance* 46,1575–1618.

[Feigenbaum and Freund, 1998] Feigenbaum, J.A. and P.G.O. Freund, 1996. Discrete scale invariance in stock markets before crashes. *Int. J. Mod. Phys.* 10, 3737–3745.

[Feigenbaum and Freund, 1998] Feigenbaum, J.A. and P.G.O. Freund, 1998. Discrete Scale Invariance and the "Second Black Monday". *Int. J. Mod. Phys.* 12, 57–60.

[Gluzman and Yukalov, 1998] Gluzman, S. and V. I. Yukalov, 1998. Renormalization Group Analysis of October Market Crashes. *Mod. Phys. Lett. B* 12, 75–84.

[Goldenfeld, 1992] Goldenfeld, N., 1992. Lectures on phase transitions and the renormalization group. *Addison-Wesley Publishing Company.* Reading, Massachussets.

[Krugman, 1998] Krugman, P., 1998. The Confidence Game, How Washington worsened Asia's crash. The New Republic, October 5.

[Laherrére and Sornette, 1998] Laherrère, J. and D. Sornette, 1998. Stretched exponential distributions in Nature and Economy: "Fat tails" with characteristic scales. *European Physical Journal B* 2, 525–539.

[Laloux et al., 1998] Laloux, L., M. Potters, R. Cont, J.-P. Aguilar and J.-P. Bouchaud, 1998. Are Financial Crashes Predictable? preprint cond-mat/9804111

[Leland, 1988] Leland, H. (1988). Portfolio insurance and October 19th. *California Management Review* 30, 80–89.

[Liggett, 1985] Liggett, T.M., 1985. Interacting particle systems. New York : Springer-Verlag.

[Liggett, 1997] Liggett, T.M., 1997. Stochastic models of interacting systems. *The Annals of Probability* 25, 1-29.

[Onsager, 1944] Onsager, L., 1944. *Physics Review* 65, 117.

[Pomeau, 1982] Pomeau, Y., 1982. Symétrie des fluctuations dans le renversement du temps. *J. Physique* 43, 859-867.

[Potters et al., 1998] Potters, M., Cont, R. and Bouchaud, J.P., 1998. Financial markets as adaptive systems. *Europhysics Letters* 41(3-, 239–244.

[Press W.H. *et al.* , 1992] Press W.H. *et al.* (1992). Numerical Recipes Cambridge University Press.

[Rubinstein, 1988] Rubinstein, M. (1988). Portfolio insurance and the market crash. *Financial Analysts Journal* 44, 38–47.

[Shiller, 1989] Shiller, R.J., 1989. Market volatility. Cambridge, Mass. : MIT Press.

[Sornette, 1998] Sornette, D. (1998). Discrete scale invariance and complex dimensions. *Physics Reports* 297, 239–270.

[Sornette and Johansen, 1997] Sornette, D. and Johansen, A. (1997). Large financial crashes. *Physica A* 245, 411–422.

[Sornette and Johansen, 1998] Sornette, D. and A. Johansen, 1998. A Hierarchical Model of Financial Crashes. *Physica A*, in press.

[Sornette et al., 1996] Sornette, D., Johansen, A., and Bouchaud, J.-P. (1996). Stock market crashes, precursors and replicas. *Journal of Physics I France* 6,167–175.

[Vandewalle et al., 1998a] Vandewalle, N., Ph. Boveroux, A. Minguet and M. Ausloos, 1998. The krach of October 1987 seen as a phase transition: amplitude and universality. *Physica A* 255(1-2), 201–210.

[Vandewalle et al., 1998b] Vandewalle, N., M. Ausloos, P. Boveroux and A. Minguet, 1998. How the financial crash of October 1997 could have been predicted. *European Physical Journal B* 4(2), 139–41.

[Van Norden and Schaller, 1994] Van Norden, S. and H. Schaller, 1994. Speculative behavior, regime-switching and stock market crashes. working paper, june 1994.

[Wilson, 1979] Wilson, K. G., 1979. Problems in Physics with many scales of length. *Scientific American* 241(2), 158–179.

