

Switching processes in financial markets

Tobias Preis^{a,b,c,1}, Johannes J. Schneider^d, and H. Eugene Stanley^{a,1}

^aCenter for Polymer Studies and Department of Physics, Boston University, 590 Commonwealth Avenue, Boston, MA 02215; ^bETH Zurich (Swiss Federal Institute of Technology), Chair of Sociology, in particular of Modeling and Simulation, CLU E 5, Clausiusstrasse 50, 8092 Zurich, Switzerland; ^cArtemis Capital Asset Management GmbH, Gartenstrasse 14, 65558 Holzheim, Germany; and ^dDepartment of Physics, Mathematics, and Computer Science, Johannes Gutenberg University, Staudingerweg 7, 55128 Mainz, Germany

Contributed by H. Eugene Stanley, December 30, 2010 (sent for review October 24, 2010)

For an intriguing variety of switching processes in nature, the underlying complex system abruptly changes from one state to another in a highly discontinuous fashion. Financial market fluctuations are characterized by many abrupt switchings creating upward trends and downward trends, on time scales ranging from macroscopic trends persisting for hundreds of days to microscopic trends persisting for a few minutes. The question arises whether these ubiquitous switching processes have quantifiable features independent of the time horizon studied. We find striking scale-free behavior of the transaction volume after each switching. Our findings can be interpreted as being consistent with time-dependent collective behavior of financial market participants. We test the possible universality of our result by performing a parallel analysis of fluctuations in time intervals between transactions. We suggest that the well known catastrophic bubbles that occur on large time scales—such as the most recent financial crisis—may not be outliers but single dramatic representatives caused by the formation of increasing and decreasing trends on time scales varying over nine orders of magnitude from very large down to very small.

complex systems | econometrics | switching phenomena | phase transitions | futuriCT

The study of dramatic crash events is limited by the fortunately rare number of such events. However, there is a truly gargantuan amount of preexisting precise financial market data already collected, many orders of magnitude more than for other complex systems. Accordingly, financial markets are becoming a paradigm of complex systems (1, 2), and increasing numbers of scientists are analyzing market data (3–18) and modeling financial markets (19–29). The probability distribution function and the time autocorrelation function reveal interesting features, such as long-range power-law correlations in volatility (30) and fat tails in the price change probability distribution function (31, 32).

Increasingly, one seeks to understand the current financial crisis by comparisons with the depression of the 1930s. Here we ask if the smaller financial crises also provide information of relevance to large crises. If such “microbubbles” are relevant, then the larger abundance of data on smaller crises should provide quantifiable statistical laws for bubble formation and financial collapse on various scales.

Data Analyzed

To answer this question, we perform parallel analyses of trend switching on two quite different time scales: (i) from ≈ 10 ms to $\approx 10^6$ ms, and (ii) from $\approx 10^8$ ms to $\approx 10^{10}$ ms.

- German market: For the first analysis, we use a price time series of the German DAX Future (FDAX) traded at EUREX, which is one of the world's largest derivatives exchanges. The time series comprises $T_1 = 13,991,275$ trades of three disjoint three-month periods (March 16, 2007–June 15, 2007; June 20, 2008–September 19, 2008; and September 19, 2008–December 19, 2008). The data base contains the transaction prices, the volumes, and the corresponding time stamps (33–36), with a large liquidity and intertrade times—time intervals between consecutive transactions—down to 10 ms, which allows us to perform an analysis of microtrends (Fig. 1A).

- US market: For the second analysis, which focuses on macro-trends, we use price time series of daily closing prices of all stocks of the S&P500 index. The time series comprises overall $T_2 = 2,592,531$ closing prices. Our oldest closing prices date back to January 2, 1962. Latest closing prices were recorded on June 16, 2009. The data base of closing prices we analyze contains the daily closing prices and the daily cumulative trading volumes.

In addition, we perform a parallel analysis of the 30 assets contributing to the Dow Jones Industrial Average (DJIA)*. The time series stem from the Trade and Quote (TAQ)[†] data base and cover 2,623,445,866 transactions.

