Predicting financial crashes using discrete scale invariance

Anders Johansen, Didier Sornette, and Olivier Ledoit

The authors present a synthesis of all the available empirical evidence in the light of recent theoretical developments for the existence of characteristic log-periodic signatures of growing bubbles in a variety of markets, including eight unrelated crashes from 1929 to 1998 on stock markets as diverse as the US, Hong Kong, and Russian markets, and on currencies. It is suggested that no major financial crash preceded by an extended bubble has occurred in the past two decades without exhibiting such log-periodic signatures.

1. INTRODUCTION

Two recent papers (Johansen and Sornette 1999a, Johansen *et al*. 2000) have presented increasing evidence on the October 1929 and 1987 Wall Street and the October 1987 Hong Kong crashes, on the August 1998 global market events, and on the 1985 Forex event on the US dollar for the following hypothesis (Sornette *et al*. 1996) describing stock market crashes: that they are caused by a slow build-up of long-range time correlations reflecting those between traders, leading to a collapse of the stock market in one critical instant. This build-up manifests itself as an overall power law acceleration in the price decorated by 'log-periodic' precursors. Here, 'log-periodicity' refers to a sequence of oscillations with progressively shorter cycles of a period decaying according to a geometrical series. In addition, extensive statistical tests have been performed (Johansen *et al*. 2000, Johansen 1997) to show that the reported 'log-periodic' structures essentially never occurred in $\sim 10^5$ years of synthetic trading following a 'classical' time-series model, the GARCH(1,1) model with Student *t*-statistics, often used as a benchmark in academic circles as well as by practitioners. Thus, the null hypothesis that log-periodicity could result simply from random fluctuations is strongly rejected.

From a theoretical view point, we have proposed a rational expectation model of bubbles and crashes with two main components: (1) we assume that a crash may be caused by *local* self-reinforcing imitation processes between noise traders that can be quantified by the theory of critical phenomena developed in the physical sciences; (2) we allow for a remuneration of the risk of a crash by a higher rate of growth of the bubble, which is a consequence of the fact that the crash is not a certain deterministic outcome of the bubble and, as a consequence, it remains rational for traders to remain invested provided they are suitably compensated.

Since our initial proposition (Sornette *et al*. 1996), several works have extended the empirical investigation (Feigenbaum and Freund 1996, 1998,

5

Gluzman and Yukalov 1998, Vandewalle *et al.* 1998a, 1998b, Johansen and Sornette 1999b, Drozdz *et al.* 1999, Johansen and Sornette 1999c) as well as the theoretical framework of critical phenomena (Sornette and Johansen 1997, 1998). However, a recent comment (Laloux *et al.* 1999) and a short assessment note (Ilinski 1999) have raised concerns about the reality of log-periodic signatures in bubbles preceding crashes and the relevance of the concept of critical phenomena to financial crashes in general.

The purpose of the present paper is primarily to present a complete up-to-date synthesis of the available evidence containing all previously reported cases, together with four new cases, which include the correction of the US dollar against the Canadian dollar and the Japanese yen starting in August 1998 as well as the bubble on the Russian market and its ensuing collapse in 1997–98. We are able to show a remarkable *universality* of the results for all events, with approximately the same value of the fundamental scaling ratio λ characterizing the log-periodic signatures. We also address briefly the concern raised by Ilinski (1999) about the validity of our model (Johansen and Sornette 1999a, Johansen *et al.* 2000). Considering the robustness of the results presented here, our analysis opens the road to financial crash forecasting; see, for example, footnote 12 of Stauffer and Sornette (1998) and Vandewalle *et al.* (1998b).

A similar analysis of log-periodicity for 'antibubbles' have been used to issue such a forecast in December 1998 for the recovery of the Nikkei in 1999 (Johansen and Sornette 1999c). At the time of writing (May 1999), the forecast, performed at a time when the Nikkei was at its lowest, seems to have correctly captured the change of regime and the overall upward trend since the beginning of this year.

2. THEORETICAL FRAMEWORK

In order to put the empirical analysis in perspective, it is useful to recapitulate the main features of the model (Johansen and Sornette 1999a, Johansen *et al.* 2000). This will also allow us to stress its robustness with respect to other choices in the model construction and address the criticism made by Ilinski (1999).

2.1 *Ingredients*

- Our key assumption is that a crash may be caused by *local* self-reinforcing imitation between traders (Shiller 1990, 1995, Censi *et al.* 1996, Cutler *et al.* 1990, Keim and Madhavan 1995, Nelson 1994). This self-reinforcing imitation process leads to the blossoming of a bubble. If the tendency for traders to 'imitate' their 'friends' increases up to a certain point called the 'critical' point, many 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 requires a probabilistic description: a crash is *not* a certain outcome of the bubble but can be characterized by its

hazard rate $h(t)$, i.e. the probability per unit time that the crash will happen in the next instant provided it has not happened yet.

- 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, because there is a finite probability of 'landing smoothly', i.e. of attaining the end of the bubble without crash. In this model, the ability to predict the critical date is perfectly consistent with the behavior of the rational agents: they all know this date, the crash may happen anyway, and they are unable to make any abnormal risk-adjusted profits by using this information.

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 have some degree of randomness. 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 only the most probable time of the crash.

The model does not impose any constraint on the amplitude of the crash. If we assume it proportional to the current price level, then the natural variable is the logarithm of the price. If, instead, we assume that the crash amplitude is a finite fraction of the gain observed during the bubble, then the natural variable is the price itself (Johansen and Sornette 1999a). We are aware that the standard economical proxy is the logarithm of the price, and not the price itself, since only relative variations should play a role. However, as we shall see below, different price dynamics give both possibilities.

In the stylized framework of a purely speculative asset that pays no dividends and in which we ignore the interest rate, risk aversion, information asymmetry, and the market-clearing condition, rational expectations are simply equivalent to the familiar martingale hypothesis:

$$\mathsf{E}_t[p(t')] = p(t) \quad \text{for all } t' > t, \tag{1}$$

where $p(t)$ denotes the price of the asset at time t and $\mathsf{E}_t[\cdot]$ denotes the expectation conditional on information revealed up to time t. In the following, we will relax the hypothesis of no risk aversion and show that the predictions of the model are very robust. This is important because risk aversion is often considered as a key ingredient controlling the crash dynamics and the time dependence of the risk perception by market participants also controls in part the precrash dynamics.

We model the occurrence of a crash as a jump process j whose value is 0 before the crash and 1 afterwards. The random nature of the crash occurrence is modeled by the cumulative distribution function $Q(t)$ of the time of the crash, the probability density function $q(t) = \mathrm{d}Q/\mathrm{d}t$ and the hazard rate $h(t) = q(t)/[1 - Q(t)]$. The hazard rate is the probability per unit of time that the crash will happen in the next instant provided that it has not happened yet.

Assume for simplicity that, during a crash, the price drops by a fixed

percentage $\kappa \in (0, 1)$, say between 20 and 30% of the price increase above a reference value p_1. Then, the dynamics of the asset price before the crash are given by

$$dp = \mu(t)p(t)\,dt - \kappa[p(t) - p_1]\,dj, \tag{2}$$

where j denotes a jump process whose value is 0 before the crash and 1 afterwards. Taking $\mathsf{E}_t[dp]$ and using the fact that $\mathsf{E}_t[dj] = h(t)\,dt$ leads to

$$\mu(t)p(t) = \kappa[p(t) - p_1]h(t). \tag{3}$$

In other words, if the crash hazard rate $h(t)$ increases, the return μ increases to compensate the traders for the increasing risk. Plugging (3) into (2), we obtain an ordinary differential equation. For $p(t) - p(t_0) < p(t_0) - p_1$, its solution is

$$p(t) \approx p(t_0) + \kappa[p(t_0) - p_1]\int_{t_0}^{t} h(t')\,dt' \quad \text{before the crash.} \tag{4}$$

The integral $\int_{t_0}^{t} h(t')\,dt'$ is the cumulative probability of a crash until time t. We have found that this regime, where the price itself is the relevant observable, applies to relatively short time scales of approximately two to three years prior to the crash studied here.

If, instead, the price drops by a fixed percentage $\kappa \in (0, 1)$ of the price, the dynamics of the asset price before the crash are given by

$$dp = \mu(t)p(t)\,dt - \kappa p(t)\,dj. \tag{5}$$

We then get

$$\mathsf{E}_t[dp] = \mu(t)p(t)\,dt - \kappa p(t)h(t)\,dt = 0, \tag{6}$$

which yields

$$\mu(t) = \kappa h(t). \tag{7}$$

The corresponding equation for the price is

$$\log\!\left(\frac{p(t)}{p(t_0)}\right) = \kappa \int_{t_0}^{t} h(t')\,dt' \quad \text{before the crash.} \tag{8}$$

This gives the logarithm of the price as the relevant observable. It has been applied successfully to the 1929 and 1987 Wall Street crashes up to about 7.5 years prior to the crash (Sornette and Johansen 1997, Johansen *et al.* 2000).

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 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, p. 389). It may go against 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 crash hazard rate $h(t)$.

Ilinski (1999) raises the concern that the martingale condition (1) leads to a model that "assumes a zero return as the best prediction for the market". He continues: "No need to say that this is not what one expects from a perfect model of market bubble! Buying shares, traders expect the price to rise and it is reflected (or caused) by their prediction model. They support the bubble and the bubble supports them!"

In other words, Ilinski (1999) criticizes a key economic element of our model: market rationality. We have captured this by assuming that the market level is expected to stay constant, as written in equation (1). Ilinski (1999) claims that this equation (1) is wrong because the market level does not stay constant in a bubble: it rises, almost by definition.