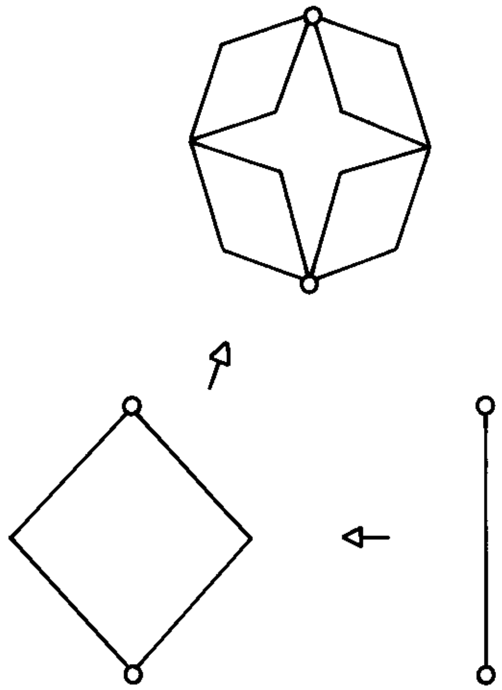

Figure 1: The first three steps of the recursive construction of the hierarchical Diamond Lattice.

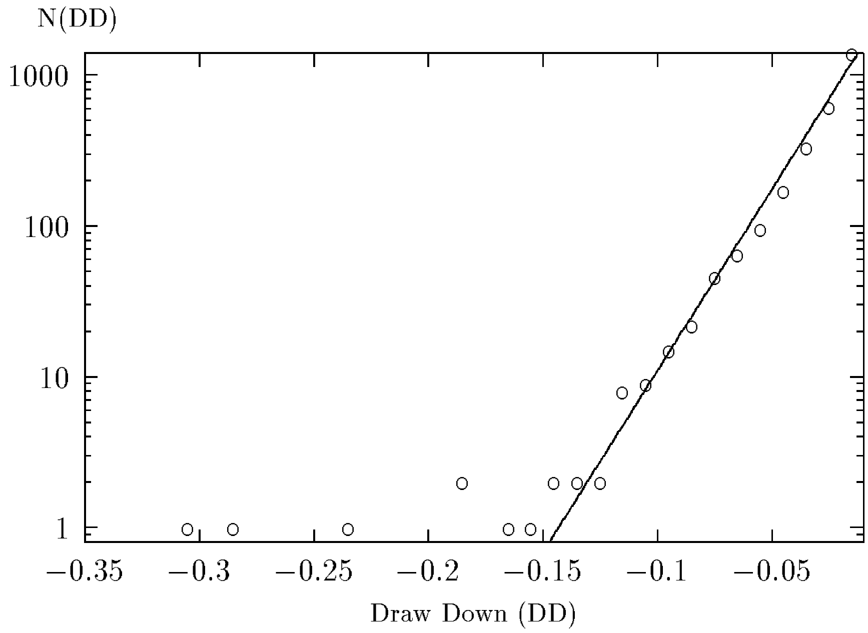

Figure 2: Number of times a given level of draw down has been observed in this century in the Dow Jones Average. The bin-size is 1%. A threshold of 1% has been applied. The fit is equation (23) with $N_0 = 2789$ and $\mathrm{DD}_c \approx 0.018$.