Renormalization Method

To analyze switching processes of financial fluctuations, we first propose how a switching process can be quantitatively analyzed. Let $p(t)$ be the transaction price of trade t , which is a discrete variable $t = 1, \dots, T$. A transaction price $p(t)$ is defined to be a local maximum of order Δt if there is no higher transaction price in the interval $t - \Delta t \leq t \leq t + \Delta t$, and is defined to be a local minimum of order Δt if there is no lower transaction price in this interval (Fig. 1B).

Here, we perform an analysis of the volume fluctuations $v(t)$ from one price extremum to the next. The volume is the number of contracts traded in each individual transaction in case of microtrends for the German market and the number of traded stocks per day in case of macro-trends for the US market. For the analysis, we introduce a renormalized time scale ϵ between successive extrema (37). Thus, $\epsilon = 0$ corresponds to the beginning of a trend and $\epsilon = 1$ indicates the end of a trend (Fig. 1C). We analyze a range of ϵ for the interval $0 \leq \epsilon \leq 2$, so we can analyze trend switching processes both before as well as after the critical value $\epsilon = 1$ (Fig. 1). The renormalization is essential to assure that trends of various lengths can be aggregated and that all switching points have a common position in the renormalized time.

Results of Analysis

Fig. 2A provides the volume $v^*(\epsilon)$ averaged over all increasing and decreasing microtrends in the full time series of $T_1 = 13,991,275$ records, and is normalized by the average volume of all microtrends studied. In order to remove outliers, e.g., overnight gaps, only these microtrends are collected in which the time intervals between successive trades $\tau(t)$ (38) are not longer than

Author contributions: T.P., J.J.S., and H.E.S. designed research; T.P., J.J.S., and H.E.S. performed research; T.P., J.J.S., and H.E.S. analyzed data; and T.P., J.J.S., and H.E.S. wrote the paper.

The authors declare no conflict of interest.

*We analyze all US stocks which were included in the DJIA in March 2010. This data base covers the time period from January 4, 1993 to December 31, 2009.

[†]The New York Stock Exchange's TAQ database contains data on all trades and quotes for stocks from The New York Stock Exchange (NYSE), The American Stock Exchange (AMEX) and The National Association of Securities Dealers Automated Quotation (NASDAQ).

¹To whom correspondence may be addressed. E-mail: mail@tobiaspreis.de or hes@bu.edu.

This article contains supporting information online at www.pnas.org/lookup/suppl/doi:10.1073/pnas.1019484108/-DCSupplemental.

