Detection of crashes and rebounds in major equity markets

Wanfeng Yan, Reda Rebib and Ryan Woodard

Department of Management, Technology and Economics,
ETH Zurich, Kreuzplatz 5, CH-8032 Zurich, Switzerland
E-mail: wyan@ethz.ch
E-mail: reda.rebib@gmail.com
E-mail: rwoodard@ethz.ch

Didier Sornette*

Department of Management, Technology and Economics,
ETH Zurich, Kreuzplatz 5, CH-8032 Zurich, Switzerland
and
Swiss Finance Institute,
c/o University of Geneva,
40 Blvd. Du Pont d'Arve, CH-1211 Geneva 4, Switzerland
E-mail: dsornette@ethz.ch
*Corresponding author

Abstract: Financial markets are well known for their dramatic dynamics and consequences that affect much of the world's population. Consequently, much research has aimed at understanding, identifying and forecasting crashes and rebounds in financial markets. The Johansen-Ledoit-Sornette (JLS) model provides an operational framework to understand and diagnose financial bubbles from rational expectations and was recently extended to negative bubbles and rebounds. Using the JLS model, we develop an alarm index based on an advanced pattern recognition method with the aim of detecting bubbles and performing forecasts of market crashes and rebounds. Testing our methodology on ten major global equity markets, we show quantitatively that our developed alarm performs much better than chance in forecasting market crashes and rebounds. We use the derived signal to develop elementary trading strategies that produce statistically better performances than a simple buy-and-hold strategy.

Keywords: the JLS model; financial bubbles; crashes; rebounds; log-periodic power law; LPPL; pattern recognition method; alarm index; prediction; error diagram; trading strategy.

Biographical notes: Wanfeng Yan is currently a Visiting Scholar at Stern School of Business of New York University. He obtained his PhD at ETH Zurich under the supervision of Prof. Dr. Didier Sornette. His research interests are identifying and forecasts of financial bubbles, complex networks and human dynamics.

Reda Rebib currently works in the Geneva office of the Boston Consulting Group. He graduated from the Swiss Federal Institute of Technology in Zurich (ETHZ) with a Master in Management, Technology and Economics (MTEC). He mainly focused his master studies on financial mathematics and economics. This focus led him to write his master thesis on 'The detection of equity market crashes and recoveries' in collaboration with Prof. Sornette's chair and the alternative investment firm Man Group. Before his Master, he completed his Bachelor in Mathematics at the Swiss Federal Institute of Technology in Lausanne (EPFL).

Ryan Woodard is a Senior Researcher at the Swiss Federal Institute of Technology (ETH Zurich). His research focuses on the prediction of crises and extreme events in complex systems, including financial markets and social organisations. He specialises in identifying and forecasting financial bubbles and crashes and in analysing the statistics of deadly terrorist events. In 2008, he helped launch the Financial Crisis Observatory to test the hypothesis that financial bubbles can be diagnosed in real-time and their termination can be predicted probabilistically.

Didier Sornette is a Professor of Entrepreneurial Risks in the Department of Management, Technology and Economics at the Swiss Federal Institute of Technology (ETH Zurich), and a Professor of Finance at the Swiss Finance Institute. He is the author 450+ research papers and five books. His research focuses on the prediction of crises and extreme events in complex systems and in particular of financial bubbles and crashes, and the diagnostic of systemic instabilities. In 2008, he launched the Financial Crisis Observatory to test the hypothesis that financial bubbles can be diagnosed in real-time and their termination can be predicted probabilistically.

1 Introduction

Hundreds of millions of people are influenced by bubbles and crashes in the financial system. Although many academics, practitioners and policy makers have studied these extreme events, there is little consensus yet about their causes or even definitions [see ref. Kaizoji and Sornette (2009) for a review of existing models and references therein]. More than a decade ago, the Johansen-Ledoit-Sornette (JLS) model (Johansen et al., 1999, 2000; Johansen and Sornette, 1999a; Sornette, 2003) has been developed to detect bubbles and crashes in financial markets. This model states that imitation and herding behaviour of noise traders during the bubble regime may lead to transient accelerating price growth that may end with a possible crash. Before the crash, the price grows at a faster-than-exponential rate rather than an exponential rate, which is normally used in other bubble models (Watanabe et al., 2007; Abreu and Brunnermeier, 2003). This faster-than-exponential growth is due to positive feedback in the valuation of assets among the investors. Further consideration of tension and competition between the

fundamentalists and the noise traders leads the real prices to oscillate about this growth rate with increasing frequency. This oscillation is actually periodic in the logarithm of time before the most likely end of the bubble growth and, therefore, the model describes log-periodic power law (LPPL) behaviour.

Since positive feedback mechanisms in the markets may also lead to transient periods of accelerating *decrease* in price, followed in turn by rapid reversals, the JLS model was recently extended to study 'negative bubbles' and rebounds (Yan et al., 2010, 2012a). There, it was shown that, given the existence of extended crashes and rebounds, negative bubbles could be identified and their endings forecast in the same spirit as for 'traditional' bubbles such as the 2006–2008 oil bubble (Sornette et al., 2009), the Chinese index bubble in 2009 (Jiang et al., 2010), the real estate market in Las Vegas (Zhou and Sornette, 2008), the UK and USA real estate bubbles (Zhou and Sornette, 2003, 2006), the Nikkei index anti-bubble in 1990–1998 (Johansen and Sornette, 1999b), the S&P 500 index anti-bubble in 2000–2003 (Sornette and Zhou, 2002), the Dow Jones industrial average historical bubbles (Vandewalle et al., 1999), the corporate bond spreads (Clark, 2004), the Polish stock market bubble (Gnaciński and Makowiec, 2004), the Western stock markets (Bartolozzi et al., 2005), the Brazilian real (R\$) – US dollar (USD) exchange rate (Matsushita et al., 2006), the 2000–2010 world major stock indices (Drożdż et al. (2008), the South African stock market bubble (Zhou and Sornette, 2009) and the US repurchase agreements market (Yan et al., 2012b). Moreover, new experiments in ex-ante bubble detection and forecast were recently launched in the Financial Crisis Observatory at ETH Zurich (Sornette et al., 2010a, 2010b; Woodard et al., 2011).

To test this idea and to implement a systematic forecast procedure based on the JLS model, Sornette and Zhou (2006) adapted a pattern recognition method to detect bubbles and crashes. This method has two major advantages: first, it makes full use of the historical results and therefore increases the forecast quality; second, it enables the development of *systematic* forecasts. This method was originally developed by mathematician I.M. Gelfand and his collaborators in the mid-1970s as an earthquake prediction scheme. Since then, this method has been widely used in many kinds of predictions, ranging from uranium prospecting (Briggs and Press, 1977) to unemployment rates (Keilis-Borok et al., 2005; Yan et al., 2010, 2012a) extended this method to negative bubbles and rebounds of financial markets. They also improved the method in Sornette and Zhou (2006) by separating the learning period and prediction period to enable a pure causal prediction.

All these previous works are limited by the fact that they focused on the detection of either crashes or rebounds of only one stock market index (the S&P 500). The present paper improves on the current research by studying:

1   the general applicability of the pattern recognition method to global equity markets

2   at the same time, the overall predictability offered by this method in detecting market crashes and rebounds.

In the hope of addressing these questions, in this paper, we expand this method (the Gelfand pattern recognition methodology combined with LPPL calibrations) to ten major global equity markets, with the goal of detecting and forecasting crashes and rebounds. Our results indicate that the performance of the predictions on both crashes and rebounds for most of the indices is better than chance. That is, the end of large

drawups and drawdowns and the subsequent crashes and rebounds can be successfully forecast. We also extend the previous works, which use error diagrams as the only metric of success, by using a simple trading strategy to test for the predictability of the method. We show that the trading strategy outperforms a simple buy-and-hold benchmark.