This misunderstanding addresses a rather subtle point of the model and stems from the difference between two different types of return:

1. The unconditional return is indeed zero, as seen from (1), and reflects the fair game condition.

2. The conditional return, conditioned upon no crash occurring between time t and time t', is nonzero and is given by equations (3) and (7), respectively. If the crash hazard rate is increasing with time, the conditional return will be accelerating precisely because the crash becomes more probable and the investors need to be remunerated for their higher risk.

Thus, the expectation which remains constant in equation (1) takes into account the probability that the market *may* crash. Therefore, *conditionally* on staying in the bubble (no crash yet), the market must rationally rise to compensate buyers for having taken the risk that the market *could* have crashed.

The market price reflects the equilibrium between the greed of buyers who hope the bubble will inflate and the fear of sellers that it may crash. A bubble that goes up is just one that could have crashed but did not. Our model is well summarized by borrowing the words of another economist: "...the higher probability of a crash leads to an acceleration of [the market price] while the bubble lasts." Interestingly, this citation is culled from the very same article by Blanchard (1979) that Ilinski (1999) cites as an alternative model that is more realistic than ours. We see that this is in fact more of an endorsement than an alternative.

To go into detail, Blanchard's (1979) model is slightly more general because it incorporates risk aversion through

$$\nu \mathsf{E}_t[p(t')] = p(t), \tag{9}$$

where $\nu \in (0, 1]$ is an appropriate discount factor. We have been aware of this since the beginning, and the only reason why we did not take it into account is that it obviously makes *no* difference in our substantive predictions (log-periodic oscillations and power law acceleration) as long as ν remains bounded away from zero and infinity.

Another way to incorporate risk aversion into our model is to say that the

probability of a crash in the next instant is perceived by traders as being K times bigger than it objectively is. This amounts to multiplying our hazard rate $h(t)$ by K, and once again this makes no substantive difference as long as K is bounded away from zero and infinity.

The point here is that ν and K both represent a general aversion of fixed magnitude against a risk. Risk aversion is a central feature of economic theory, and it is generally thought to be stable within a reasonable range associated with slow-moving secular trends such as changes in education, social structures, and technology.

Ilinski (1999) correctly points out that risk perceptions are constantly changing in the course of real-life bubbles, but wrongly claims that our model violates this intuition. In our model, risk perceptions do oscillate dramatically throughout the bubble, even though subjective aversion to risk remains stable, simply because it is the *objective degree of risk that the bubble may burst* that goes through wild swings.

For these reasons, the criticisms put forward by Ilinski, far from making a dent in our economic model, serve instead to show that it is robust, flexible, and intuitive.

2.2 Crash Hazard Rate and the Critical Imitation Model

The crash hazard rate $h(t)$ quantifies the probability that a large group of agents place sell orders simultaneously and create enough of an imbalance in the order book for market makers to be unable to absorb the other side without lowering prices substantially. Most of the time, market 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 to determine by what mechanism did they suddenly manage to organize a coordinated sell-off?

We have proposed the following answer (Johansen *et al.* 2000). All the traders in the world are organized into a network (of family, friends, colleagues, etc.) and they influence each other *locally* through this network. For instance, an active trader is constantly on the phone exchanging information and opinions with a set of selected colleagues. In addition, there are indirect interactions mediated, for instance, by different parts of the media. Specifically, if I am directly connected with k other traders, then there are only two forces that influence my opinion: (a) the opinions of these k people and of the global information network and (b) an idiosyncratic signal that I alone generate. Our working hypothesis here is that agents tend to *imitate* the opinions of their connections. The force (a) will tend to create order, while force (b) will tend to create disorder. The main issue here is the conflict 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). This mechanism does not require an overarching coordination mechanism since macro-level coordination can arise from micro-level imitation and it relies on a realistic model of how agents form opinions by constant interaction.

Many models of interaction and imitation between traders have been developed. To cut a long story short, the upshot is that the conflict between order and disorder often leads to a regime where order may win. When this occurs, the bubble ends. Many models (but not all) undergo this transition in a 'critical' manner (Goldenfeld 1992, Dubrulle *et al*. 1997). The sensitivity of the market reaction to news or external influences accelerate on the approach to this transition in a specific way characterized by a power law divergence at the critical time t_c of the form $F(t) = (t_c - t)^{-z}$, where z is called a critical exponent. This form amounts to the property that

$$\frac{d \ln F}{d \ln(t_c - t)} = -z \tag{10}$$

is a constant, namely that the behavior of the observable F becomes self-similar close to t_c with respect to dilation of the distance $t_c - t$ to the critical point at t_c. The symmetry of self-similarity in the present context refers to the fact that the relative variations $d \ln F = dF/F$ of the observable with respect to relative variations $d \ln(t_c - t) = d(t_c - t)/(t_c - t)$ of the time-to-crash are independent of time t, as expressed by the constancy of the exponent z.

Accordingly, the crash hazard rate follows a similar dependence, namely

$$h(t) = B(t_c - t)^{-b}, \tag{11}$$

where B is a positive constant and t_c is the critical point or the theoretical date of the bubble death. The exponent b must lie between 0 and 1 for an important economic reason: otherwise, the price would go to infinity when approaching t_c (if the bubble has not burst yet).

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. The time t_c is the most probable time of the crash. There exists a residual probability

$$1 - \int_{t_0}^{t_c} h(t)\, dt > 0 \tag{12}$$

of attaining the critical date without a crash. This residual probability is crucial for the coherence of the theory: otherwise, the whole model would unravel because rational agents would anticipate the crash.

Plugging equation (11) into equation (4) gives the following law for price:

$$p(t) \approx p_c - \frac{\kappa B}{\beta}(t_c - t)^{\beta} \quad \text{before the crash,} \tag{13}$$

where $\beta = 1 - b \in (0, 1)$ and p_c is the price at the critical time t_c (conditioned on no crash having been triggered). We see that the price before the crash also follows a power law with a finite upper bound p_c. The slope of the price, which is the expected return per unit of time, becomes unbounded as we approach the

critical date t_c. This is to compensate for an unbounded hazard rate approaching t_c.

In the contemporary information environment, one could argue that the media can play a more efficient role in the creation of an opinion than the network of 'friends' and will thus modify if not destroy the critical nature of the crash resulting from the increasing cooperativeness between the network of 'friends'. There are two ways to address this issue. First, consider the situation of a typical trader. Since all her competitors have access to the same globally shared information, she does not feel that this gives her an edge, but only serves to keep her abreast of the market moves. The information that she considers really valuable are those that she can share and exchange with a selected and narrow circle of trusted colleagues. This argument brings us back to the model of an imitation network with 'local' connections. A second argument is that the media does nothing but reflect a kind of consensus resulting precisely from the collective action of all the participants in the market and hence only mirrors the interactions between the traders. In addition, this revelation of the general consensus and of the global market sentiment reinforces the effective interaction between the traders. In the language of the statistical physics of critical points, the media can thus be compared with an effective mean field created endogenously and may thus effectively strengthen the existence of a critical point as described here.

2.3 Log-periodicity

The power law dependence (11) of the hazard rate is the hallmark of self-similarity of the market across scales as the critical time t_c is approached. 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. This critical state occurs 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 become highly correlated.

As we have said, scale invariance of a system near its critical point must be represented by power law dependences of the observables (Dubrulle *et al.* 1977). Formally, these power laws are the solution of the renormalization group equations (Wilson 1979, Goldenfeld 1992) that describe the correlations between agents at many scales. It turns out that the most general mathematical solutions of these renormalization group equations are power laws with *complex* exponents which as a consequence exhibit log-periodic corrections to scaling (Sornette 1998).

A straightforward mechanism for these complex exponents to appear is to define the dynamics on a hierarchical structure displaying a discrete scale invariance. Schematically, we can think of the stock market made up of

participants that differ in size by many orders of magnitude, ranging from individuals to gigantic professional investors, such as pension funds. Furthermore, structures at even higher levels, such as currency influence spheres (US dollar, euro, yen, . . .), exist and, with the current globalization and deregulation 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 has features which resemble that of hierarchical systems 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. Models (Johansen *et al.* 2000, Sornette and Johansen 1998) of imitative interactions on hierarchical structures recover the power law behavior (13) and predict that the critical exponents b and β may be complex numbers!

Recently, it has been realized that discrete scale invariance and its associated complex exponents may appear 'spontaneously' in nonhierarchical systems, i.e. without the need for a pre-existing hierarchy (see Sornette (1998) for a review). There are many dynamical mechanisms that can operate to produce these complex exponents, such as nonlocal geometry, frozen heterogeneity (Saleur and Sornette 1996), cascades of instabilities, intermittent amplifications (see Johansen and Sornette (1999c) for such a mechanism in the context of 'antibubbles' proposed recently for the Japanese stock market), and so on.

As a consequence, we are led to generalize (11) and give the first-order expansion of the general solution for the hazard rate:

$$h(t) \approx B(t_{c} - t)^{-b} + C(t_{c} - t)^{-b} \cos[\omega_1 \log(t_{c} - t) + \psi]. \tag{14}$$

The crash hazard rate now displays log-periodic oscillations. This can easily be seen to be the generalization of equation (10) by taking the exponent z to be complex with a nonzero imaginary part, since the real part of $(t_{c} - t)^{-z+i\omega}$ is $(t_{c} - t)^{-z} \cos[\omega \ln(t_{c} - t)]$. The evolution of the price before the crash and before the critical date is then given by

$$p(t) \approx A_1 + B_1(t_{c} - t)^{\beta} + C_1(t_{c} - t)^{\beta} \cos[\omega_1 \log(t_{c} - t) + \phi_1], \tag{15}$$

where ϕ_1 is another phase constant. The key feature is that accelerating oscillations appear in the price of the asset 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 such that the ratio of consecutive time intervals is a constant:

$$\lambda \equiv e^{2\pi/\omega_1}. \tag{16}$$

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}.