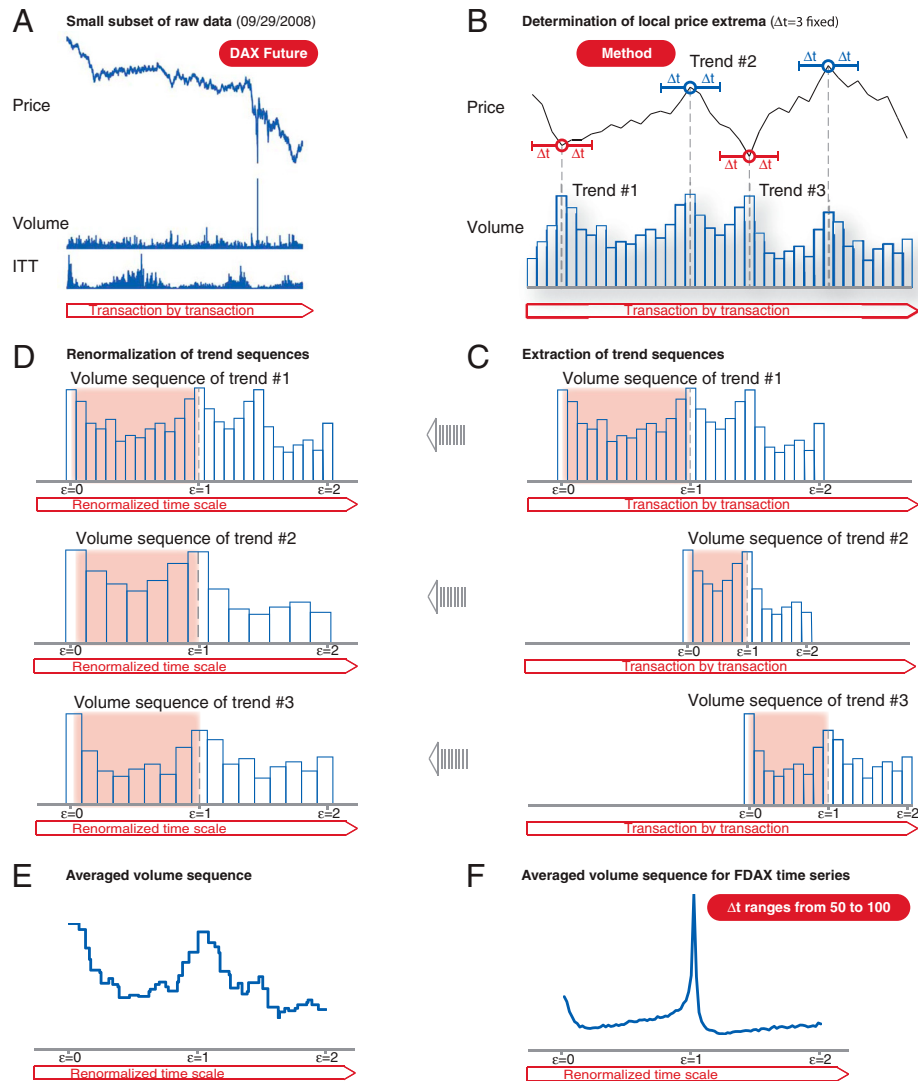


Fig. 1. Segregation and rescaling of trend sequences in a multivariate time series in order to analyze financial market quantities on the path from one price extremum to the next. (A) Small subset comprising 121,400 transactions of the full dataset (13,991,275 transactions) analyzed, extracted from the German DAX future (FDAX) time series which provides transaction prices, transaction volumes, and time intervals between transaction—intertransaction times (ITT). This subset recorded on September 29, 2008 documents the volatile reaction of stock markets as the US government's \$700 billion financial bailout plan was rejected by the House of Representatives on that day. (B) Schematic visualization of trend segregation for $\Delta t = 3$. Positive trends start at local price minima (red circles) and end at local maxima (blue circle)—and vice versa. A transaction price $p(t)$ is a local maximum if there is no higher transaction price in the interval $t - \Delta t \leq t \leq t + \Delta$. Analogously, $p(t)$ is a local minimum if there is no lower transaction price in the interval $t - \Delta t \leq t \leq t + \Delta$. (C) Segregated sequences of transaction volumes belonging to the three trends identified in (B). We assign $\epsilon = 0$ to the start of each trend, and $\epsilon = 1$ to the end of each trend. In order to study trend switching processes—both before as well as after the end of a trend—we consider additionally the subsequent volume sequences of identical length. (D) Visualization of the volume sequences in the renormalized time scale. The renormalization assures that trends of various lengths can be aggregated as all switching points have a common position in this renormalized scale. (E) Averaged volume sequence derived from the summation of the three trend sequences. (F) Average volume sequence $v^*(\epsilon)$ for all trends in the full FDAX time series derived from summation over various values of Δt . Extreme values of the price coincide with peaks in the time series of the volumes.

1 min, which is roughly 60 times longer than the average intertrade time (≈ 0.94 s). Furthermore, transaction volumes have not to be larger than 100 contracts (the average transaction volume is 2.55 contracts). As expected, new price extrema are linked with peaks in the volume time series. In Fig. 2D, we show the averaged volume $v^*(\epsilon)$ vs. $|\epsilon - 1|$ as a log—log histogram. Surprisingly, the averaged volume falls on straight lines and thus indicates a power-law scaling behavior of the form

$$v^*(|\epsilon - 1|) \sim |\epsilon - 1|^{\beta_v} \quad [1]$$

with scaling parameters $\beta_v^- = -0.068 \pm 0.001$ (t -test, p -value $< 2 \times 10^{-16}$) before, and $\beta_v^+ = -0.155 \pm 0.004$ (t -test, p -value $= 9.2 \times 10^{-16}$) after a price extremum. Such an extraction of slopes

by performing least-squares linear regressions is not sufficient for the claim that the averaged volume is drawn from a power-law distribution. However, additional performed statistical tests (see *SI Appendix*) enable us to conclude that our observations are indeed consistent with the hypothesis that v^* is drawn from a power-law distribution.