The structure of the paper is as follows. In Section 2, we briefly introduce the JLS model and the bubble/negative bubble versions of the model. Then, we present the pattern recognition method for the prediction of crashes and rebounds in Section 3. The quality of the prediction is tested in Section 4 using error diagrams to compare missed events versus total alarm time. We next introduce the trading strategy based on the alarm index and test its performance in Section 5. We summarise our results and conclude in Section 6.

2   The JLS model for detecting bubbles and negative bubbles

The JLS model (Johansen and Sornette, 1999a; Johansen et al., 1999, 2000; Sornette, 2003) is an extension of the rational expectation bubble model of Blanchard and Watson (1983). In this model, a financial bubble (negative bubble) is modelled as a regime of super-exponential power law growth (decline) punctuated by short-lived corrections organised according to the symmetry of discrete scale invariance (Sornette, 1998). The super-exponential power law is argued to result from positive feedback resulting from noise trader decisions that tend to enhance deviations from fundamental valuation in an accelerating spiral. That is, the price of a stock goes higher (lower) than the fundamental value during the bubble (negative bubble) period, ending in a sudden regime change.

In the JLS model, the dynamics of asset prices is described as

$$\frac{dp}{p} = \mu(t)dt + \sigma(t)dW - \kappa dj, \tag{1}$$

where p is the stock market price, μ is the drift (or trend) and dW is the increment of a Wiener process (with zero mean and unit variance). The term dj represents a discontinuous jump such that $dj = 0$ before the crash or rebound and $dj = 1$ after the crash or rebound occurs. The change amplitude associated with the occurrence of a jump is determined by κ. The parameter κ is positive for bubbles and negative for negative bubbles. The assumption of a constant jump size is easily relaxed by considering a distribution of jump sizes, with the condition that its first moment exists. Then, the no-arbitrage condition is expressed similarly with κ replaced by the mean of the distribution. Each successive crash corresponds to a jump of dj by one unit. The dynamics of the jumps is governed by a hazard rate $h(t)$. For the bubble (negative bubble) regime, $h(t)dt$ is the probability that the crash (rebound) occurs between t and $t + dt$ conditional on the fact that it has not yet happened. We have $E_t[dj] = 1 \times h(t)dt + 0 \times (1 - h(t)dt)$ and therefore

$$\mathrm{E}_t[dj] = h(t)dt. \tag{2}$$

Under the assumption of the JLS model, noise traders exhibit collective herding behaviours that may destabilise the market. The JLS model assumes that the aggregate

effect of noise traders should have a discrete scale invariant property (Sornette, 1998, 2003). Therefore, it can be accounted for by the following dynamics of the hazard rate

$$h(t) = B'(t_c - t)^{m-1} + C'(t_c - t)^{m-1} \cos(\omega \ln(t_c - t) - \phi') \ . \tag{3}$$

The no-arbitrage condition reads $E_t[dp] = 0$, where the expectation is performed with respect to the risk-neutral measure and in the frame of the risk-free rate. This is the standard condition that the price process is a martingale. Taking the expectation of expression (1) under the filtration (or history) until time t reads

$$E_t[dp] = \mu(t)p(t)dt + \sigma(t)p(t)E_t[dW] - \kappa p(t)E_t[dj] \ . \tag{4}$$

Since $E_t[dW] = 0$ and $E_t[dj] = h(t)dt$ [equation (2)], together with the no-arbitrage condition $E_t[dp] = 0$, this yields

$$\mu(t) = \kappa h(t) \ . \tag{5}$$

This result (5) expresses that the return $\mu(t)$ is controlled by the risk of the crash quantified by its hazard rate $h(t)$.

Now, conditioned on the fact that no crash or rebound has occurred, equation (1) is simply

$$\frac{dp}{p} = \mu(t)dt + \sigma(t)dW = \kappa h(t)dt + \sigma(t)dW \ . \tag{6}$$

Its conditional expectation leads to

$$E_t\left[\frac{dp}{p}\right] = \kappa h(t)dt \ . \tag{7}$$

Substituting with the expression (3) for $h(t)$ and integrating yields the so-called LPPL equation:

$$\ln E[p(t)] = A + B(t_c - t)^m + C(t_c - t)^m \cos(\omega \ln(t_c - t) - \phi) \tag{8}$$

where $B = -\kappa B'/m$ and $C = -\kappa C'/\sqrt{m^2 + \omega^2}$. Note that this expression (8) describes the average price dynamics only up to the end of the bubble or negative bubble. The JLS model does not specify what happens beyond t_c. This critical t_c is the termination of the bubble/negative bubble regime and the transition time to another regime. This regime could be a big crash/rebound or a change of the growth rate of the market. The dynamics of the Merrill Lynch European Monetary Union (EMU) Corporates Non-Financial Index in 2009 (Sornette et al., 2010a) provides a vivid example of a change of regime characterised by a change of growth rate rather than by a crash or rebound. For $m < 1$, the crash hazard rate accelerates up to t_c but its integral up to t, which controls the total probability for a crash (or rebound) to occur up to t, remains finite and less than 1 for all times $t \leq t_c$. It is this property that makes it rational for investors to remain invested knowing that a bubble (or negative bubble) is developing and that a crash (or rebound) is looming. Indeed, there is still a finite probability that no crash or rebound will occur during the lifetime of the bubble. The

excess return $\mu(t) = \kappa h(t)$ is the remuneration in the case of a bubble (the cost in the case of a negative bubble) that investors require (accept to pay) to remain invested in the bubbly asset, which is exposed to a crash (rebound0 risk. The condition that the price remains finite at all time, including t_c, imposes that $m > 0$.

Within the JLS framework, a bubble is qualified when the hazard rate accelerates. According to (3), this imposes $0 < m < 1$. Since, by definition, the hazard rate should be non-negative, an additional constraint is (van Bothmer and Meister, 2003)

$$b \equiv -Bm - |C|\sqrt{m^2 + \omega^2} \geq 0 \;, \tag{9}$$

for bubbles and

$$b \equiv -Bm - |C|\sqrt{m^2 + \omega^2} \leq 0 \;, \tag{10}$$

for negative bubbles.

3   Prediction method

We adapt the pattern recognition method of Gelfand et al. (1976) to generate predictions of crashes and rebound times in financial markets on the basis of the detection and calibration of bubbles and negative bubbles. Here we give a brief summary of the method, which is decomposed into five steps.

3.1   Fit the time series with the JLS model

Given a historical price time series of an index (such as the S&P 500 or the Dow Jones industrial average, for example), we first divide it into different sub-windows (t_1, t_2) of length $dt \equiv t_2 - t_1$ according to the following rules:

1   The earliest start time of the windows is t_{10}. Other start times t_1 are calculated using a step size of $dt_1 = 50$ calendar days.

2   The latest end time of the windows is t_{20}. Other end times t_2 are calculated with a negative step size $dt_2 = -50$ calendar days.

3   The minimum window size $dt_{\min} = 110$ calendar days.

4   The maximum window size $dt_{\max} = 1500$ calendar days.

For each sub-window generated by the above rules, the log of the index is fit with the JLS equation (8). The fitting procedure is a combination of a preliminary heuristic selection of the initial points and a local minimising algorithm (least squares). The linear parameters are slaved by the non-linear parameters before fitting. Details of the fitting algorithm can be found in Yan et al. (2011) and Johansen et al. (2000). We keep the best ten parameter sets for each sub-window and use these parameter sets as the input to the pattern recognition method.