We note that these log-periodic structures bear some similarity to patterns classified empirically by chartists and others using methods of technical analysis,

such as 'Elliott waves' (Frost and Prechter 1998) and 'log-spirals' (Erman 1999). For instance, a logarithmic spiral in the plane obeys the equation $r = r_0 e^{a\theta}$ in polar coordinates. The intersections with the x-axis occur for $\theta = n\pi$ (recall that $x = r\cos\theta$) and the corresponding coordinates are $x_n = (-1)^n e^{a\pi n}$. The intersections with the y-axis occur for $\theta = \frac{1}{2}\pi + n\pi$ (recall that $y = r\sin\theta$) and the corresponding coordinates are $y_n = (-1)e^{a\pi/2}e^{a\pi n}$. Both series indeed form discrete geometrical series with a scaling ratio $\lambda = e^{a\pi}$ similar to (16). We do not need, however, to stress the difference between a pattern recognition approach (Erman 1999) and our economically motivated approach incorporating a rational model of imitative agents as well as well-established tools from statistical physics.

3. EMPIRICAL RESULTS

3.1 *Are Large Crashes Special?*

Crashes are extreme events. There are two possible ways to describe them:

1. The distribution of returns is stationary and the extreme events can be extrapolated as lying in its far tail. Within this framework, recent works in finance and insurance have recently investigated the relevance of the body of theory known as 'Extreme Value Theory' to extreme events and crashes (Embrechts *et al.* 1997, 1998, 1999).

2. Crashes cannot be accounted for by an extrapolation of the distribution of smaller events to the regime of extremes and belong intrinsically to another regime, another distribution, and are thus outliers.

To test which one of these two alternatives provides the most accurate description, a statistical analysis of market fluctuations (Johansen and Sornette 1998) has provided significant indications that large crashes may be outliers. Indeed, it was established that the distribution of draw-downs of the Dow Jones average daily closing is well described by an exponential distribution with a decay constant of about 2%. This exponential distribution holds only for draw-downs smaller than about 15%. In other words, this means that all draw-downs of amplitudes of up to approximately 15% are well described by the same exponential distribution with characteristic scale 2%. This characteristic decay constant means that the probability of observing a draw-down larger than 2% is about 37%. Following hypothesis 1 and extrapolating this description to, for example, the three largest crashes on Wall Street this century (1914, 1929, and 1987) yields a recurrence time of about 50 centuries for *each* single crash. In reality, the three crashes occurred in less than one century. This suggests that hypothesis 2 is preferable. In the following, we use the generic term 'crash' to refer to the significant drop of a market occurring over a few days to a few weeks which follows an extended period of bullish behavior and signifies the end of the bubble.

As an additional null-hypothesis, 10 000 synthetic data sets, each covering a time-span close to a century and hence adding up to about 10^6 years, have been generated using a GARCH(1,1) model estimated from the true index with a Student t-distribution with four degrees of freedom (Bollerslev *et al*. 1992). This model is often used in academic and practitioner circles as a reasonable description of the market and includes both the nonstationarity of volatilities and the fat tail nature of the price returns. Our analysis (Johansen *et al*. 2000) showed that only two data sets had three draw-downs above 22% and none had four. However, three of these six 'crashes' were preceded by a draw-up of comparable size of the draw-down and hence showed an abnormal behavior not found for real crashes. This means that in approximately 1 million years of 'Garch-trading', with a reset every century, *never* did three crashes occur in a single century.

This suggests that different mechanisms may be responsible for large crashes and that hypothesis 2 is the correct description of crashes. This also raises the issue whether precursory patterns exists decorating the speculative bubble preceding the crash. The point is simply that while the GARCH(1,1) model does a reasonable job of reproducing fluctuations in 'normal' trading, it cannot capture the fluctuations connected with large crashes, and hence another type of model is necessary for this special behavior. Of course, these simulations do not prove that our model is the correct one, only that one of the standard models of the 'industry' (which makes a reasonable null hypothesis) is utterly unable to account for the stylized facts associated with large financial crashes. A better model is thus called for and this is exactly the goal of the present paper to provide for such a framework.

3.2 Log-periodic Analysis of Large Crashes

The numerical procedure of fitting equation (15), as well as equations (19), (20), and (21), is a minimization of the variance

$$Var = \frac{1}{N} \sum_{i=1}^{N} [y_i - f(t_i)]^2 \qquad (17)$$

between the N data points y_i and the fit function $f(t)$ using the downhill simplex algorithm of Press *et al*. (1994). Tests using the maximum likelihood method with the Student t-distribution show that our results are not significantly sensitive to the estimation procedure (Lin 1998). In order to reduce the complexity of the fit, all linear variables A, B, C are determined from the nonlinear variables by requiring that a solution has zero derivative of the variance as a function of the linear variables. This means that equation (15), is an effective four-parameter fit with β, ϕ, t_c, ω. Further details on our numerical procedure are given by Johansen and co-workers (Johansen *et al*. 2000, Johansen 1997).

A remaining question concerns what data interval prior to the crash to fit. The procedure used was the following. The last point used for all crashes was the highest value of the price before the crash; the first point used was the lowest

value of the price when the bubble started. For all the cases discussed here, except the Russian antibubble analyzed in Section 3.5,[1] this procedure always gave a convincing result and it was never necessary to change the interval fitted.

Figures 1–4 show the behavior of the market index prior to the four stock market crashes of October 1929 (Wall Street), October 1987 (Wall Street), October 1997 (Hong Kong), and August 1998 (Wall Street), as well as the collapse of the US dollar against the Deutschmark and Swiss franc in 1985 and against the Canadian dollar and the yen in 1998. A fit with equation (15) is shown as a continuous line for each event. Table 1 lists the parameters of the fit to the data. Note the small fluctuations in the value of the scaling ratio $2.2 \leqslant \lambda \leqslant 2.8$ for all data sets except the Swiss franc. This agreement cannot be accidental and constitutes one of the key tests of our framework.

In order to qualify further the significance of the log-periodic oscillations in a nonparametric way, as well as providing an independent test on the value of the frequency of the log-periodic oscillations and hence the preferred scaling ratio λ, we have eliminated the leading trend from the price data by the transformation

$$p(t) \rightarrow \frac{p(t) - A_1 - B_1(t_c - t)^\beta}{C_1(t_c - t)^\beta}. \tag{18}$$

This transformation should produce a pure $\cos[\omega_1 \log(t_c - t) + \phi_1]$ *if* equation (15) was a perfect description. In Figure 5, we show the residual for the bubble prior to the 1987 crash. We see very convincing periodic behavior as a function of $\log[(t_c - t)/t_c]$. In Figure 6, we see the corresponding data as a function of the true time to the crash, $t_c - t$. Looking at Figures 5 and 6, two things especially strike the eye. First, we see that in the latter figure the distance

TABLE 1. Time t_c is the critical time predicted from the fit of the market index to the equation (15). The other parameters of the fit are also shown. The error *Var* is the variance between the data and the fit (equation (17)) and has units *price*^{2}. Each fit is performed up to the time t_{max} at which the market index achieved its highest maximum before the crash. Time t_{min} is the time of the lowest point of the market disregarding smaller 'plateaux'. The percentage drop is calculated from the total loss from t_{max} to t_{min}.

Crash	t_c	t_{max}	t_{min}	Drop (%)	β	ω_1	λ	A	B	C	*Var*
1929 (WS)	30.22	29.65	29.87	47	0.45	7.9	2.2	571	−267	14.3	56
1985 (DEM)	85.20	85.15	85.30	14	0.28	6.0	2.8	3.88	1.16	−0.77	0.0028
1985 (CHF)	85.19	85.18	85.30	15	0.36	5.2	3.4	3.10	−0.86	−0.055	0.0012
1987 (WS)	87.74	87.65	87.80	30	0.33	7.4	2.3	411	−165	12.2	36
1997 (HK)	97.74	97.60	97.82	46	0.34	7.5	2.3	20077	−8241	−397	190360
1998 (WS)	98.72	98.55	98.67	19	0.60	6.4	2.7	1321	−402	19.7	375
1998 (YEN)	98.78	98.61	98.77	21	0.19	7.2	2.4	207	−84.5	2.78	17
1998 (CAN$)	98.66	98.66	98.71	5.1	0.26	8.2	2.2	1.62	−0.23	−0.011	0.00024

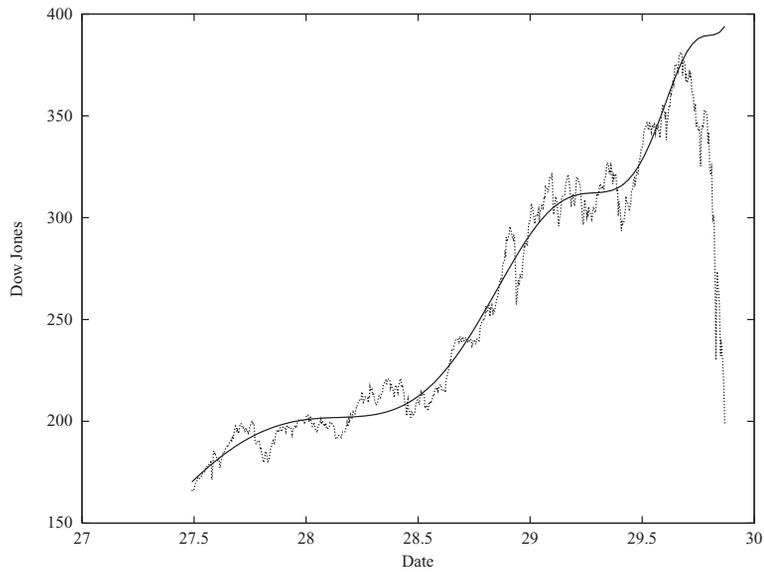[1] Here 'lowest' and 'highest' are, of course, interchanged.