Next we test the possible universality of our result by performing a parallel analysis for trends on long time scales using the daily closing price data base of S&P500 stocks. Note that for our parallel analysis on macroscopic time scales the order of a extremum Δt is measured in units of days, and that $v^*(\epsilon)$ is averaged over all trends and all closing price time series of all S&P500 components. In order to avoid biased contributions for the rescaled averaging caused by inflation based drifts over more than 47 y, the

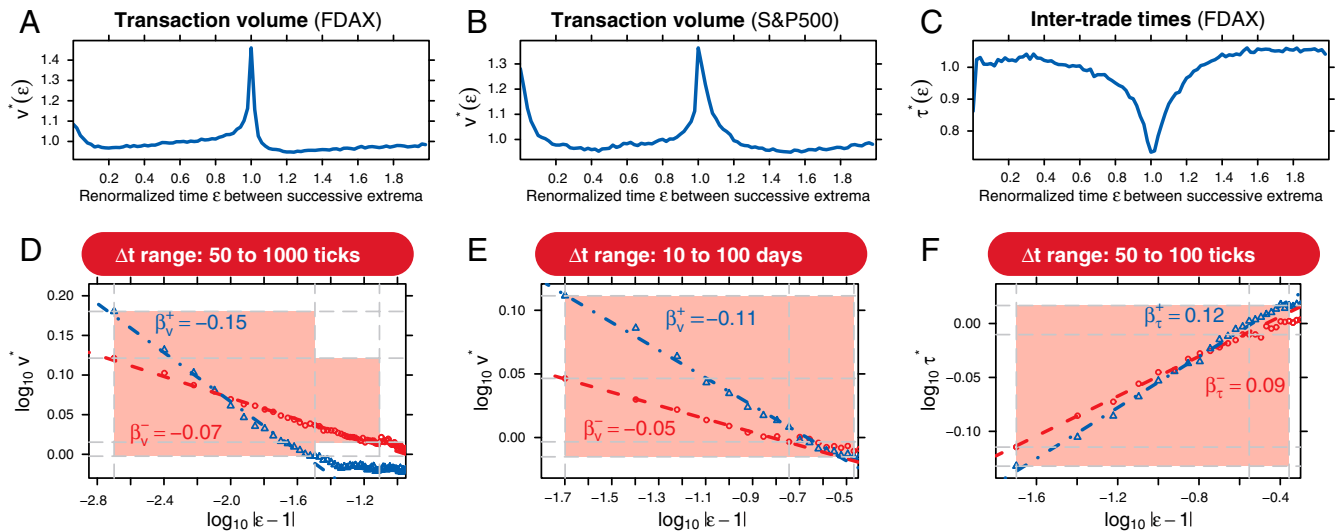


Fig. 2. Renormalization time analysis and log–log plots of quantities with scale-free properties. (A) Averaged volume sequence $v^*(\epsilon)$ of the German DAX Future time series. Δt ranges from 50 to 100 transactions (ticks). Extreme values of the price coincide with sharp peaks in the volume time series. (B) A very similar behavior is obtained for the averaged volume sequence $v^*(\epsilon)$ of S&P500 stocks. Here, Δt ranges from 10 days to 100 days. (C) Averaged intertrade time sequence $\tau^*(\epsilon)$ of the German DAX Future time series. Extreme values of the price time series are reached with a significant decay of intertrade times (50 ticks $\leq \Delta t \leq 100$ ticks). Averaged volume and averaged intertrade time sequences are asymmetric for two main reasons: Only $\epsilon = 0$ and $\epsilon = 1$ correspond to extreme values in the price time series. In order to study the behavior before and after a trend switching point, we extend individual sequences from $\epsilon = 0$ to $\epsilon = 1$ by an identical amount of transactions or intertrade times. In addition, the end of a trend, $\epsilon = 1$, does not necessarily correspond to the starting point of the next trend due to filter criteria mentioned in the main text. (D) Log–log plot of the FDAX transaction volumes (50 ticks $\leq \Delta t \leq 1,000$ ticks) before reaching an extreme price value ($\epsilon < 1$, circles) and after reaching an extreme price value ($\epsilon > 1$, triangles). The straight lines correspond to power-law scaling with exponents $\beta_v^+ = -0.155 \pm 0.004$ (t -test, p -value = 9.2×10^{-16}) and $\beta_v^- = -0.068 \pm 0.001$ (t -test, p -value $< 2 \times 10^{-16}$). The shaded intervals mark the region in which the empirical data are consistent with a power-law behavior. The left border of the shaded regions is given by the first measuring point closest to the switching point. The right borders stem from statistical tests of