3.2 Definition of crash and rebound

We refer to a crash as 'Crh' and to a rebound as 'Rbd'. A day d begins a crash (rebound) if the price on that day is the maximum (minimum) price in a window of 100 days before and 100 days after. That is,

$$\text{Crh} = \{d \mid P_d = \max\{P_x\}, \forall x \in [d - 100, d + 100]\} \tag{11}$$

$$\text{Rbd} = \{d \mid P_d = \min\{P_x\}, \forall x \in [d - 100, d + 100]\} \tag{12}$$

where P_d is the adjusted closing price on day d. Our task is to diagnose such crashes and rebounds in advance. We could also use other windows instead of ± 100 to define a rebound. The results are stable with respect to a change of this number because we learn from the 'learning set' with a certain rebound window width and then try to predict the rebounds using the same window definition. The reference (Yan et al., 2012a) shows the results for ± 200 days and ± 365 days type of rebounds.

3.3 Learning set, class, group and informative parameter

As described above, we obtain a set of parameters that best fit the model (8) for each window. Then we select a subset of the whole set which only contains the fits of crashes and rebounds with critical times found within the window (that is, where parameter t_c is *not* calculate to be beyond the window bounds). We learn the properties of historical rebounds from this set and develop the predictions based on these properties. We call this set the *learning set*. In this paper, a specific day for each index is chosen as the 'present time' (for backtesting purposes). All the fit windows before that day will be used as the learning set and all the fit windows after that will be used as the *testing set*, in which we will predict future rebounds. The quality of the predictability of this method can be quantified by studying the predicted results in the testing set using only the information found in the learning set.

Each of the sub-windows generated by the rules in Section 3.1 will be assigned one of two *classes* and one of 14 *groups*. Classes indicate how close the modelled critical time t_c is to a historical crash or rebound, where Class I indicates 'close' and Class II indicates 'not close' ('close' will be defined below as a parameter). Groups of windows have similar window widths. For each fit, we create a set of six parameters: m, ω, ϕ and B from equation (8), b fromequation (9) and q as the residual of the fit. We will compare the probability density functions (pdf) of these parameters among the different classes and groups. The main goal of this technique is *to identify patterns of parameter pdf's that are different between windows with crashes or rebounds and windows without crashes or rebounds*. Given such a difference and a new, out-of-sample window, we can probabilistically state that a given window will or will not end in a crash or rebound.

A figure is very helpful for understanding these concepts. We show the selection of the sub-windows and sort the fits by classes and groups in Figure 1(a). Then we create the pdf's of each of these parameters for each fit and define *informative parameters* as those parameters for which the pdf's differ significantly according to a Kolmogorov-Smirnov test. For each informative parameter, we find the regions of the abscissa of the pdf for which the Class I pdf (fits with t_c close to an

extremum) is greater than the Class II pdf. This procedure has been performed in Figure 1(b).

3.4 Questionnaires, traits and features

Using the informative parameters and their pdf's described above, we can generate a *questionnaire* for each day of the learning or testing set. A questionnaire is a quantitative inquiry into whether or not a set of parameters is likely to indicate a bubble or negative bubble. Questionnaires will be used to identify bubbles (negative bubbles) which will be followed by crashes (rebounds). In short, the length of a questionnaire tells how many informative parameters there are for a given window size. An informative parameter implies that the pdf's of that parameter are very different for windows with a crash/rebound than windows without. The more values of '1' in the questionnaire, the more likely it is that the parameters are associated with a crash/rebound. One questionnaire is constructed for each day t_{scan} in our learning set. We first collect all the fits which have a critical time near that day ('near' will be defined). Then we create a string of bits whose length is equal to the number of informative parameters found. Each bit can take a value –1, 0 or 1 (a balanced ternary system). Each bit represents the answer to the question: are more than half of the collected fits more likely to be considered as Class I? If the answer is 'yes', we assign 1 in the bit of the questionnaire corresponding to this informative parameter. Otherwise, we assign 0 when the answer cannot be determined or −1 when the answer is 'no'. A visual representation of this questionnaire process is shown in Figure 1(c).

The concept of a *trait* is developed to describe the property of the questionnaire for each t_{scan}. Each questionnaire can be decomposed into a fixed number of traits if the length of the questionnaire is fixed. We will not give the details in how the traits are generated in this paper. For a clear explanation of the method, please refer to Section 3.9 of the reference (Yan et al., 2012a). Think of a trait as a sub-set (like the 'important' short section from a very long DNA sequence) of a fixed-length questionnaire that is usually found in windows that show crashes/rebounds. Conversely, a trait can indicate windows where a crash/rebound is not found.

Assume that there are two sets of traits T_I and T_{II} corresponding to Class I and Class II, respectively. Scan day by day the date t before the last day of the learning set. If t is 'near' an extreme event (crash or rebound), then all traits generated by the questionnaire for this date belong to T_I. Otherwise, all traits generated by this questionnaire belong to T_{II}. 'Near' is defined as at most 20 days away from an extreme event. The same definition will be used later in Section 4 when we introduce the error diagram.

Using this threshold, we declare that an alarm starts on the first day that the unsorted crash alarm index time series exceeds this threshold. The duration of this alarm D_a is set to 41 days, since the longest distance between a crash and the day with index greater than the threshold is 20 days.

Count the frequencies of a single trait τ in T_I and T_{II}. If τ is in T_I for more than α times and in T_{II} for less than β times, then we call this trait τ a *feature F_I* of Class I. Similarly, if τ is in T_I for less than α times and in T_{II} for more than β times, then we call τ a *feature F_{II}* of Class II. The pair (α, β) is defined as a *feature qualification*. Figure 1(d) shows the generation process of traits and features. We would like to clarify

that by definition some of the traits are not from any type of feature since they are not 'extreme' and we cannot extract clear information from them.

Figure 1   Pattern recognition procedure, (a) Step 1: create sub-windows, fit each window with the JLS model, classify the fits in groups and classes (b) Step 2: for each group compare fits in Class I with those in Class II and find out the informative parameters, (c) Step 3: generate the questionnaire for each trading day, (d) Step 4: create traits from questionnaires, obtain features for each class by traits statistics, (e) Step 5: construct the alarm index for each day by decomposing the questionnaire into traits for that day and compare these traits to the features of each class (see online version for colours)

(a)

(b)

(c)

(d)