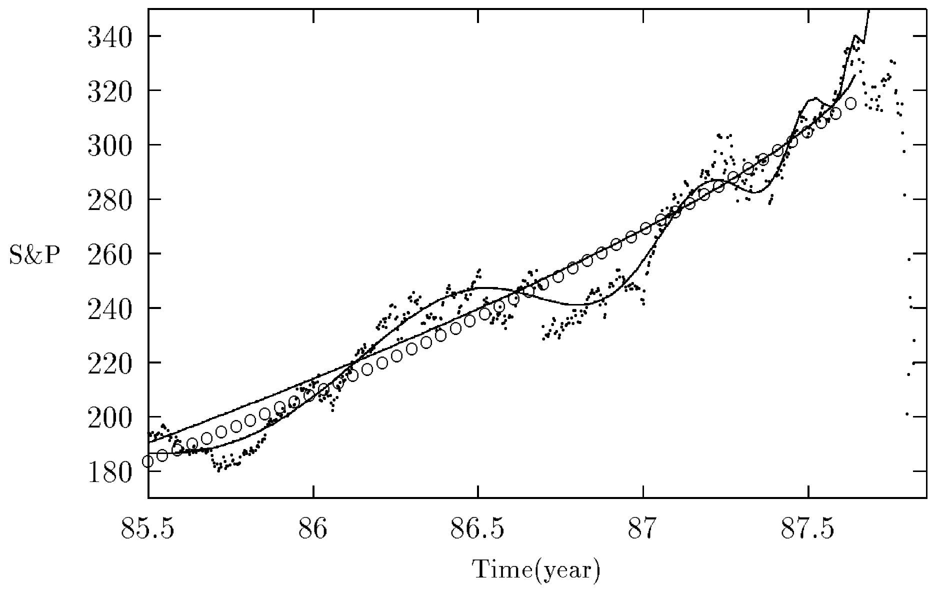

Figure 3: The New York stock exchange index $S\&P500$ from July 1985 to the end of 1987 corresponding to 557 trading days. The $\circ$ represent a constant return increase in terms an exponential with a characteristic increase of ≈ 4 years^{-1} and var $(F_{exp}) \approx 113$. The best fit to a pure power-law gives $A \approx 327$, $B \approx -79$, $t_c \approx 87.65$, $\alpha \approx 0.7$ and var$_{pow} \approx 107$. The best fit to (18) gives $A \approx 412$, $B \approx -165$, $t_c \approx 87.74$, $C \approx 12$, $\omega \approx 7.4$, $\phi = 2.0$, $\alpha \approx 0.33$ and var$_{lp} \approx 36$.

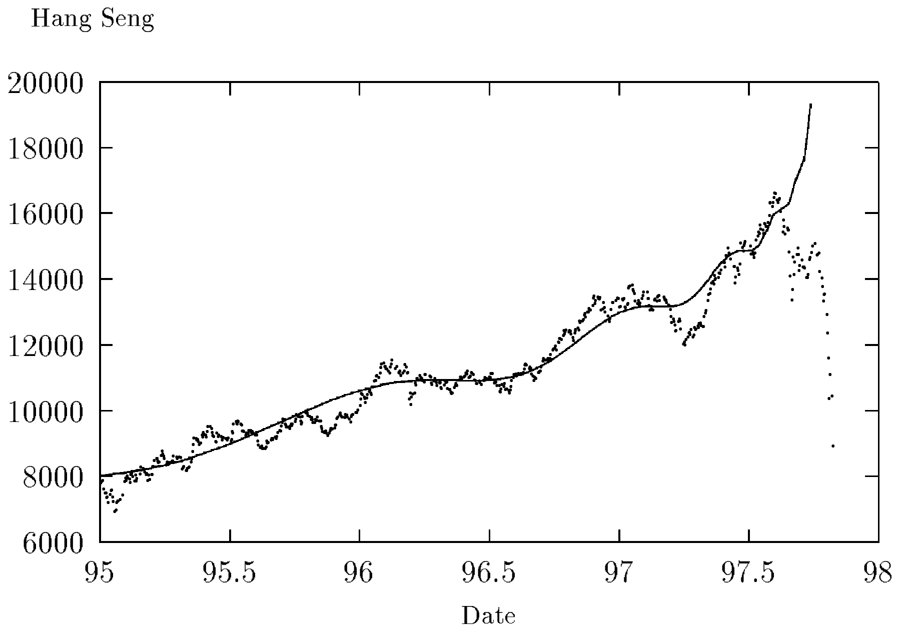

Figure 4: The Hang Seng index of the Hong Kong stock exchange approximately two and half years prior to the crash of October 1997. The best fit to (18) gives $A \approx 2. \cdot 10^4$, $B \approx -8240$, $t_c \approx 97.74$, $C \approx -397$, $\omega \approx 7.5$, $\phi \approx 1.0$, $\alpha \approx 0.34$ and $r.m.s. \approx 436$.