FIGURE 1. The Dow Jones index prior to the October 1929 crash on Wall Street. The fit is equation (15) with $A_1 \approx 571$, $B_1 \approx -267$, $C_1 \approx 14.3$, $\beta \approx 0.45$, $t_c \approx 30.22$, $\phi_1 \approx 1.0$, $\omega_1 \approx 7.9$.
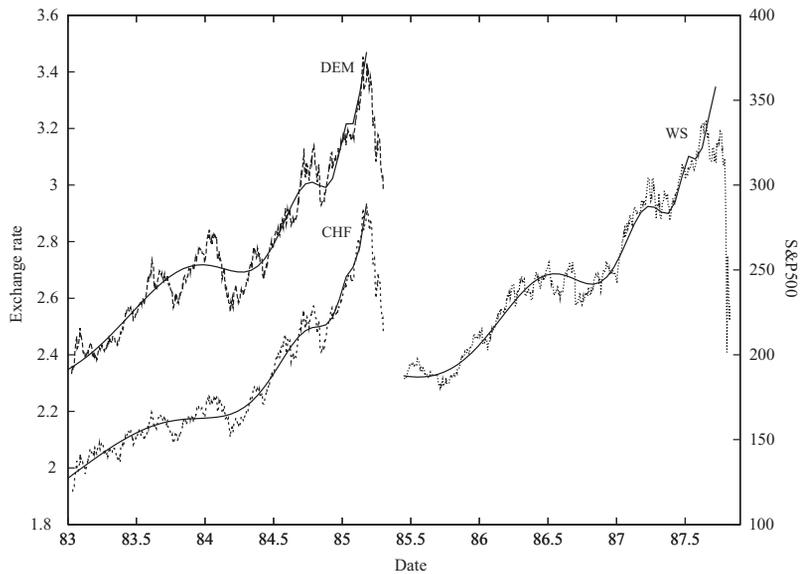

FIGURE 2. The S&P500 US index prior to the October 1987 crash on Wall Street and the US dollar against the Deutschmark and Swiss franc prior to the collapse in mid-85. The fit to the S&P500 is equation (15) with $A_1 \approx 412$, $B_1 \approx -165$, $C_1 \approx 12.2$, $\beta \approx 0.33$, $t_c \approx 87.74$, $\phi_1 \approx 2.0$, $\omega_1 \approx 7.4$. The fits to the Deutschmark and Swiss franc currencies against the US dollar gives $A_1 \approx 3.88$, $B_1 \approx -1.2$, $C_1 \approx 0.08$, $\beta \approx 0.28$, $t_c \approx 85.20$, $\phi_1 \approx -1.2$, $\omega_1 \approx 6.0$ and $A_1 \approx 3.1$, $B_1 \approx -0.86$, $C_1 \approx 0.05$, $\beta \approx 0.36$, $t_c \approx 85.19$, $\phi_1 \approx -0.59$, $\omega_1 \approx 5.2$, respectively.
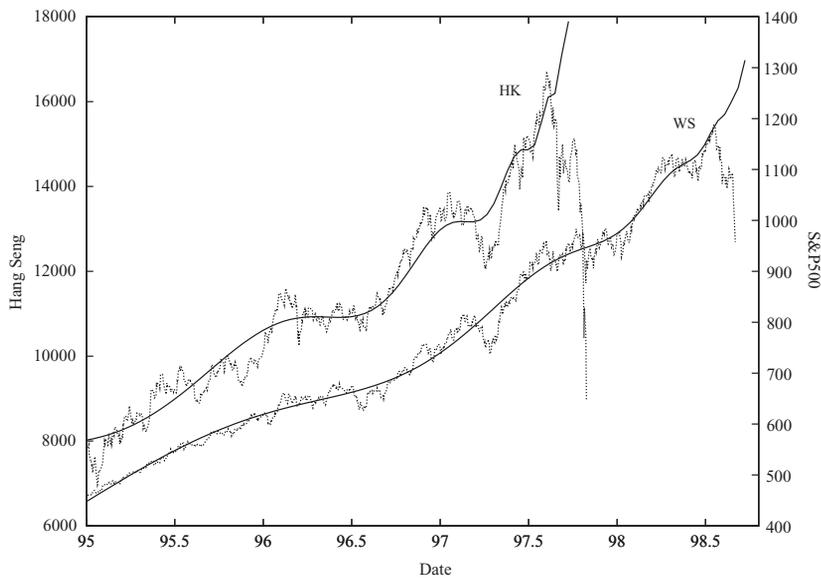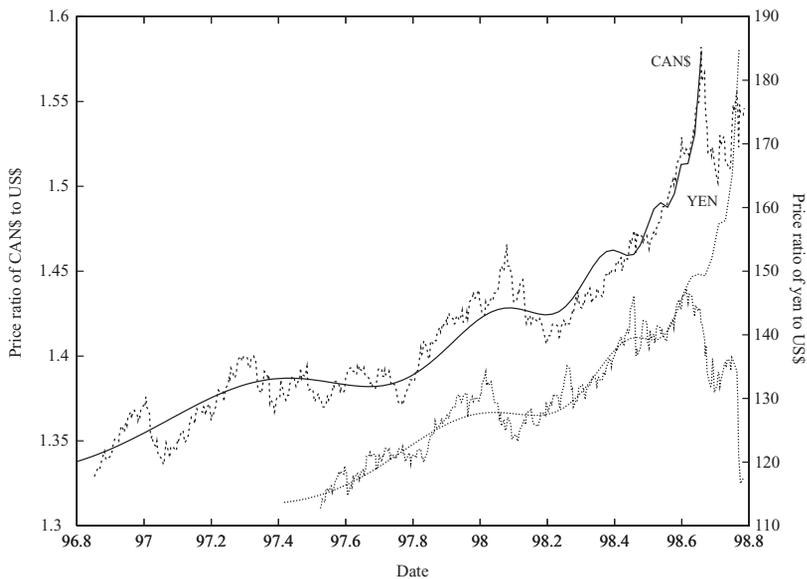
FIGURE 3. The Hang Seng index prior to the October 1997 crash on the Hong Kong stock exchange and the S&P500 stock market index prior to the recent crash on Wall Street in August 1998. The fit to the Hang Seng index is equation (15) with $A_1 \approx 20077$, $B_1 \approx -8241$, $(\kappa/\beta)C_1 \approx -397$, $\beta \approx 0.34$, $t_c \approx 97.74$, $\phi_1 \approx 0.78$, $\omega_1 \approx 7.5$. The fit to the S&P500 index is equation (15) with $A_1 \approx 1321$, $B_1 \approx -402$, $C_1 \approx 19.7$, $\beta \approx 0.60$, $t_c \approx 98.72$, $\phi_1 \approx 0.75$, $\omega_1 \approx 6.4$.

FIGURE 4. The Canadian dollar and yen currencies against the US dollar prior to the drop starting in August 1998. The fit with equation (15) to the two exchange rates gives $A_1 \approx 1.62$, $B_1 \approx -0.22$, $C_1 \approx -0.011$, $\beta \approx 0.26$, $t_c \approx 98.66$, $\phi_1 \approx -0.79$, $\omega_1 \approx 8.2$ and $A_1 \approx 207$, $B_1 \approx -85$, $C_1 \approx 2.8$, $\beta \approx 0.19$, $t_c \approx 98.78$, $\phi_1 \approx -1.4$, $\omega_1 \approx 7.2$, respectively.
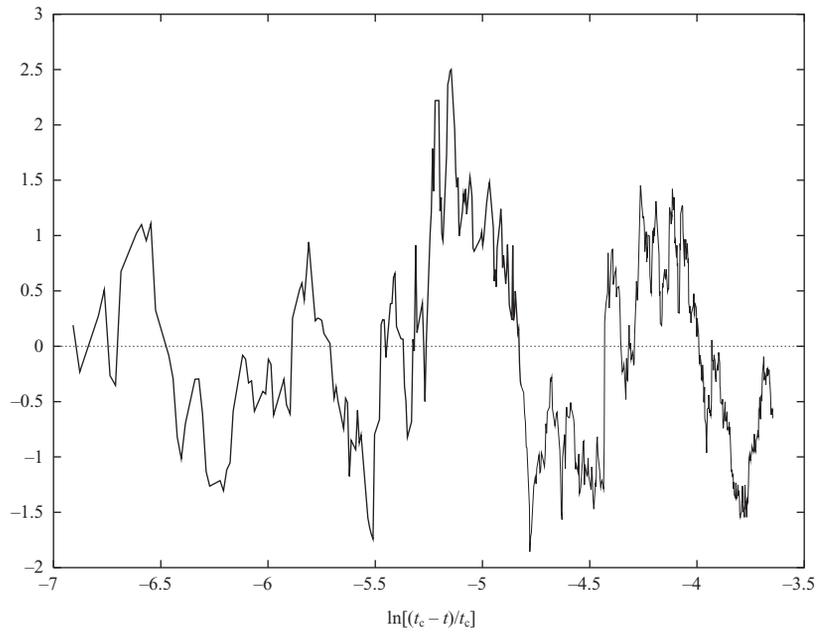
FIGURE 5. The residual as defined by the transformation (18) as a function of $\log[(t_c - t)/t_c]$ for the 1987 crash.
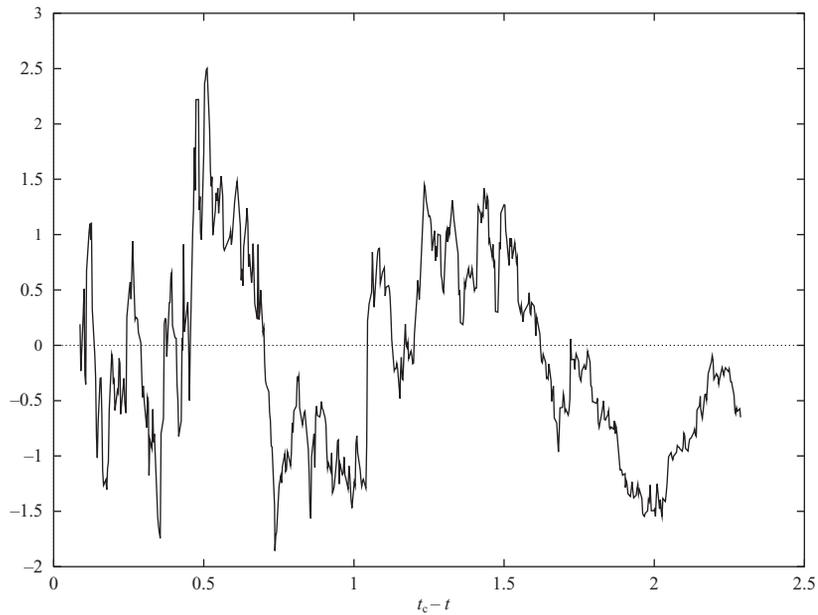