(e)
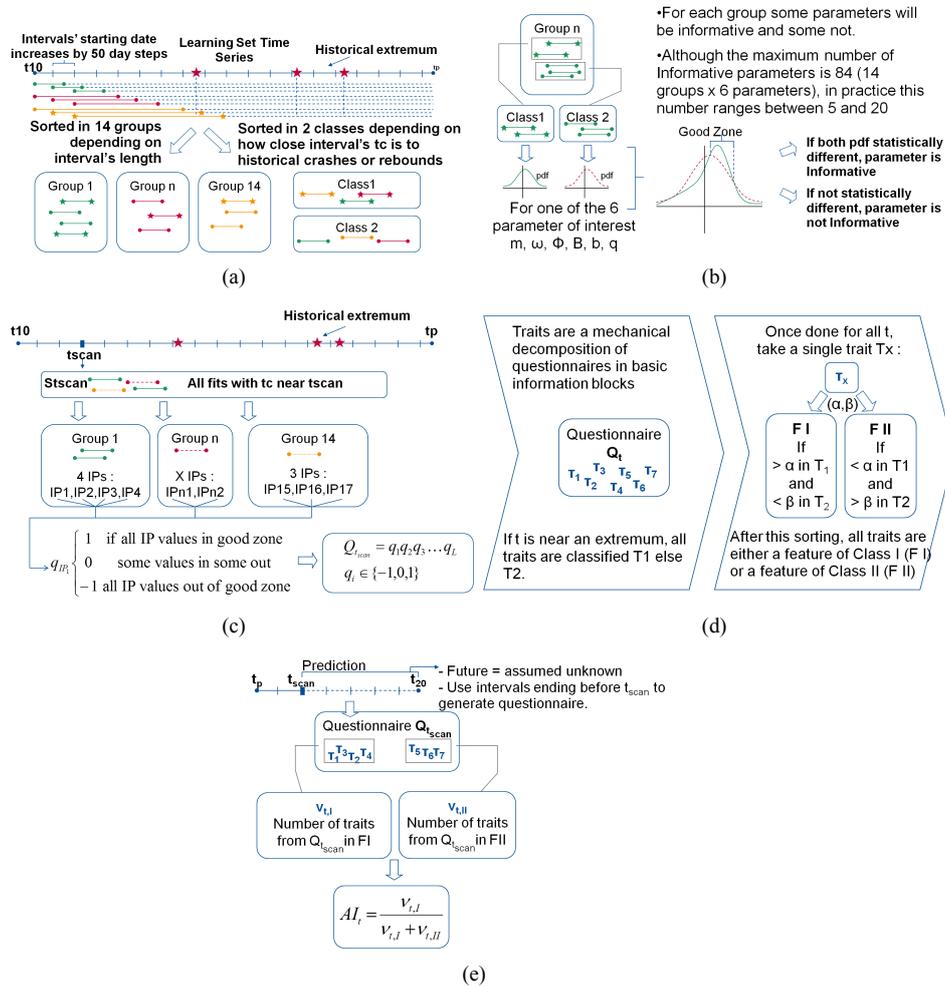
3.5   Alarm index

The final piece in our methodology is to define an *alarm index* for both crashes and rebounds. An alarm index is developed based on features to show the probability that a

certain day is considered to be a rebound or a crash. We first collect all the fits which have a predicted critical time near this specific day and generate questionnaires and traits from these fits. The rebound (crash) alarm index for a certain day is just a ratio quantified by the total number of traits from feature type F_I (a set of traits which have high probability to represent rebound (crash)) divided by the total number of traits from both F_I and F_{II}. Note that F_{II} is a set of traits which have low probability to represent rebound (crash). The principles for the generation of the alarm index are summarised in Figure 1(e).

Two types of alarm index are developed. One is for the back tests in the learning set, as we have already used the information before this time to generate informative parameters and features. The other alarm index is for the prediction tests. We generate this prediction alarm index using only the information before a certain time and then try to predict crashes and rebounds in the 'future'.

Although the pattern recognition method provides a systematic approach for detecting financial rebounds and crashes, and the forecast qualify is very good as we will show in the next section, the implementation of the current method is not straightforward due to its complexity. Besides, this method requires a relatively large amount of computation. Taking the work presented in the present paper as an example, we used 200 CPU cores for about one week to obtain the results shown below. However, with the development of faster and faster computers, and in particular of GPUs (Preis, 2011), we expect that this high computation cost will become less and less an issue in the future.

4 Prediction in major equity markets

We perform systematic detections and forecasts on both the market crashes and rebounds for ten major global equity markets using the method we discussed in Section 3. The ten indices are S&P 500 (USA), Nasdaq composite (USA), Russell 2000 (USA), FTSE 100 (UK), CAC 40 (France), SMI (Switzerland), DAX (German), Nikkei 225 (Japan), Hang Seng (Hong Kong) and ASX (Australia). The basic information for these indices used in this study is listed in Table 1. Due to space constraints, we cannot show all results for these ten indices here. The complete results for three indices from different continents are shown in this paper. They are Russell 2000 (USA), SMI (Europe) and Nikkei 225 (Asia). Partial results for the remaining indices are shown in Table 2, which will be discussed in details later in Section 5.

The alarm index depends on the features which are generated using information from the learning set. Thus, the alarm index before the end of the learning set uses 'future' information. That is, the value of the alarm index on a certain day t_0 in the learning set uses prices found at $t > t_0$ to generate features. The feature definitions from the learning set are then used to define the alarm index in the *testing* set using *only* past prices. That is, the value of the alarm index on a certain day t_0 in the *testing* set uses only prices found at $t \leq t_0$ in the testing set (and the *definitions* of features found in the *learning* set). We do not use 'future' information in the testing set. In this case, the alarm index predicts crashes and rebounds in the market.

The crash and the rebound alarm index for Russell 2000, SMI and Nikkei are shown in Figures 2 to 4. Figures 3(a), 4(a) and 5(a) show the back testing results for these three indices and Figures 3(b), 4(b) and 5(b) present the prediction results. In all of these

results, the feature qualification pair $(7, 100)$ is used. This means that a certain trait must appear in trait Class I (crash or rebound) at least seven times *and* must appear in trait Class II (no crash or rebound) less than 100 times. If so, then we say that this trait is a feature of Class I. If, on the other hand, the trait appears seven times or less in Class I *or* appears 100 times or more in Class II, then this trait is a feature of Class II. Tests on other feature qualification pairs are performed also. Due to the space constraints, we do not show the alarm index constructed by other feature qualification pairs here, but later we will present the predictability of these alarm indices by showing the corresponding error diagrams. In the rest of this paper, if we do not mention otherwise, we use $(7, 100)$ as the feature qualification pair.

Table 1 Information for the tested indices

Index name	Yahoo ticker	Learning start	Prediction start	Prediction end	win. #
S&P 500	^GSPC	5-Jan-1950	26-Mar-1999	3-Jun-2009	11662
Nasdaq	^IXIC	13-Dec-1971	20-Mar-1999	30-Jul-2010	7209
Russell 2000	^RUT	30-Sep-1987	17-Apr-1999	27-Aug-2010	4270
FTSE 100	^FTSE	3-May-1984	17-Apr-1999	27-Aug-2010	4970
CAC 40	^FCHI	1-Mar-1990	17-Apr-1999	27-Aug-2010	3766
SMI	^SSMI	9-Nov-1990	17-Apr-1999	27-Aug-2010	3626
DAX	^GDAXI	26-Nov-1990	17-Apr-1999	27-Aug-2010	3626
Nikkei 225	^N225	4-Jan-1984	17-Apr-1999	27-Aug-2010	5026
Hang Seng	^HSI	31-Dec-1986	17-Apr-1999	27-Aug-2010	4410
ASX	^AORD	6-Aug-1984	17-Apr-1999	27-Aug-2010	4914

Note: Yahoo ticker of each index, starting time of learning and prediction periods, ending time of prediction and number of sub-windows.

Figure 2 Alarm index and log-price of the Russell 2000 Index, (a) *learning* set (before 17-Apr-1999) (b) *predicting* set (after 17-Apr-1999)

(a)   (b)
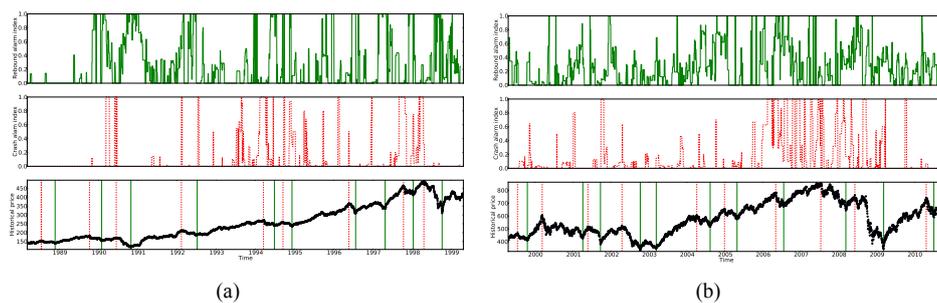
Notes: For each figure: (upper) rebound alarm index using feature qualification pair (7, 100), (middle) crash alarm index using feature qualification pair (7, 100). The rebound (crash) alarm index is in the range [0, 1]. The higher the rebound (crash) alarm index, the more likely is the occurrence of a rebound (crash). (lower) Plot of price versus time of Russell index (shown in blue cycles). Green solid vertical lines indicate rebounds defined by local minima within plus and minus 100 days around them. Red dashed vertical lines indicate crashes defined by local maxima within plus and minus 100 days around them. Note that these rebounds and crashes are the historical 'change of regime' rather than only the jump-like reversals.