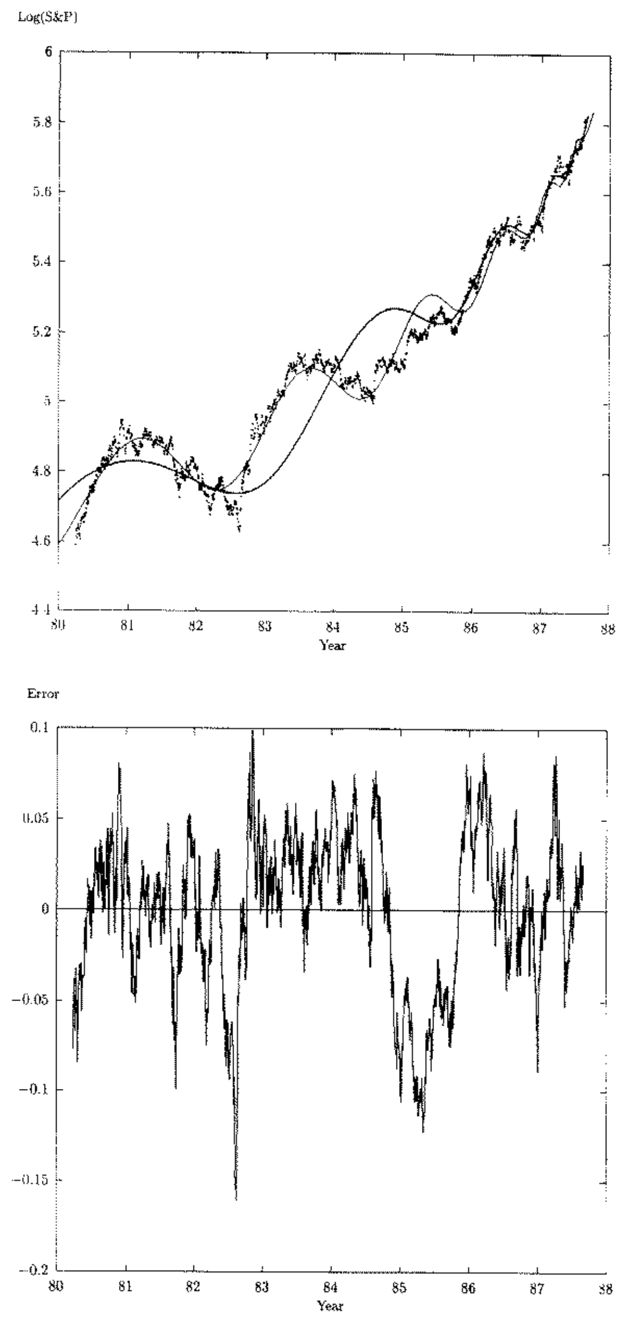

a)

b)

Figure 5: a) The $S\&P500$ from January 1980 to September 1987 and best fit by (22) (thin line) with $r.m.s. \approx 0.043$, $t_c \approx 1987.81$, $\alpha \approx 0.68$, $\omega \approx 8.9$, $\Delta_\omega \approx 18$, $\Delta_t \approx 11$ years, $A \approx 5.9$, $B \approx -0.38$, $C \approx 0.043$. The thick line is the fit shown in figure 3 extended to the full time interval. The comparison with the thin line allows one to visualise the frequency shift described by equation (22).
b) The relative error of the fit.

a)

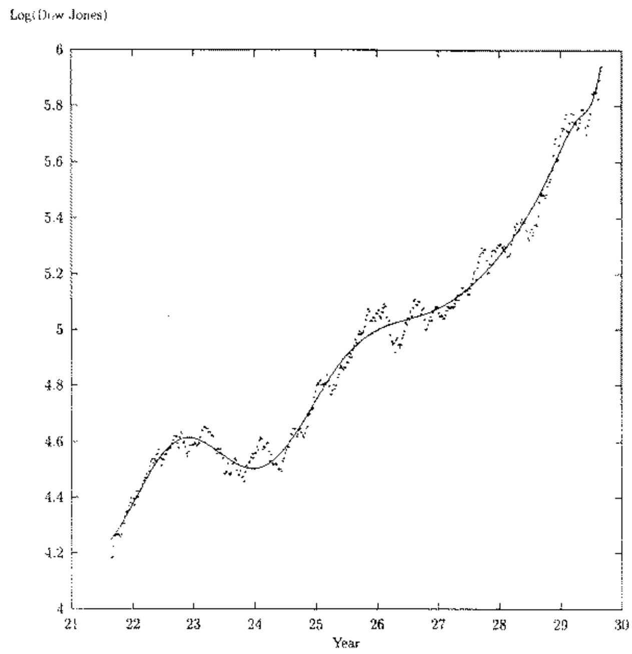

b)