the power-law hypothesis (see SI Appendix). (E) Log–log plot of the transaction volumes shown in (B) indicates a power-law behavior with exponents $\beta_v^+ = -0.109 \pm 0.003$ (t -test, p -value $< 2 \times 10^{-16}$) and $\beta_v^- = -0.052 \pm 0.001$ (t -test, p -value = 1.7×10^{-9}) which are similar to our results on short time scales. (F) Log–log plot of the intertrade times on short time scales (50 ticks $\leq \Delta t \leq 100$ ticks) exhibits a power-law behavior with exponents $\beta_\tau^+ = 0.118 \pm 0.002$ (t -test, p -value $< 2 \times 10^{-16}$) and $\beta_\tau^- = 0.092 \pm 0.002$ (t -test, p -value = 1.8×10^{-15}). An equivalent analysis on long time scales is not possible as daily closing prices are recorded with equidistant time steps.

analyzed price time series $p(t)$ contains the logarithm of the daily closing prices. A log–log histogram of our parallel analysis for the US market on large time scales (Fig. 2 B and E) provides evidence for a similar behavior with scaling parameters $\beta_v^- = -0.052 \pm 0.001$ (t -test, p -value = 1.7×10^{-9}) before, and $\beta_v^+ = -0.109 \pm 0.003$ (t -test, p -value $< 2 \times 10^{-16}$) after a price extremum. Statistical tests confirm the consistency with a power-law distribution.

In order to verify a possible universality, we analyze the behavior of the intertrade times $\tau(t)$ of the German market during the short time interval from one price extremum to the next. The cross-correlation function between price changes and intertrade times exhibits no reasonable correlation values as well. Thus, one can conjecture that the tendency to decreased intertrade times for the end of positive microtrends is counteracted by the tendency to decreased intertrade times for the end of negative microtrends. The crucial issue is to distinguish between positive and negative microtrends realized by the renormalized time ϵ between successive extrema. In Fig. 2C, the averaged intertrade times $\tau^*(\epsilon)$ reflects the link between intertrade times and price extrema. Fig. 2F shows $\tau^*(\epsilon)$ vs. $|\epsilon - 1|$ as a log–log histogram supporting a power-law behavior of the form

$$\tau^*(|\epsilon - 1|) \sim |\epsilon - 1|^{\beta_\tau} \quad [2]$$

with scaling parameters $\beta_\tau^- = 0.092 \pm 0.002$ (t -test, p -value = 1.8×10^{-15}) before, and $\beta_\tau^+ = 0.118 \pm 0.002$ (t -test, p -value $< 2 \times 10^{-16}$) after a price extremum. Statistical tests confirm the consistency with a power-law distribution as well.