Figure 3   Alarm index and price of the SMI index, (a) *learning* set (before 17-Apr-1999)
(b) *predicting* set (after 17-Apr-1999)

(a)   (b)
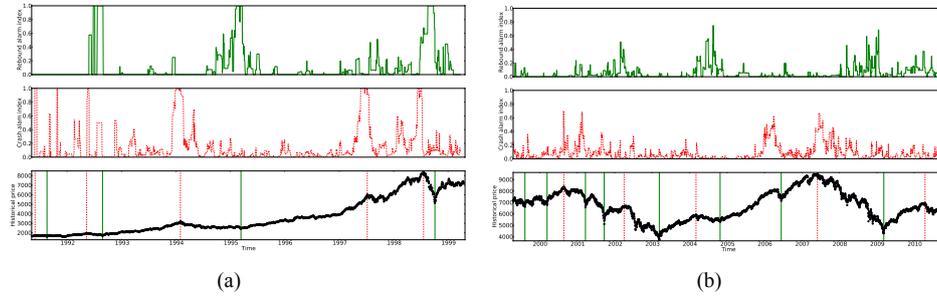
Note: The format is the same as Figure 2.

Figure 4   Alarm index and price of the Nikkei index, (a) *learning* set (before 17-Apr-1999)
(b) *predicting* set (after 17-Apr-1999)

(a)   (b)
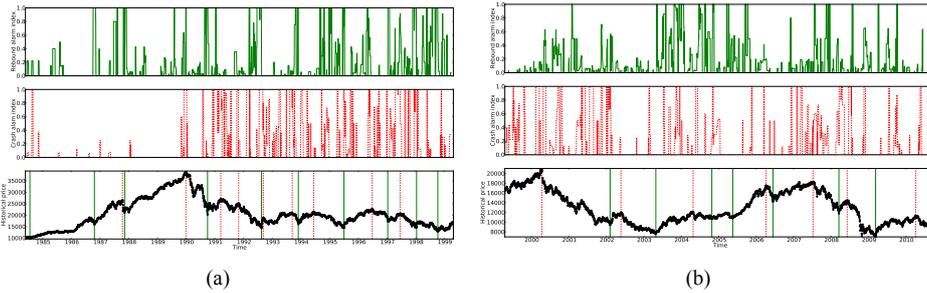
Note: The format is the same as Figure 2.

To check the quality of the alarm index quantitatively, we introduce error diagrams (Mochan, 1997; Mochan and Kagan, 1992). Using Nikkei 225 as an example, we create an error diagram for crash predictions after 17-Apr-1999 with a certain feature qualification in the following way:

1   Calculate features and define the alarm index using the *learning set* between 04-Jan-1984 and 17-Apr-1999.

2   Count the number of crashes after 17-Apr-1999 as defined in Section 3.2 and expression (11). There are seven crashes.

3   Take the crash alarm index time series (after 17-Apr-1999) and sort the set of all alarm index values in decreasing order. There are 4,141 points in this series and the sorting operation delivers a list of 4,141 index values, from the largest to the smallest one.

4   The largest value of this sorted series defines the first threshold.

5   Using this threshold, we declare that an alarm starts on the first day that the unsorted crash alarm index time series exceeds this threshold. The duration of this alarm D_a is set to 41 days, since the longest distance between a crash and the day with index greater than the threshold is 20 days. This threshold is consistent with the previous classification of questionnaires in Section 3.4, where we define a predicted critical time as 'near' the real extreme events when its distance is less than 20 days. Then, a prediction is deemed successful when a crash falls inside that window of 41 days.

6   If there are no successful predictions at this threshold, move the threshold down to the next value in the sorted series of alarm index.

7   Once a crash is predicted with a new value of the threshold, count the ratio of unpredicted crashes (unpredicted crashes/total crashes in set) and the ratio of alarms used (duration of alarm period/4,141 prediction days). Mark this as a single point in the error diagram.

In this way, we will mark seven points in the error diagram for the seven Nikkei 225 crashes after 17-Apr-1999.

The aim of using such an error diagram in general is to show that a given prediction scheme performs better than random. A random prediction follows the line $y = 1 - x$ in the error diagram. A set of points below this line indicates that the prediction is better than randomly choosing alarms. The prediction is seen to improve as more error diagram points are found near the origin point $(0, 0)$. The advantage of error diagrams is to avoid discussing how different observers would rate the quality of predictions in terms of the relative importance of avoiding the occurrence of false positive alarms and of false negative missed rebounds. By presenting the full error diagram, we thus sample all possible preferences and the unique criterion is that the error diagram curve be shown to be statistically significantly below the anti-diagonal $y = 1 - x$.

Figure 5   Error diagram for back tests and predictions of crashes and rebounds for Russell 2000 index with different types of feature qualifications, (a) back tests of crashes (b) back tests of rebounds (c) predictions of crashes (d) predictions of rebounds
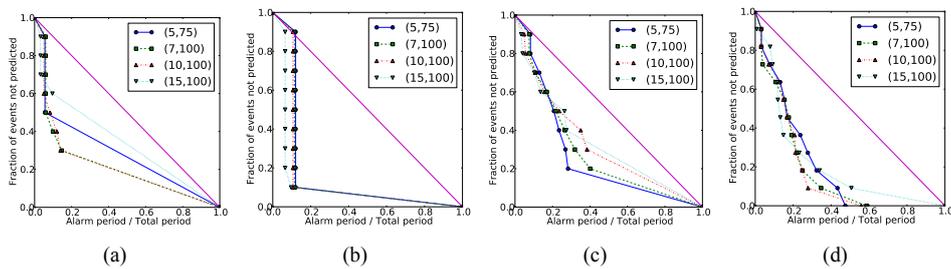
Notes: The value of the feature qualifications are shown in the legend. The fact that all the curves lie under the line $y = 1 - x$ indicates better performance than chance in detecting crashes and rebounds using our method. Feature qualification (α, β) means that, if the occurrence of a certain trait in Class I is larger than α and less than β, then we call this trait a feature of Class I and vice versa. See text for more information.

Figure 6   Error diagram for back tests and predictions of crashes and rebounds for SMI index with different types of feature qualifications, (a) back tests of crashes (b) back tests of rebounds (c) predictions of crashes (d) predictions of rebounds

(a)   (b)   (c)   (d)
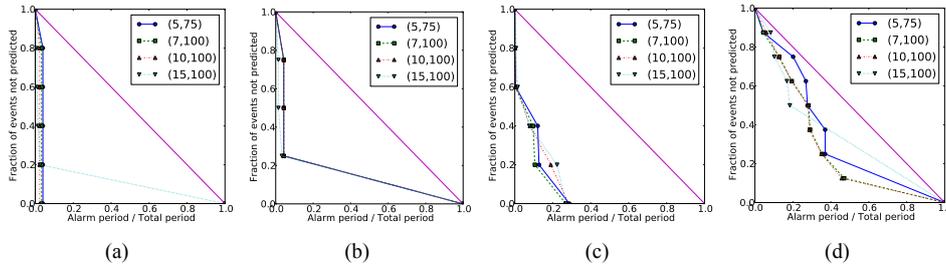
The format is the same as Figure 5.