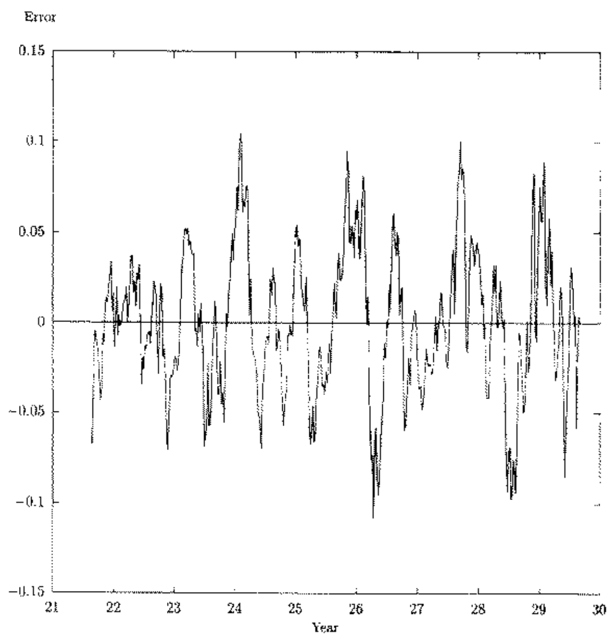

Figure 6: a) The Dow Jones Average from June 1921 to September 1929 and best fit by equation (22). The parameters of the fit are r.m.s.≈ 0.041, $t_c \approx 1929.84$ year, $\alpha \approx 0.63$, $\omega \approx 5.0$, $\Delta_\omega \approx -70$, $\Delta_t \approx 14$ years, $A \approx 61$, $B \approx -0.56$, $C \approx 0.08$.
b) The relative error of the fit.

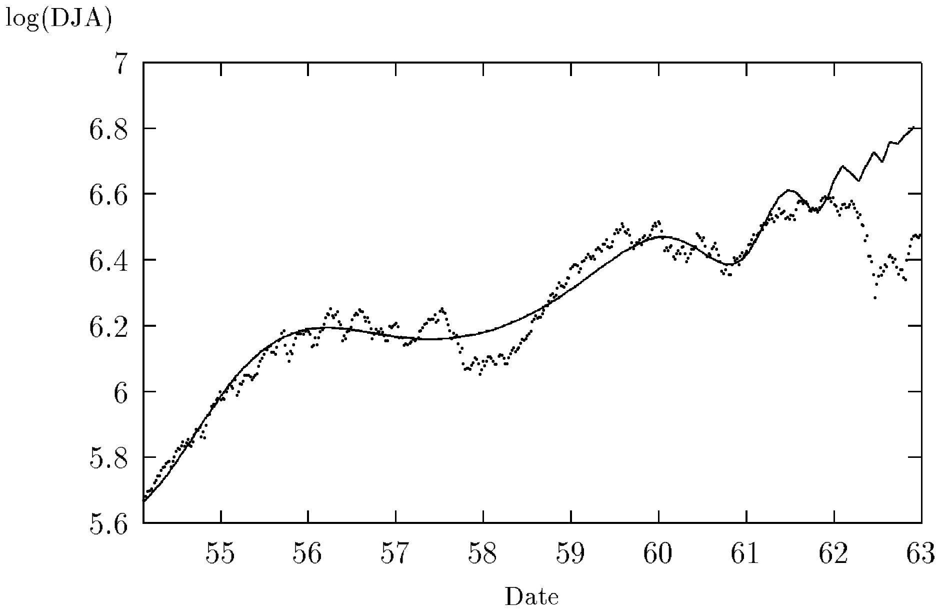

Figure 7: The "slow crash" of 1962 and the best fit by equation (22).