FIGURE 6. The residual as defined by the transformation (18) as a function of $t_c - t$ for the 1987 crash.
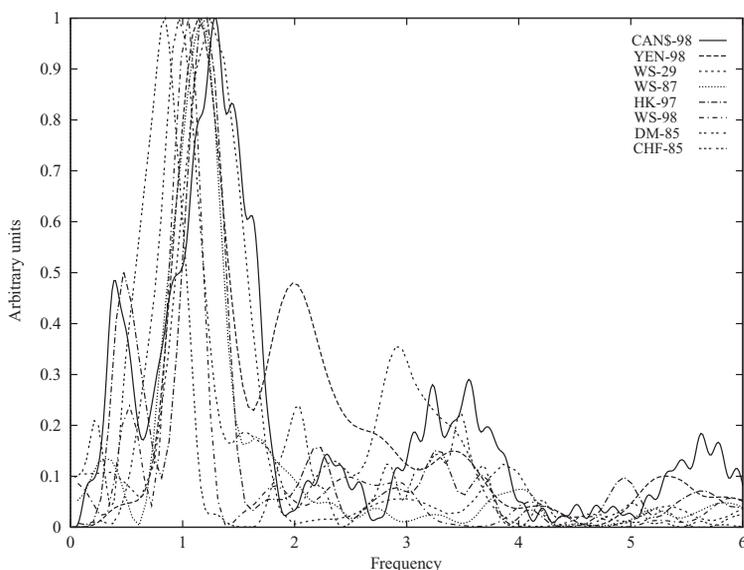
FIGURE 7. The Lomb periodogram for the 1929, 1987, and 1998 crashes on Wall Street, the 1997 crash on the Hong Kong stock exchange, the US dollar currency crash in 1985 against the Deutschmark and Swiss franc and in 1998 against the yen, and the 5.1% correction against the Canadian dollar. For each periodogram, the significance of the peak should be estimated against the noise level.

between consecutive extrema of the data is clearly accelerating; in the former we clearly see that a transformation to the logarithm of the time to the crash exhibits an overall periodic trend, hence demonstrating log-periodic behavior. Second, since the detrended data oscillates nicely around zero, the two figures clearly show that the power law in equation (15) is an excellent approximation to the overall rise in the data. That the residual oscillates beyond ± 1 only indicates that the cosine used in equation (15) is not an optimal choice for a periodic function in the present context. However, it has the advantage of being familiar to the majority, as well as straightforward to implement numerically.

The significance of this oscillating trend has been further estimated by performing a spectral analysis, designed to test for the statistical significance of the log-periodic oscillations. To obtain this, we use a so-called Lomb periodogram (Press *et al.* 1994) rather than the more standard power spectrum, since the data is not equidistantly sampled in time. The Lomb periodogram is a local fit of a cosine (with a phase) using some user-chosen range of frequencies. In Figure 7, we see a peak around $f \approx 1.1$ for all eight cases shown in Figures 1–4, corresponding to $\omega_1 = 2\pi f \approx 7$, in perfect agreement with the results listed in Table 1. Since the 'noise' spectrum is unknown and very likely different for each crash, we cannot estimate the confidence interval of the peak in the usual manner (Press *et al.* 1994) and compare the results for the different crashes. Therefore, only the relative level of the peak *for each separate* periodogram can be taken as a measure of the significance of the oscillations, and the periodograms have hence been normalized. We note that *if* the noise were to have a

white Gaussian distribution, the confidence given by the periodograms increases exponentially with the value of the abscissa and would be well above 99.99% for all cases shown (Press *et al*. 1994). Note also that the strength of the oscillations is ∼5–10% of the leading power law behavior for all eight cases signifying that they cannot be neglected.

To summarize thus far, our spectral analysis demonstrates that the observed log-periodic oscillations have a very strong power spectrum, much above noise level. It would be very difficult and much less parsimonious to account for these structures by another model.

3.3 *Are Log-periodic Signatures Statistically Significant?*

In the case of the 1929 and 1987 crashes on Wall Street, log-periodic oscillations in the index could be identified as long as ∼7.5 years prior to the crashes (Sornette and Johansen 1997). In order to investigate the significance of these results, we picked at random fifty 400-week intervals in the period 1910 to 1996 of the Dow Jones average and launched the same fitting procedure as on the time period prior to the 1929 and 1987 crashes on Wall Street. The results were encouraging. Of 11 fits with a quality (specifically, the variance of the fit of equation (20) with $\tau = t_c - t$ to the data) comparable with that of the two crashes, only six data sets produced values for β, ω_1, ω_2, and T_1, as defined by equation (20), which were in the same range. However, all of these fits belonged to the periods prior to the crashes of 1929, 1962, and 1987. The existence of a 'crash' in 1962 was, prior to these tests, unknown to us, and the identification of this crash naturally strengthens the case. We refer the reader to Johansen *et al*. (2000) for a presentation of the best fit obtained for this 'crash'.

We also generated 1000 synthetic data sets of length 400 weeks using the same GARCH(1,1) model used in Section 3.1. Each of these 1000 data sets was analyzed in the same manner as the real crashes. In 66 cases, the best minima had values similar to the real crashes, but all of these fits, except two, did not resemble the true stock market index prior to the 1929 and 1987 crashes on Wall Street, the reason primarily being that they only contained one or two oscillations. However, two fits looked rather much like that of the 1929 and 1987 crashes (Johansen *et al*. 2000, Johansen 1997).

This result corresponds approximately to the usual 95% confidence interval for *one* event. In contrast, we have here provided five examples of log-periodic signatures before large stock market crashes. We are aware that this is not proof of log-periodic signatures in bubbly markets (only a thorough data analysis can establish this without reasonable doubt), but it provides a very good indication that these signatures are not easily generated accidentally.

3.4 *Log-periodicity in Bearish Markets?*

Stock market jargon divides stock market trends into either 'bullish' or 'bearish'. If log-periodic signals, apparently, are present in bullish markets, the obvious question to ask is whether this is also the case in bearish markets.
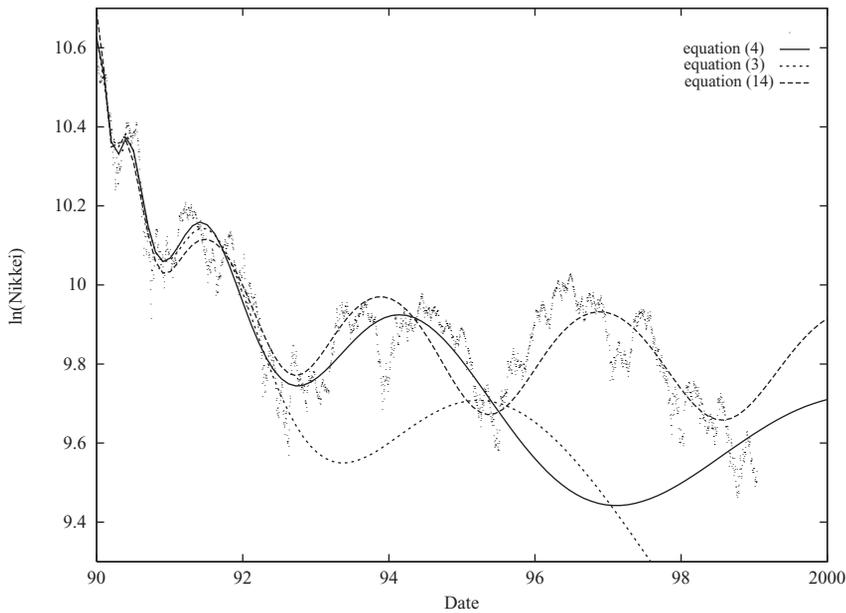
FIGURE 8. The natural logarithm of the Nikkei stock market index after the start of the decline on 1 January 1990 until 31 December 1998. The lines are equation (19) (dotted line) fitted over an interval of $\sim$2.6 years, equation (20) (continuous line) over $\sim$5.5 years, and equation (21) (dashed line) over 9 years. The parameter values of the first fit of the Nikkei are $A_1 \approx 10.7$, $B_1 \approx -0.54$, $C_1 \approx -0.11$, $\beta \approx 0.47$, $t_c \approx 89.99$, $\phi_1 \approx -0.86$, $\omega_1 \approx 4.9$ for equation (19). The parameter values of the second fit of the Nikkei are $A_2 \approx 10.8$, $B_2 \approx -.70$, $C_2 \approx -0.11$, $\beta \approx 0.41$, $t_c \approx 89.97$, $\phi_2 \approx 0.14$, $\omega_1 \approx 4.8$, $T_1 \approx 9.5$, $\omega_2 \approx 4.9$ for equation (20). The third fit uses the entire time interval and is performed by adjusting only T_1, T_2, ω_2, and ω_3, while β, t_c, and ω_1 are fixed at the values obtained from the previous fit. The values obtained for these four parameters are $T_1 \approx 4.3$ years, $T_2 \approx 7.8$ years, $\omega_2 \approx -3.1$, and $\omega_3 \approx 23$. Note that the values obtained for the two time scales T_1 and T_2 confirm their ranking. This last fit predicts a change of regime and that the Nikkei should increase in 1999.