Figure 7   Error diagram for back tests and predictions of crashes and rebounds for Nikkei 225 index with different types of feature qualifications, (a) back tests of crashes (b) back tests of rebounds (c) predictions of crashes (d) predictions of rebounds

(a)   (b)   (c)   (d)
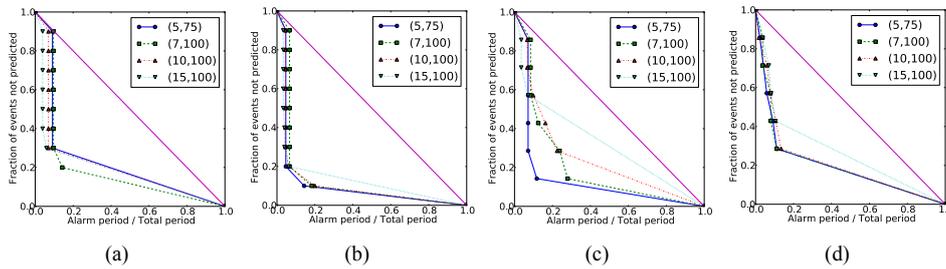
The format is the same as Figure 5.

In Figures 5 to 7, we show the results on predictions and back tests in terms of error diagrams for crashes and rebounds in each of the indices. The results for different feature qualification pairs (α, β) are shown in each figure. All these figures show that our alarm index for crashes and rebounds in either back testing or prediction performs much better than random, which strengthens the validation of our methodology.

5   Trading strategy

One of the most powerful methods to test the predictability of a signal is to design simple trading strategies based on it. We do so with our alarm index by using simple moving average strategies, which keep all the key features of the alarm index and avoid parametrisation problems. The strategies are kept as simple as possible and can be applied to any indices.

The trading strategies are designed as follows: the daily exposure of our strategy θ is determined by the average value of the alarm index for the past n days. The rest of our wealth, $1 - \theta$, is invested in a three-month US treasury bill.

Let us denote the average rebound and crash alarm index of the past n days as AI_R and AI_C respectively. We create three different strategies:

- A long strategy using only the rebound alarm index. We will take a long position in this strategy only. The daily exposure of our strategy is based on the average value of the past n days rebound alarm index: $\theta = AI_R$.

- A short strategy using only the crash alarm index. We will take a short position in this strategy only. The daily absolute exposure of our strategy is based on the average value of the past n days crash alarm index: $|\theta| = AI_C$.

- A long-short strategy linearly combining both strategies above. When the average rebound alarm index is higher than the average crash alarm index, we take a long position and vice versa. The absolute exposure $|\theta| = |AI_R - AI_C|$.

These strategies have the advantage of having few parameters, as only the duration n needs to be determined. Despite their simplicity, they capture the two key features of the alarm index. First, we see that the alarms are clustered around certain dates. The more clustering seen, the more likely that a change of regime is coming and, therefore, the more we should be invested. Second, we see that a strong alarm close to 1 should be treated as more important than a weaker alarm while at the same time the smaller alarms still contain some information and should not be discarded.

Table 2 summarises the Sharpe ratios for long-short strategies on the out-of-sample period (testing set) for each index. The strategy is calculated with four different moving average look-backs: $n = 20, 30, 40, 60$ days. We use the Sharpe ratios of the market during this period as the benchmark in this table. Recall that the Sharpe ratio is a measure of the excess return (or risk premium) per unit of risk in an investment asset or a trading strategy. It is defined as:

$$S = \frac{R - R_f}{\sigma} = \frac{R - R_f}{\sqrt{Var[R - R_f]}} \, , \tag{13}$$

where R is the return of the strategy and R_f is the risk free rate. We use the US three-month treasury bill rate here as the risk free rate. The Sharpe ratio is used to characterise how well the return of an asset compensates the investor for the risk taken: the higher the Sharpe ratio number, the better. When comparing two assets with the same expected return against the same risk free rate, the asset with the higher Sharpe ratio gives more return for the same risk. Therefore, investors are often advised to pick investments with high Sharpe ratios. From Table 2, we can find that, for seven out of ten global major indices, the Sharpe ratios of our strategies (no matter which look-back duration n is chosen) are much higher than the market, which means that our strategies perform better than the simple buy-and-hold strategy. This result indicates that the JLS model combined with the pattern recognition method has a statistically significant power in systematic detection of rebounds and of crashes in financial markets.

Table 2   Summary of Sharpe ratios for the market and the long-short strategies with different moving average duration n

Index	Strategy (duration n)				Market
	20	30	40	60	
S&P 500	0.32	0.38	0.43	0.39	−0.28
Nasdaq	0.18	0.41	0.48	0.34	−0.11
Russell 2000	0.29	0.3	0.27	0.27	0.03
FTSE 100	−0.07	−0.06	0.01	0.05	−0.22
CAC 40	−0.18	−0.35	−0.37	−0.24	−0.19
SMI	0.28	0.22	0.14	0.13	−0.25
DAX	−0.26	−0.24	−0.27	−0.19	−0.06
Nikkei 225	0.07	0.19	0.39	0.59	−0.33
Hang Seng	0.32	0.38	0.31	0.23	0.06
ASX	−0.38	−0.41	−0.33	−0.24	0.02

Notes: The start and end dates of the strategies are 26-Mar-1999–3-Jun-2009 for S&P 500, 20-Mar-1999–30-Jul-2010 for Nasdaq and 17-Apr-1999–27-Aug-2010 for others. The feature qualification pairs of (7, 100) for both crash and rebound alarm indices are used in this calculation.

The long-short strategies for CAC 40, DAX and ASX perform not as well as the market. However, this is not a statement against the prediction power of our method, but instead supports the evidence that our method detects specific signatures preceding rebounds and crashes that are essentially different from high volatility indicators. Contrary to common lore and to some exceptional empirical cases (Schwert, 1990; Andersen and Sornette, 2004), crashes of the financial markets often happen during low-volatility periods and terminate them. By construction, both our rebound alarm index and the crash alarm index are high during such volatile periods. So if the alarm indices of both types are high at the same time, it is likely that the market is experiencing a highly volatile period. Now, we can refine the strategy and combine the evidence of a directional crash or rebound, together with a high volatility indicator. If we interpret the two co-existing evidences as a signal for a crash, we should ignore these rebound alarm index and take the short strategy mentioned before. As an application, we show the wealth trajectories of DAX based on different type of strategies in Figure 8(a). In the beginning of 2008, both the rebound alarm index and the crash alarm index for DAX are very high, therefore, we detect this period as a highly volatile period and ignore the rebound alarm index. The short strategy gives a very high Sharpe ratio $S = 0.41$ compared to the long-short and short benchmarks where $S = 0.06$. The strategy's average weight is used to compute these benchmarks. These simple benchmarks are constantly invested by a given percentage in the market so that, over the whole time period, they give the same exposure as the corresponding strategy being tested but without the genuine timing information the strategy should contain.

As before, the detailed out-of-sample performances for each sample index (Russell 2000, SMI and Nikkei 225) are also tested. Figures 8(b) to 8(d) illustrate the wealth trajectories for different strategies. In order to show the consistency of the strategies with respect to the chosen parameters in the pattern recognition method, we show the performance of the Russell 2000 index for different qualification pairs: (15, 100) for rebounds and (10,100) for crashes. From these wealth trajectories, it is very obvious that our alarm indices can catch the market rebounds and crashes efficiently.