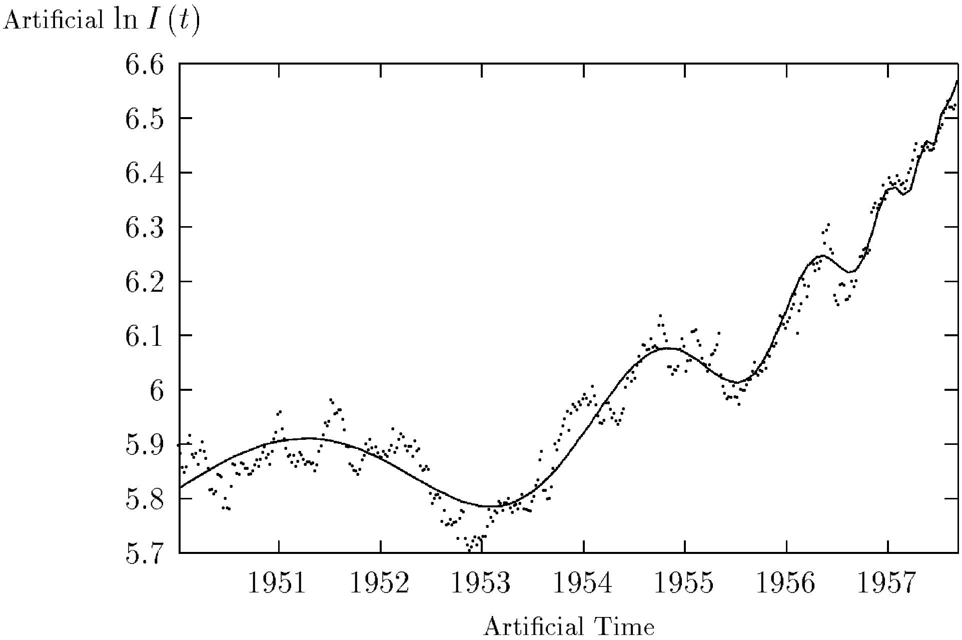

Figure 8: Surrogate data and best fit by equation (22) is $A \approx 6.6$, $B \approx -.34$, $C \approx -.047$, $\alpha \approx 0.74$, $t_c \approx 1957.7$, $\phi \approx -.13$ $\omega \approx 9.1$, $\Delta_t \approx 4.1$, $\Delta\omega \approx -2.1$ and $r.m.s. \approx 0.037$

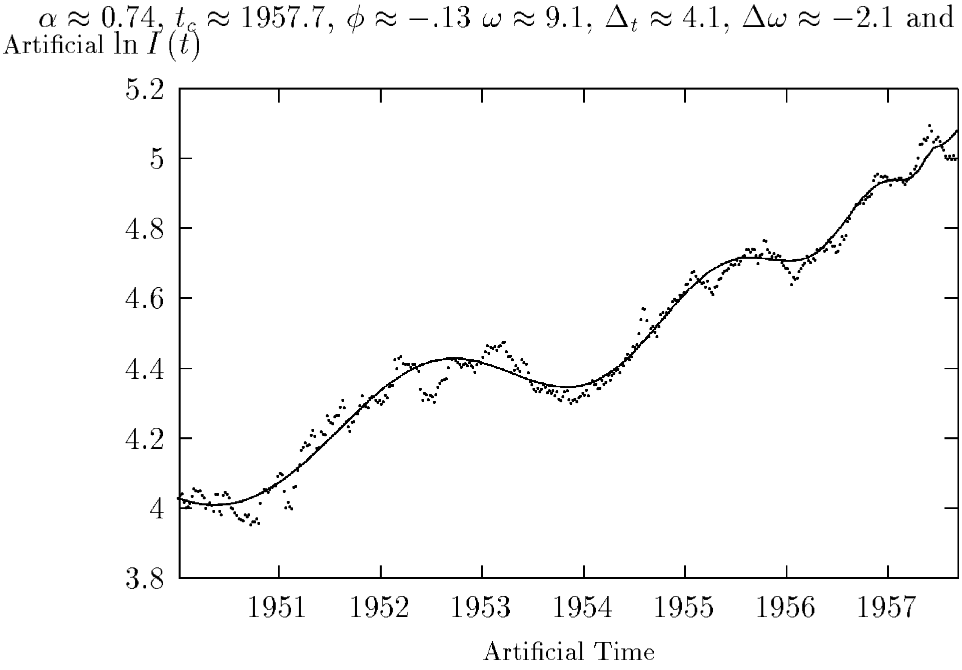

Figure 9: Surrogate data and best fit by equation (22) is $A \approx 5.1$, $B \approx -.24$, $C \approx -.037$, $\alpha \approx 0.85$, $t_c \approx 1957.7$, $\phi \approx -3.9$ $\omega \approx 5.5$, $\Delta_t \approx 7.7$, $\Delta\omega \approx 11$ and $r.m.s. \approx 0.036$