The most recent example of a genuine long-term depression comes from Japan, where the Nikkei has decreased by more than 60% in the 9 years following the all-time high of 31 December 1989. In Figure 8, we see (the logarithm of) the Nikkei index from 31 December 1989 until 31 December 1998. The three fits are equations (19), (20), and (21), respectively (Johansen and Sornette 1999c):

$$\log p(t) \approx A_1 + B_1\tau^\beta + C_1\tau^\beta \cos[\omega_1 \log \tau + \phi_1], \tag{19}$$

$$\log p(t) \approx A_2 + \frac{\tau^\beta}{\sqrt{1 + (\tau/T_1)^{2\beta}}}$$

$$\times \left[B_2 + C_2 \cos\left(\omega_1 \log \tau + \frac{\omega_2}{2\beta}\log[1 + (\tau/T_1)^{2\beta}] + \phi_2\right)\right], \tag{20}$$

$$\log p(t) \approx A_3 + \frac{\tau^{\beta}}{\sqrt{1 + (\tau/T_1)^{2\beta} + (\tau/T_2)^{4\beta}}}$$

$$\times \left[B_3 + C_3 \cos\left(\omega_1 \log \tau + \frac{\omega_2}{2\beta} \log[1 + (\tau/T_1)^{2\beta}] + \frac{\omega_3}{4\beta} \log[1 + (\tau/T_2)^{4\beta}] + \phi_3 \right) \right],$$

(21)

where $\tau = t - t_c$. Equation (21) predicts the transition from the log-frequency ω_1 close to t_c to $\omega_1 + \omega_2$ for $T_1 < \tau < T_2$ and to the log-frequency $\omega_1 + \omega_2 + \omega_3$ for $T_2 < \tau$. The equations (20) (resp. (21)) extend the renormalization group approach to the second (resp. third) order of perturbation (Johansen and Sornette 1999c, Sornette and Johansen 1997).

The interval used for equation (19) is until mid-1992 and for equation (20) until mid-1995. The parameter values produced by the different fits with equations (19) and (20) agree remarkably well and the values for the exponent β also agree well, as shown in Table 2. The fit with (21) was, owing to the large number of free variables, performed differently and the parameter values for t_c, β, and ω_1 determined from the fit with (20) was kept fixed and only T_1, T_2, ω_2, ω_3, and ϕ_3 were allowed to be adjusted freely. What lends credibility to the fit with equation (21) is that, despite its complex form, we get values for the two cross-over time scales T_1 and T_2 which correspond very nicely to what is expected from the theory: T_1 has moved down to $\sim$4.4 years in agreement with the time interval used for equation (20) and T_2 is $\sim$7.8 years, which is compatible with the 9-year interval of the fit. This does not mean that fitting is not very

TABLE 2. Time t_c is the predicted time of the crash from the fit of the market index to the equation (15). The other parameters of the fits to the preceding bubble are also given. The error *Var* is the variance between the data and the fit (equation (17)) and has units *price*^{2} except for the Nikkei, where the units are [log(*price*)]^2. The fit to the bubble is performed up to the time at which the market index achieved its highest maximum before the crash. The parameters t_c, β, ω_1, and λ correspond to the fit with equation (15), where t_c and t have been interchanged. Here t_{max} and t_{min} represent the endpoints of the interval fitted.

Bubble	t_c	t_{max}	t_{min}	Drop (%)	β	ω_1	λ	A	B	C	*Var*
ASPMT	97.61	97.61	97.67	17	0.37	7.5	2.3	1280	−1025	59.5	907
IRTS	97.61	97.61	97.67	17	0.39	7.6	2.3	633	−483	38.8	310
ROSI	97.61	97.61	97.67	20	0.40	7.7	2.3	4254	−3166	246	12437
ASPGEN	97.62	97.60	97.67	8.9	0.25	8.0	2.2	2715	−2321	72.1	1940
Antibubble	t_c	t_{max}	t_{min}	Drop (%)	β	ω_1	λ	A	B	C	*Var*
ROSI	97.72	97.77	98.52	74	0.32	7.9	2.2	4922	−3449	472	59891
Nikkei											
eqn. (19)	89.99	90.00	92.63	63	0.47	4.9	3.6	10.7	−0.54	−0.11	0.0029
eqn. (20)	89.97	90.00	95.51	63	0.41	4.8	3.7	10.8	−0.70	−0.11	0.0600

degenerate; it is—but the ranking of T_1 and T_2 is always the same and the values given do not deviate much from those in the caption of Figure 8, i.e. by ± 1 year.

The value $\omega_1 \approx 4.9$ correspond to $\lambda \approx 3.6$, which is not in the range found for the bubbles (see Table 1). An additional difference between the antibubble in the Nikkei and these other cases is the strength of the oscillations compared with the leading behavior. For the Nikkei, it is $\sim 20\%$, i.e. 2–4 times as large as the amplitude obtained for the stock market and Forex bubbles.

3.5 Log-periodic Analysis of the Russian Stock Market

Intrigued by the claims of Ilinski (1999) of a genuine stock market bubble, followed by a crash, *without* log-periodic signatures on the Russian stock market, we have analyzed four Russian stock market indices from their start through to the present. The four indices were the Russian Trading System Interfax Index (IRTS), the Agence Skate Press Moscow Times Index (ASPMT), the Agence Skate Press General Index (ASPGEN), and the Credit Suisse First Boston Russia Index (ROSI). The ROSI is generally considered the best of the four and we will put an emphasis on the results obtained with that index. The Russian stock market is highly volatile (a fact that we will return to later), which means that great care must be taken in maintaining a representative stock market index. This is the primary reason for using four indices in the analysis.

In Figure 9, we see the ROSI fitted with equation (15) in the interval [96.21 : 97.61]. The interval is chosen by identifying the start of the bubble and the end represented by the date of the highest value of the index before the crash similarly to the seven major market crashes discussed previously. For all four indices, the same start-day and end-day could be identified ± 1 day.

In Figures 11 and 12, we see the detrended data using the transformation (18). The conclusions drawn in relation to the corresponding figures for the 1987 Wall Street crash (Figures 5 and 6) are again valid, supporting both an overall power law rise as well as log-periodic signatures for the Russian crash of 1997.

As can be seen from Table 2, the dimensionless parameters β, ω_1, and λ, as well as the predicted time of the crash t_c, for the fit to the different indices agree very well, except for the exponent β obtained from the ASPGEN index. In fact, the value obtained for the preferred scaling ratio λ fluctuates by no more than 5% for the four fits, thereby showing numerical stability.

The origin of this bubble is well known. In 1996 large international investors (read US, German, and Japanese) began to invest heavily in the Russian markets believing that the financial situation of Russia had finally stabilized. Nothing could have been further from the truth (Intriligator 1998, Malki 1999), but the belief and hope in a new investment haven with large returns led to herding and bubble development. This means that the same herding which created the log-periodic bubbles on Wall Street (1929, 1987, 1998), Hong Kong (1997), and the Forex (1985, 1998) entered an emerging market and brought along the same log-periodic trends characterizing the global markets. The fact that the consistent values of λ obtained for the four indices of the Russian market are comparable with those obtained for the Wall Street, Hong Kong,
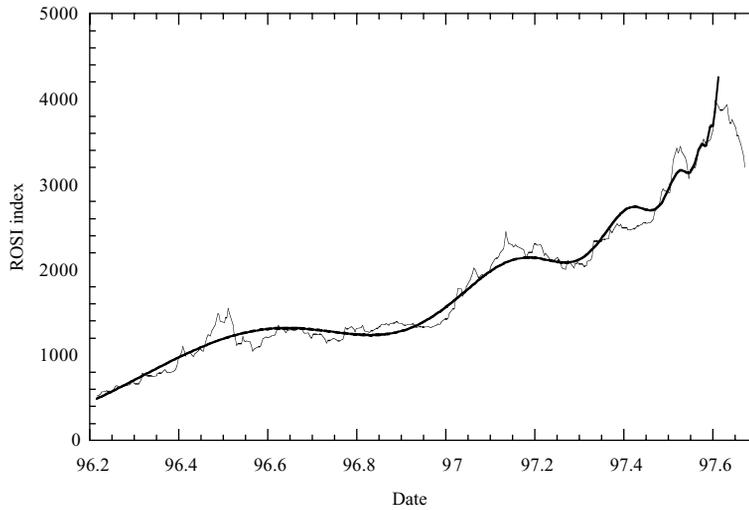
FIGURE 9. The ROSI index fitted with equation (15) over the interval shown. The parameter values of the fit with equation (15) are $A_1 \approx 4254$, $B_1 \approx -3166$, $C_1 \approx 246$, $\beta \approx 0.40$, $t_c \approx 97.61$, $\phi_1 \approx 0.44$, $\omega_1 \approx 7.7$.
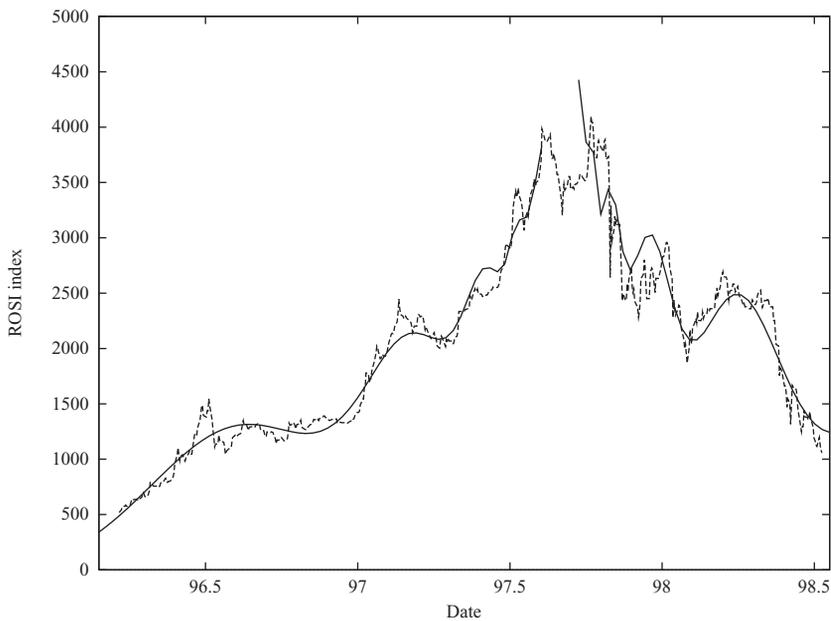