Figure 8 Performance of trading strategy using our technique, for each figure: (upper) the exposures for different strategies, where the olive solid line represents the long-short strategy, the green dotted line and red dashed line are for long and short strategy respectively, (lower) the historical price and wealth trajectories of the strategies, (a) DAX index 60 days moving average strategy using the feature qualification pair (15,100) for crashes and (10, 200) for rebounds (b) Russell 2000 index 30 days moving average strategy using the feature qualification pair (15,100) for crashes and (10, 100) for rebounds (c) SMI index 20 days moving average strategy (d) Nikkei 225 index 60 days moving average strategy

(a)   (b)   (c)   (d)
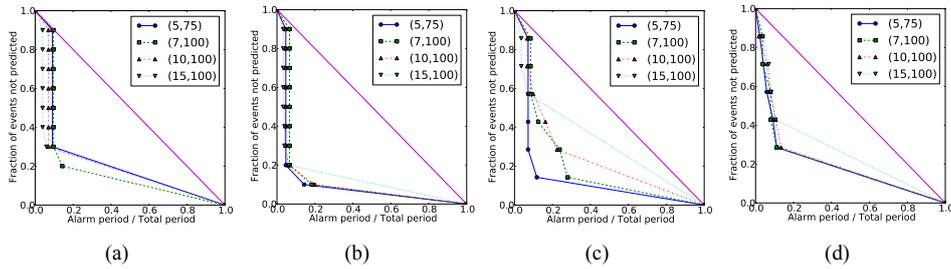
Notes: The blue circles represent the historical price of the index while the others are the wealth trajectories consistent with the upper figure (olive solid – long-short, green dotted – long, red dashed – short). The Sharpe ratios for the strategies are 0.07, 0.47, 0.28 and 0.59 (long-short), –0.66, 0.18, -0.04 and 0.11 (long), and 0.41, 0.26, 0.65 and 0.59 (short). The Sharpe ratio of the corresponding benchmarks, which consist of constant position in the market with exposure equal to the strategy over the whole period, are 0.06, 0.03, 0.2 and 0.33 (long-short), –0.06, 0.03, –0.2 and –0.33 (long), and 0.06, –0.03, 0.2 and 0.33 (short). And the Sharpe ratio of the index in this period is –0.06, 0.03, –0.2 and –0.33. Note that for DAX, the short strategy performs better than long-short or long strategies as discussed in the text.

The detailed performances for these three stock indices are listed in Tables 3 to 6. We also provide the performance of the Russell 2000 index for the 'normal' qualification pair: $(7, 100)$ in Table 4 as a reference. These tables confirm again that strategies mostly succeed in capturing big changes of regime. Compared to the market, the strategies based on our alarm index perform better than the market for more than eleven years in all the important measures: Annual returns and Sharpe ratios are larger, while volatilities, downside deviations and maximum drawdowns are smaller than the market performance.

To summarise, the trading strategies based on the alarm index generated from the JLS model and the pattern recognition method shows better performances compared to the simple buy-and-hold strategy. This again strengthens the validation of our methodology.

Table 3   Russell 2000 index long-short strategies out-of-sample performance table

	Strategy (duration n)				Market
	20	*30*	*40*	*60*	
Ann Ret	5.3%	5.4%	4.8%	4.2%	0.1%
Vol	6.4%	5.9%	5.3%	4.1%	26.3%
Downside dev	4.2%	3.6%	3.4%	2.7%	18.9%
Sharpe	0.42	0.47	0.41	0.37	0.03
Max DD	10%	7%	7%	6%	65%
Abs Expo	16%	15%	14%	12%	
Ann turnover	328%	220%	167%	123%	

Notes: Start date: 17-Apr-1999, end date 27-Aug-2010. Qualification pairs: (15, 100) for
rebounds and (10,100) for crashes.

Table 4   Russell 2000 index long-short strategies out-of-sample performance table

	Strategy (duration n)				Market
	20	*30*	*40*	*60*	
Ann Ret	4.8%	4.7%	4.3%	4.0%	0.1%
Vol	7.9%	7.1%	6.5%	5.2%	26.3%
Downside dev	5.4%	4.8%	4.5%	3.6%	18.9%
Sharpe	0.29	0.30	0.27	0.27	0.03
Max DD	11%	12%	9%	9%	65%
Abs Expo	23%	21%	20%	18%	
Ann turnover	446%	306%	237%	172%	

Notes: Start date: 17-Apr-1999, end date 27-Aug-2010. Qualification pairs: (7, 100) for
both rebounds and crashes.

Table 5   SMI index long-short strategies out-of-sample performance table

	Strategy (duration n)				Market
	20	*30*	*40*	*60*	
Ann Ret	3.4%	3.3%	3.0%	3.0%	−3.4%
Vol	2.6%	2.6%	2.5%	2.3%	20.4%
Downside dev	1.7%	1.6%	1.6%	1.5%	14.6%
Sharpe	0.28	0.22	0.14	0.13	−0.20
Max DD	7%	6%	7%	5%	59%
Abs Expo	9%	9%	9%	9%	
Ann turnover	134%	102%	77%	55%	

Note: Start date 17-Apr-1999, end date 27-Aug-2010.

Table 6 Nikkei 225 index long-short strategies out-of-sample performance table

	Strategy (duration n)				Market
	20	*30*	*40*	*60*	
Ann Ret	3.0%	3.7%	4.8%	5.7%	−8.4%
Vol	6.3%	6.0%	5.6%	5.0%	25.4%
Downside dev	4.4%	4.1%	3.7%	3.2%	18.7%
Sharpe	0.07	0.19	0.39	0.59	−0.33
Max DD	15%	13%	10%	8%	75%
Abs Expo	19%	18%	17%	16%	
Ann turnover	417%	300%	223%	146%	

Note: Start date 17-Apr-1999, end date 27-Aug-2010.

6 Conclusions

We provided a systematic method to detect financial crashes and rebounds. The method is a combination of the JLS model for bubbles and negative bubbles, and the pattern recognition technique originally developed for earthquake predictions. The outcome of this method is a rebound/crash alarm index to indicate the probability of a rebound/crash for a certain time. The predictability of the alarm index has been tested by ten major global stock indices. The performance is checked quantitatively by error diagrams and trading strategies. All the results from error diagrams indicate that our method in detecting crashes and rebounds performs better than chance and confirm that the new method is very powerful and robust in the prediction of crashes and rebounds in financial markets. Our long-short trading strategies based on the crash and rebound alarm index perform better than the benchmarks (buy-and-hold strategy with the same exposure as the average exposure of our strategies) in seven out of ten indices. Highly volatile periods are observed in the indices of which the long-short trading strategy fails to surpass the benchmark. By construction of the alarm index and the fact that highly volatile periods are not coherent with bullish markets, we claim that we should ignore the rebound alarm index during such volatile periods. This statement has been proved by the short strategy which only consider the crash alarm index. Thus, our trading strategies confirm again that the alarm index has a strong ability in detecting rebounds and crashes in the financial markets.

References

Abreu, D. and Brunnermeier, M.K. (2003) 'Bubbles and crashes', *Econometrica*, Vol. 71, No. 1, pp.173–204.

Andersen, J.V. and Sornette, D. (2004) 'Fearless versus fearful speculative financial bubbles', *Physica A*, Vol. 337, Nos. 1–2, pp.565–585.

Bartolozzi, M., Drożdż, S., Leinweber, D., Speth, J. and Thomas, A. (2005) 'Self-similar log-periodic structures in Western stock markets from 2000', *International Journal of Modern Physics C*, Vol. 16, No. 9, pp.1347–1361.