The straight lines in Fig. 2 D–F offer insight into financial market fluctuations: (i) a clear connection between volumes, intertrade times, and price fluctuations on the path from one

extremum to the next extremum, and (ii) the underlying law, which describes the volumes and intertrade times around extrema varying over nine orders of magnitude starting from the smallest possible time scale, is a power-law with scaling parameters which quantitatively characterizes the region around the trend switching point. As a direct consequence of the consistency with power-law distributions, the behavior does not depend on the scale. Thus, we find identical behavior for other subintervals of $50 \leq \Delta t \leq 1,000$. These findings can contribute to the understanding of mechanisms causing catastrophic events. However, one should be aware of that our findings cannot be used for predicting individual transaction sequences due to their level of noise.

Confirmation Based on 2,623,445,866 Transactions

We perform a parallel analysis of the 30 assets contributing to the Dow Jones Industrial Average (DJIA). The time series stem from the Trade and Quote (TAQ) database and cover 2,623,445,866 transactions, roughly 200 times the length of the FDAX time series analyzed before.

The TAQ database makes it possible to study the switching phenomenon of transaction volume on an intraday time scale for all 30 stocks individually*. For the sake of clarity and presentation, results for all stocks considered can be found in SI Appendix. For each stock, we document the development of the stock price, the aggregated volume as a function of ϵ , and $v^*(\epsilon)$ vs. $|\epsilon - 1|$ as a log–log histogram with two power-law fits. The fitting exponents are reported in the upper right corner. Each fitting range is based on statistical tests identical to those applied in the previous analyses: The shaded intervals mark

*The quality of the time stamp is not comparable to that of the FDAX time series. Thus, only volume is studied.

the region in which the empirical data are consistent with a power-law behavior. The left border of the shaded regions is given by the first measuring point closest to the switching point. The right borders stem from statistical tests of the power-law hypothesis.

As is evident in the presented results, the *DJIA* components confirm our previous findings and provide evidence for power-law exponents between $\beta_v^- = -0.06$ and $\beta_v^- = -0.01$ before the switching point $\varepsilon = 1$ and between $\beta_v^+ = -0.46$ and $\beta_v^+ = -0.12$ after the switching point $\varepsilon = 1$.

Parallels to a Phase Transition in Physics

The straight lines in Fig. 2 *D–F* offer insight into the nature of financial market fluctuations: (i) We uncover a clear relationship between transaction volumes and price fluctuations on the path from one extremum to the next extremum. In addition, the analysis of intertrade times provides confirmation for our finding. However, it is noteworthy that transaction volumes and intertrade times are not directly linked to each other. The transaction volume per individual transaction increases dramatically before an extreme price value is reached. Parallely, time intervals between individual transactions decreases dramatically before an extreme price value is reached. (ii) We find the underlying law—a power-law—which describes transaction volumes and intertrade times around price extrema varying over nine orders of magnitude starting from the smallest possible time scale of individual transactions. This power-law with unique exponents quantitatively characterizes the region around the trend switching point. At this point a positive trend ends and a negative trend starts or a negative trend ends and a positive trend starts. As a direct consequence of the existence of power-law behavior, the phenomenon does not depend on the scale. Thus, we find an identical behavior for other subintervals of $50 \leq \Delta t \leq 1,000$. With a decreasing value of Δt , the number of local minima and maxima increases in the price time series, around which we find the same scale-free behavior. The finding of a power-law behavior of transaction volumes and intertrade times in the time domain of individual transactions supports the hypothesis that a fluctuating price time series passes through a sequence of transitions similar to phase transitions in physics.

Based on our findings we may ask what kind of transition could the end of a trend—a microtrend on scales of 10 ms or a macrotrend on scales of 100 d—correspond to, or is the end of a trend an altogether different kind of phase transition that resembles all phase transitions by displaying a regime of scale-free behavior characterized by a critical exponent? It may be premature to speculate on possible analogies, so we will limit ourselves to describing here what seems to be a promising candidate. Consider a simple Ising magnet characterized by one-dimensional spins that can point North or South. Each spin interacts with some (or even with all) of its neighbors with positive interaction strength J , such that, when J is positive neighboring spins lower their energy by being parallel. The entire system is bathed in a magnetic field that interacts with all the spins equally with a strength parametrized by H , such that when H is positive, the field points North and when H is negative the field points South. Thus, when H is positive, the system lowers its energy by each spin pointing North. Thus, there are two competing control parameters J and H . If, e.g., the system is prepared in a state with the majority of spins pointing North yet the field H points South, the competition will be between the relative effects of J and H : the J interaction motivates the spins to point North but the H interaction motivates the spin to point South. Such a system is termed *metastable* because, if each North-pointing spin suddenly flips its state to point South, the system can achieve a lower total energy. This flipping will occur in time