FIGURE 10. Symmetric 'bubble' and 'antibubble': in addition to the ascending part of the ROSI index, which is reproduced from Figure 9 with the same fit, we show the deflating part fitted with equation (19). The parameter values of the fit are $A_1 \approx 4922$, $B_1 \approx -3449$, $C_1 \approx 472$, $\beta \approx 0.32$, $t_c \approx 97.72$, $\phi_1 \approx 1.4$, $\omega_1 \approx 7.9$.
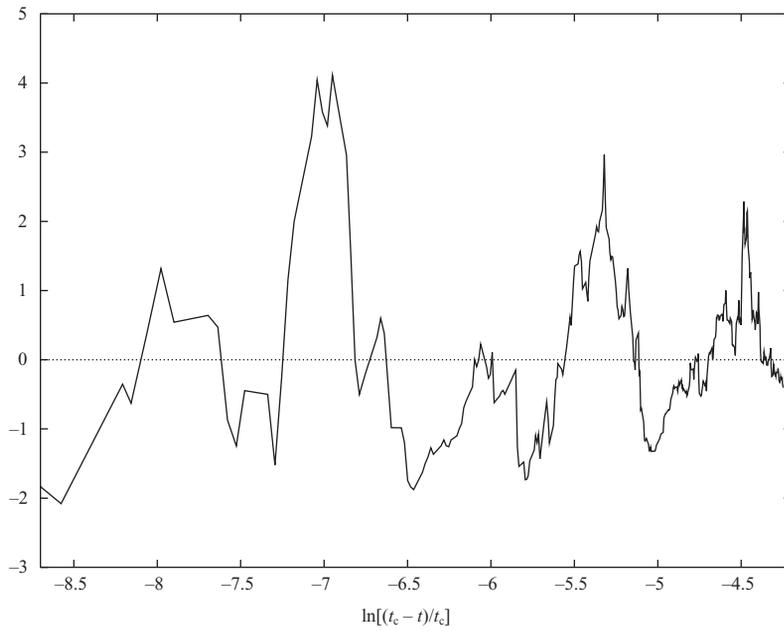
FIGURE 11. The residual as defined by the transformation (18) as a function of $\log[(t_c - t/t_c)]$ for the Russian data shown in Figure 9.
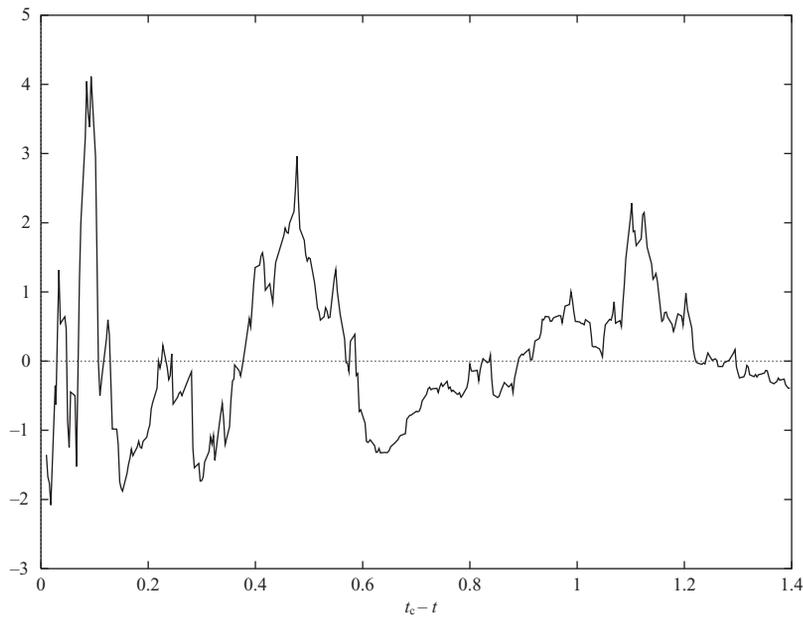

FIGURE 12. The residual as defined by the transformation (18) as a function of $t_c - t$ for the Russian data shown in Figure 9.

and Forex crashes supports this interpretation. Furthermore, it supports the idea of the stock market as a self-organizing complex system of surprising robustness. Inspired by this clear evidence of log-periodic oscillations decorating the power law acceleration signaling a bubble in the Russian stock market, we extended our analysis and searched for possible log-periodic signatures in the antibubble that followed the log-periodic bubble described above.

As described in Section 3.4, the recent decay of the Japanese Nikkei index from 1 January 1990 to the present can be excellently modeled by a log-periodically decorated power law decay. In Figure 10, we show the ROSI index for the antibubble, signaling the collapse of the Russian stock market, fitted with equation (15) with t_c and t interchanged. We use the price and not the logarithm of the price to preserve symmetry with the description of the bubble and from the consideration of the relatively short time scales involved.

In Table 2, the corresponding values for the fit of the antibubble with equation (19) are listed. The agreement between the bubble and the antibubble with respect to the values of the dimensionless variables β, ω_1, and λ, as well as t_c, is surprisingly good. However, the numerical stability of the result for the antibubble cannot be compared with that of the preceding bubble and depends on the index as well as the endpoint of the interval used in the fitting. However, the symmetry around t_c is rather striking considering the nature of the data and quite visible to the naked eye.

There are (at least) two reasons for the numerical instability of the fit to the antibubble. First, as an emerging market in decay the Russian stock market is too volatile for a smooth function such as the cosine and we are currently investigating alternative parametrizations (Johansen and Sornette 1999d). Second, the 'noise' in the Russian indices was very large in that period, in particular due to the 'depletion' of stocks. This is clearly illustrated by the heavy rearrangement that the ASPGEN and related indices went through following 17 August 1998.

In Figure 13, we see the spectral analysis for the bubble and antibubble. Whereas the agreement between the periodogram and the fit with equation (15) is excellent (within 1%), the deviation for the antibubble is ~5%. Furthermore, the residual for the antibubble used in the frequency analysis shown in Figure 13 was truncated by a few weeks because of an increasing deviation between data and fit.

It may seem odd to argue for the log-periodic oscillations while one can forcefully argue that the market is largely reflecting the vagaries of the Russian political institutions. For instance, in the antibubble case, February–April 1998 was a revival period for the market, characterized by the return of Western investors after the post-crash calm-down. This can be followed by studying the dynamics of the Russian external reserves. The timing of the return can be argued to be dictated by the risk policies of larger investors more than anything else. The next large drop of the Russian index in April 1998 originated from the decision of Mr. Yeltsin to sack Mr. Chernomyrdin's government, which destabilized the political situation and created uncertainty. Further political
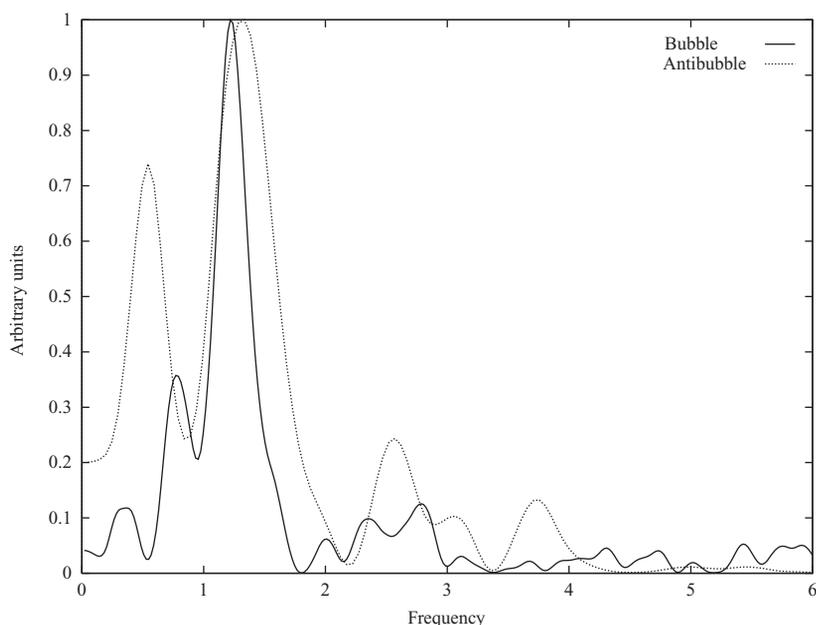
FIGURE 13. The Lomb periodogram for the bubble and antibubble shown in Figure 10. The significance of the peak should be estimated against the noise level.

disturbance was introduced twice by the Duma when it rejected Mr. Yeltsin's candidates for the prime minister's office and put itself on the brink of dissolution. In contrast, we argue that one must not mistake a global unstable situation for the specific historical action that triggered the instability. Consider a ruler put vertically on a table. Being in an unstable position, it will fall in some direction, and the specific air current or slight initial imperfection in the initial condition are of no real importance. What *is* important is the intrinsically unstable initial state of the ruler. We argue that a similar situation applies for crashes. They occur because the market has reached a state of global instability. Of course, there will always be specific events which may be identified as triggers of market motions but they will be the indicators rather than the deep sources of the instability. Furthermore, political events must also be considered as indicators of the state of the dynamical system which includes the market. There is, in principle, no decoupling between the different events. Specifically, the Russian crash may have been triggered by the Asian crises, but it was to a large extent fueled by the collapse of a banking system, which in the course of the bubble had created an outstanding debt of $19.2 billion (Malki 1999).