Blanchard, O. and Watson, M. (1983) 'Bubbles, rational expectations and speculative markets', NBER Working Paper 0945, available at http://papers.ssrn.com/abstract=226909 (accessed on 28 February 2012).

Briggs, P.L. and Press, F. (1977) 'Pattern recognition applied to uranium prospecting', *Nature*, Vol. 268, No. 5616, pp.125–127.

Clark, A. (2004) 'Evidence of log-periodicity in corporate bond spreads', *Physica A*, Vol. 338, Nos. 3–4, pp.585–595.

Drożdż, S., Kwapień, J., Oṡwiecimka, P. and Speth, J. (2008) 'Current log-periodic view on future world market development', *Acta Physica Polonica A*, Vol. 114, No. 3, pp.539–546.

Gelfand, I.M., Guberman, S.A., Keilis-Borok, V.I., Knopoff, L., Press, F., Ranzman, E.Y., Rotwain, I.M. and Sadovsky, A.M. (1976) 'Pattern recognition applied to earthquake epicenters in California', *Physics of The Earth and Planetary Interiors*, Vol. 11, No. 3, pp.227–283.

Gnaciński, P. and Makowiec, D. (2004) 'Another type of log-periodic oscillations on Polish stock market', *Physica A*, Vol. 344, Nos. 1–2, pp.322–325.

Jiang, Z-Q., Zhou, W-X., Sornette, D., Woodard, R., Bastiaensen, K., Cauwels, P. (2010) 'Bubble diagnosis and prediction of the 2005–2007 and 2008–2009 Chinese stock market bubbles', *Journal of Economic Behavior and Organization*, Vol. 74, pp.149–162.

Johansen, A. and Sornette, D. (1999a) 'Critical crashes', *Risk*, Vol. 12, No. 1, pp.91–94.

Johansen, A. and Sornette, D. (1999b) 'Financial "anti-bubbles": log-periodicity in gold and Nikkei collapses', *International Journal of Modern Physics C*, Vol. 10, No. 4, pp.563–575.

Johansen, A., Ledoit, O. and Sornette, D. (2000) 'Crashes as critical points', *International Journal of Theoretical and Applied Finance*, Vol. 3, No. 2, pp.219–255.

Johansen, A., Sornette, D. and Ledoit, O. (1999) 'Predicting financial crashes using discrete scale invariance', *Journal of Risk*, Vol. 1, No. 4, pp.5–32.

Kaizoji, T. and Sornette, D. (2009) 'Market bubbles and crashes', in *The Encyclopedia of Quantitative Finance*, Wiley, long version available at http://arxiv.org/abs/0812.2449 (accessed on 28 February 2012).

Keilis-Borok, V., Soloviev, A., Allègre, C., Sobolevskii, A. and Intriligator, M. (2005) 'Patterns of macroeconomic indicators preceding the unemployment rise in Western Europe and the USA', *Pattern Recognition*, Vol. 38, No. 3, pp.423–435.

Matsushita, R., da Silva, S., Figueiredo, A. and Gleria, I. (2006) 'Log-periodic crashes revisited', *Physica A*, Vol. 364, Nos. 1–4, pp.331–335.

Mochan, G.M. (1997) 'Earthquake prediction as a decision making problem', *Pure and Applied Geophysics*, Vol. 149, No. 1, pp.233–247.

Mochan, G.M. and Kagan, Y.Y. (1992) 'Earthquake prediction and its optimization', *Journal of Geophysical Research*, Vol. 97, No. B4, pp.4823–4838.

Preis, T. (2011) 'GPU-computing in econophysics and statistical physics', *European Physical Journal Special Topics*, Vol. 194, pp.87–119.

Schwert, G. (1990) 'Stock volatility and the crash of '87', *Review of Financial Studies*, Vol. 3, No. 1, pp.77–102.

Sornette, D. (1998) 'Discrete scale invariance and complex dimensions', *Physics Reports*, Vol. 297, No. 5, pp.239–270.

Sornette, D. (2003) *Why Stock Markets Crash (Critical Events in Complex Financial Systems)*, Princeton University Press, Princeton, NJ, USA.

Sornette, D. and Zhou, W-X. (2002) 'The US 2000-2002 market descent: how much longer and deeper?', *Quantitative Finance*, Vol. 2, No. 6, pp.468–481.

Sornette, D. and Zhou, W-X. (2006) 'Predictability of large future changes in major financial indices', *International Journal of Forecasting*, Vol. 22, No. 1, pp.153–168.

Sornette, D., Woodard, R. and Zhou, W-X. (2009) 'The 2006-2008 oil bubble: evidence of speculation and prediction', *Physica A*, Vol. 388, No. 8, pp.1571–1576.

Sornette, D., Woodard, R., Fedorovsky, M., Reimann, S., Woodard, H. and Zhou, W-X. (2010a) 'The financial bubble experiment: advanced diagnostics and forecasts of bubble terminations (the financial crisis observatory)', available at http://arxiv.org/abs/0911.0454 (accessed on 28 February 2012).

Sornette, D., Woodard, R., Fedorovsky, M., Reimann, S., Woodard, H. and Zhou, W-X. (2010b) 'The financial bubble experiment: advanced diagnostics and forecasts of bubble terminations volume II (master document)', available at http://arxiv.org/abs/1005.5675 (accessed on 28 February 2012).

van Bothmer, H-C.G. and Meister, C. (2003) 'Predicting critical crashes? A new restriction for the free variables', *Physica A*, Vol. 320, Nos. 1–4, pp.539–547.

Vandewalle, N., Ausloos, M., Boveroux, P. and Minguet, A. (1999) 'Visualizing the log-periodic pattern before crashes', *European Physics Journal B*, Vol. 9, No. 2, pp.355–359.

Watanabe, K., Takayasu, H. and Takayaso, M. (2007) 'A mathematical definition of the financial bubbles and crashes', *Physica A*, Vol. 383, No. 1, pp.120–124.

Woodard, R., Sornette, D. and Fedorovsky, M. (2011) 'The financial bubble experiment: advanced diagnostics and forecasts of bubble terminations volume III (beginning of experiment + post-mortem analysis)', available at http://arxiv.org/abs/1011.2882 (accessed on 28 February 2012).

Yan, W., Woodard, R. and Sornette, D. (2010) 'Diagnosis and prediction of tipping points in financial markets: crashes and rebounds', *Physics Procedia*, Vol. 3, No. 5, pp.1641–1657.

Yan, W., Woodard, R. and Sornette, D. (2012a) 'Diagnosis and prediction of rebounds in financial markets', *Physica A*, Vol. 391, No. 4, pp.1361–1380, available at http://arxiv.org/abs/1001.0265.

Yan, W., Woodard, R. and Sornette, D. (2012b) 'Leverage bubble', *Physica A*, Vol. 391, Nos. 1–2, pp.180–186.

Zhou, W-X. and Sornette, D. (2003) '2000-2003 real estate bubble in the UK but not in the USA', *Physica A*, Vol. 329, Nos. 1–2, pp.249–263.

Zhou, W-X. and Sornette, D. (2006) 'Is there a real estate bubble in the US?', *Physica A*, Vol. 361, No. 1, pp.297–308.

Zhou, W-X. and Sornette, D. (2008) 'Analysis of the real estate market in Las Vegas: bubble, seasonal patterns, and prediction of the CSW indexes', *Physica A*, Vol. 387, No. 1, pp.243–260.

Zhou, W-X. and Sornette, D. (2009) 'A case study of speculative financial bubbles in the South African stock market 2003-2006', *Physica A*, Vol. 388, No. 6, pp.869–880.