in a fashion not unlike the trading frequency near the end of a trend: first, one or two spins will randomly switch their state, then more, and suddenly, after in an “avalanche” of switches, the majority of spins will point South. The phase transition is termed a spinodal singularity, characterized by its own set of exponents.

Why should the end of microtrends or macrotrends have a parallel with such a metastable physical system? Presumably, near the end of a positive trend, all the market participants watching the market begin to sense that the market is metastable and that, if they do not sell soon, it could be too late to make any profit, because the price will drop. First, a few traders sell, pushing the market imperceptibly lower. Then, additional traders, sensing this microscopic downturn, may decide that now is the time to sell and they sell too. Then an “avalanche” of selling begins, with traders all hoping to protect their profits by selling before the market drops. Thus, the set of N market participants holding their position are in this sense analogous to the set of N mostly North-pointing spins, bathed in a South-pointing magnetic field.

Perhaps, the above analogy is not yet the best. It will be a future challenge to find a coherent, convincing explanation for why the end of a microtrend or macrotrend displays such striking parallels to a phase transition. In any case, a set of interacting spins appears to be analogous to a set of interacting traders. The end of the negative microtrend or macrotrend follows the same mechanism, but with everything reversed. The N Ising spins point mostly South, the magnetic field points North, and the spins flip from South to North one by one, ending with an avalanche corresponding to the spinodal singularity. Analogously, the N traders begin to suspect that the market is becoming metastable, so that they start to buy one by one. But as the traders witness the price increasing, they jump into buy mode before the price becomes too high.

Summary

In summary, we have seen that each trend—microtrend and macrotrend—in a financial market starts and ends with a unique switching process, and each extremum shares properties of macroscopic cooperative behavior (39–42). We have seen that the switching mechanism has no scale, for time scales varying over nine orders of magnitude down to the smallest possible time scale—the scale of single transactions measured in units of 10 ms. Thus, the well known catastrophic bubbles occurring on large time scales—such as the most recent financial crisis—may not be outliers but in fact single dramatic events caused by the inherent, scale-free behavior related to the formation of increasing and decreasing trends on time scales from the very large down to the very small. The larger abundance of data on smaller crises can provide quantifiable statistical laws for bubble formation and financial collapse on various scales.

Materials and Methods

Test of the Power-Law Hypothesis. Description of the statistical test confirming consistency with power-law distributions. Details are given in *SI Appendix*.

Switching Process Analysis for *DJIA* Components. Parallel analysis of the 30 assets contributing to the *Dow Jones Industrial Average (DJIA)*. Details are given in *SI Appendix*.

ACKNOWLEDGMENTS. The authors thank K. Binder, S. V. Buldyrev, C. De Grandi, S. Havlin, D. Helbing, U. Krey, H.-G. Matuttis, M. G. Mazza, I. Morgenstern, W. Paul, R. H. R. Stanley, T. Vicsek, G. M. Viswanathan, and L. Yelash for discussions. T.P. also thanks the Gutenberg Academy for financial support. This work was partially supported by the German Research Foundation Grants SCHN 1073/1-1 (to J.J.S.) and PR 1305/1-1 (to T.P.) and by the National Science Foundation (NSF) and Office of Naval Research (ONR), and the Defense Threat Reduction Agency (DTRA).

1. Axtell RL (2001) Zipf distribution of U.S. firm sizes. *Science* 293:1818–1820.
2. Stanley MHR, et al. (1995) Zipf plots and the size distribution of firms. *Economics Letters* 49:453–457.

3. Gabaix X, Gopikrishnan P, Plerou V, Stanley HE (2006) Institutional investors and stock market volatility. *Quarterly Journal of Economics* 121:461–504.