4. CONCLUSIONS

We have presented a synthesis of the available empirical evidence in the light of recent theoretical developments for the existence of characteristic log-periodic

signatures of growing bubbles in a variety of stock markets as well as currencies. We have here documented eight unrelated crashes from 1929 to 1998, on stock markets as diverse as the US, Hong Kong, and Russian markets, and on currencies. In addition, we have discovered a significant bubble on Wall Street ending in 1962 (Johansen *et al*. 2000), as well as 'antibubbles' on the Nikkei index since 1990 and the Gold price (after the 1980 bubble maximum) (Johansen and Sornette 1999c). Quite unexpectedly, we have shown that the Russian bubble crashing in August 1997 had a close to identical power law and log-periodic behavior to the bubbles observed on Wall Street and Hong Kong stock markets and on currencies. To our knowledge, no major financial crash preceded by an extended bubble has occurred in the past two decades without exhibiting a log-periodic signature. In this context, note that the novel analysis of the Russian index presented here was motivated by Ilinski's claim of a crash without log-periodic signature, which we have shown to be incorrect.

All these results, taken together with the remarkable robustness and consistency of the estimation of the exponent β as well as the more important statistics such as the scaling ratio λ, make the case for power law acceleration and log-periodicity very strong. In our opinion, one can no longer ignore these very specific and strong signatures that are characteristic of developing bubbles, and this calls for further investigations to unravel in greater depth the underlying economic, financial, and behavioral mechanisms.

These different cases, together with the 1962 'slow event', as well as the antibubbles, show that the log-periodic critical theory applies both to bubbles ending in a sudden crash as well as to bubbles landing smoothly. This is in fact a strong prediction of our rational model of imitative behavior.

What we have attempted here is not to explain why crashes happen or bubbles exist, but to quantify the process taking place during extended bubbles that very often lead to the rapid regime-switching a crash represents. We have offered evidence that bubbles and antibubbles have log-periodic and power law characteristics and hence provide a quantification of extended 'moods' on the markets. We have, furthermore, provided a model which contains the observed signatures. Whether, as proposed, discrete scale invariance or some other mechanism (Sornette 1998) is responsible for the prominent log-periodic signatures observed, only a more detailed microscopic model can answer.

REFERENCES

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

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.

Cenci, M., Cerquetti, A., and Peccati, L. (1996). Imitation and stability in a stock market. *European Journal of Operational Research*, **91**(2), 301–305.

Cutler, D. M., Poterba, J. M., and Summers, L. H. (1990). Speculative dynamics and the role of feedback traders. *American Economic Review*, **80**, 63–68.

Drozdz, S., Ruf, F., Speth, J., and Wojcik, M. (1999). Imprints of log-periodic self-similarity in the stock market. To appear in: *European Physical Journal* B. Preprint available at `http://xxx.lanl.gov/abs/cond-mat/9901025`.

Dubrulle, B., Graner, F., and Sornette, D. (1997). *Scale Invariance and Beyond*. EDP Sciences and Springer.

Embrechts, P., Klüppelberg, C., and Mikosch, T. (1997). *Modeling Extremal Events for Insurance and Finance*. Springer.

Embrechts, P., Resnick, S., and Samorodnitsky, G. (1998). Living on the edge. *Risk*, **11**(1), 96–100.

Embrechts, P., Resnick, S. I., and Samorodnitsky, G. (1999). Extreme value theory as a risk management tool. *North American Actuarial Journal*, **3**, 32–41.

Erman, W. T. (1999). Log spirals in the stock market. *Technical Analysis of Stock Commodities*, February, 16–34.

Feigenbaum, J. A., and Freund, P. G. O. (1996). Discrete scale invariance in stock markets before crashes. *International Journal of Modern Physics*, **10**, 3737–3745.

Feigenbaum, J. A. and Freund, P. G. O. (1998). Discrete scale invariance and the "Second Black Monday". *International Journal of Modern Physics*, **12**, 57–60.

Frost, A. J. and Prechter, R. (1998). *Elliott Wave Principle: Key to Market Behavior*, 7th edn. New Classics Library, Gainesville, Georgia.

Gluzman, S., and Yukalov, V. I. (1998). Renormalization group analysis of October market crashes. *Modern Physics Letters* B, **12**, 75–84.

Goldenfeld, N. (1992). *Lectures on Phase Transitions and the Renormalization Group*. Addison-Wesley.

Ilinski, K. (1999). Critical crashes? *International Journal of Modern Physics* C, **10**(4), in the press. Preprint available at `http://xxx.lanl.gov/abs/cond-mat/9903142`.

Intriligator, M. D. (1998). Russia: Lessons of the economic collapse. *New York Times*, 8 August. Presented to the World Bank, 4 September 1998.

Johansen, A. (1997). Discrete scale invariance and other cooperative phenomena in spatially extended systems with threshold dynamics. Ph.D. thesis, Niels Bohr Institute, Copenhagen. Available at `http://www.nbi.dk/~johansen/pub.html`.

Johansen, A., Ledoit, O., and Sornette, D. (2000). Crashes as critical points. To appear in: *International Journal of Theoretical and Applied Finance*, **3**(1). Preprint available at `http://www.nbi.dk/~johansen/pub.html`.

Johansen, A., and Sornette, D. (1998). Stock market crashes are outliers. *European Physics Journal* B, **1**, 141–143. Available at `http://www.nbi.dk/~johansen/pub.html`.

Johansen, A., and Sornette, D. (1999a). Critical crashes. *Risk*, **12**(1), 91–94.

Johansen, A., and Sornette, D. (1999b). Modeling the stock market prior to large crashes. *European Physics Journal* B, **9**, 167–174. Available at `http://www.nbi.dk/~johansen/pub.html`.

Johansen, A., and Sornette, D. (1999c). Financial ''anti-bubbles'': Log-periodicity in Gold and Nikkei collapses. To appear in: *International Journal of Modern Physics* C. Preprint available at `http://www.nbi.dk/~johansen/pub.html`.

Johansen, A., and Sornette, D. (1999d). In preparation.

Keim, D. B., and Madhavan, A. (1995). Anatomy of the trading process: Empirical evidence on the behavior of institutional traders. *Journal of Financial Economics* **37**, 371–398.

Laloux, L., Potters, M., Cont, R., Aguilar, J. P., and Bouchaud, J. P. (1999). Are financial crashes predictable? *Europhysics Letters*, **45**, 1–5.

Lin, P. C. (1998). International crashes as critical points: An empirical investigation. Working paper, The Anderson School of Management, UCLA.

Malki, E. (1999). The financial crisis in Russia. Available at `http://wueconb.wustl.edu/ewp-mac/9901001`.

Nelson, P. (1994). Voting and imitative behavior. *Economic Inquiry*, **32**, 92–102.

Press, W. H., Teulolsky, S. A., Vetterlong, W. T., and Flannery, B. P. (1994). *Numerical Recipes in Fortran*. Cambridge University Press.

Saleur, H., and Sornette, D. (1996). Complex exponents and log-periodic corrections in frustrated systems. *Journal de Physique I, France*, **6**, 327–355.

Shiller, R. J. (1990). Market volatility and investor behavior. *American Economic Review*, **80**, 58–62.

Shiller, R. J. (1995). Conversation, information, and herd behavior. *American Economic Review*, **85**, 181–185.

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

Sornette, D., and Johansen, A. (1997). Large financial crashes. *Physica* A, **245**, 411–422. Available at `http://www.nbi.dk/~johansen/pub.html`.

Sornette, D., and Johansen, A. (1998). A hierarchical model of financial crashes. *Physica* A, **261**, 581–598. Available at `http://www.nbi.dk/~johansen/pub.html`.

Sornette, D., Johansen, A., and Bouchaud, J.-P. (1996). Stock market crashes, Precursors and replicas. *Journal de Physique I, France*, **6**, 167–175. Available at `http://www.nbi.dk/~johansen/pub.html`.

Stauffer, D., and Sornette, D. (1998). Log-periodic oscillations for biased diffusion on 3D random lattice. *Physica* A, **252**, 271–277.

Vandewalle, N., Boveroux, Ph., Minguet, A., and Ausloos, M. (1998a). The crash of October 1987 seen as a phase transition: Amplitude and universality. *Physica* A, **255**, 201–210.

Vandewalle, N., Ausloos, M., Boveroux, Ph., and Minguet, A. (1998b). How the financial crash of October 1997 could have been predicted. *European Physics Journal* B, **4**, 139–141.

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

A. Johansen, D. Sornette, and O. Ledoit
UCLA, Los Angeles