end-date	total # minima	"physical" minima	t_c of "physical" minima
85.00	33	1	86.52
85.14	25	4	4 in [86.7 : 86.8]
85.30	26	7	5 in [86.5 : 87.0], 2 in [87.4 : 87.6]
85.46	29	8	7 in [86.6 : 86.9] ,1 with 87.22
85.62	26	13	12 in [86.8 : 87.1] ,1 with 87.65
85.78	23	7	87.48, 5 in $[87.0 - 87.25]$, 87.68
85.93	17	4	$87.25, 87.01, 87.34, 86.80$
86.09	18	4	$87.29, 87.01, 86, 98, 87.23$
86.26	28	7	5 in [87.2 : 87.4], $86.93, 86.91,$
86.41	24	4	$87.26, 87.36, 87.87, 87.48$
86.57	20	2	$87.67, 87.34$
86.73	28	7	4 in [86.8 : 87.0], $87.37, 87.79, 87.89$
86.88	22	1	87.79
87.04	18	2	$87.68, 88.35$
87.20	15	2	$87.79, 88.03$
87.36	15	2	$88.19, 88.30$
87.52	14	3	$88.49, 87.92, 88.10$
87.65	15	3	$87.81, 88.08, 88.04$

Table 1: Number of minima obtained by fitting different truncated versions of the S& P500 time series shown in figure 5, using the procedure described in the text.

end-date	t_c	α	ω	τ
86.88	87.79	0.66	5.4	7.8
87.04	$87.68, 88.35$	$0.61, 0.77$	$4.1, 13.6$	$12.3, 10.2$
87.20	$87.79, 88.03$	$0.76, 0.77$	$9.4, 11.0$	$10.0, 9.6$
87.36	$88.19, 88.30$	$0.66, 0.79$	$7.3, 12.2$	$7.9, 8.1$
87.52	$88.49, 87.92, 88.10$	$0.51, 0.71, 0.65$	$12.3, 9.6, 10.3$	$10.2, 9.8, 9.8$
87.65	$87.81, 88.08, 88.04$	$0.68, 0.69, 0.67$	$8.9, 10.4, 10.1$	$10.8, 9.7, 10.2$

Table 2: For the last five time intervals shown in table 1, we show the corresponding parameter values for the other three variables β, ω, τ.

end-date	t_c	α	ω	τ
86.88	87.79	0.66	5.4	7.8
87.04	88.02	0.69	8.6	11.3
87.20	87.91	0.77	10.20	9.8
87.36	88.25	0.73	9.6	8.0
87.52	88.17	0.62	10.7	9.9
87.65	87.98	0.68	9.8	10.2

Table 3: The average of the values listed in table 2.

end-date	total # minima	"physical" minima	t_c of "physical" minima
27.37	12	1	31.08
27.56	14	2	$30.44, 30.85$
27.75	24	1	30.34
27.94	21	1	31.37
28.13	21	4	$29.85, 30.75, 30.72, 30.50$
28.35	23	4	$30.29, 30.47, 30.50, 36.50$
28.52	18	1	31.3
28.70	18	1	31.02
28.90	16	4	$30.40, 30.72, 31.07, 30.94$
29.09	19	2	$30.52, 30.35$
29.28	33	1	30.61
29.47	24	3	$29.91, 30.1, 29.82$
29.67	23	1	29.87

Table 4: Same as table 1 for the Oct. 1929 crash

end-date	t_c	α	ω	τ
28.90	$30.40, 30.72, 31.07, 30.94$	$0.60, 0.70, 0.70, 0.53$	$7.0, 7.6, 10.2, 13.7$	$12.3, 9.5, 9.0, 11.6$
29.09	$30.52, 30.35$	$0.54, 0.62$	$11.0, 7.8$	$12.6, 10.2$
29.28	30.61	0.63	9.5	9.5
29.47	$29.91, 30.1, 29.82$	$0.60, 0.67, 0.69$	$5.8, 6.2, 4.5$	$15.9, 11.0, 10.9$
29.67	29.87	0.61	5.4	15.0

Table 5: Same as table 2 for the Oct. 1929 crash

end-date	t_c	α	ω	τ
28.90	30.78	0.63	9.6	10.6
29.09	30.44	0.58	9.4	11.4
29.28	30.61	0.63	9.5	9.5
29.47	29.94	0.65	5.5	12.6
29.67	29.87	0.61	5.4	15.0

Table 6: The average of the values listed in table 5.