








Time Series Analysis

Chapter 13: Speculative bubbles: LPPL models



Daniel Traian PELE

Bucharest University of Economic Studies
IDA Institute Digital Assets
Blockchain Research Center
AI4EFin Artificial Intelligence for Energy Finance
Romanian Academy, Institute for Economic Forecasting
MSCA Digital Finance

Acronyms used in this material

- **ADF**: Augmented Dickey–Fuller (test)
- **AI**: Artificial Intelligence
- **AR**: AutoRegressive (model)
- **BET**: Bucharest Exchange Trading (index)
- **BET-TR**: BET Total Return (index)
- **BSADF**: Backward Supremum Augmented Dickey–Fuller (statistic)
- **EODHD**: EOD Historical Data (market-data provider)
- **GARCH**: Generalised AutoRegressive Conditional Heteroskedasticity
- **GSADF**: Generalised Supremum Augmented Dickey–Fuller (test)
- **JLS**: Johansen–Ledoit–Sornette (model)
- **LPPL**: Log-Periodic Power Law
- **LPPLS**: Log-Periodic Power Law Singularity (model)
- **OLS**: Ordinary Least Squares
- **SADF**: Supremum Augmented Dickey–Fuller (test)
- **SE**: Standard Error
- **SSR**: Sum of Squared Residuals

The question of the chapter

- **Question**: can a speculative bubble be recognised in a price series before it bursts?
  - a **speculative bubble**: a price that rises far above the value justified by the expected cash flows of the asset, sustained by the expectation of selling it later at a higher price
  - a **crash**: a large fall of the price in a short time (here: at least 20% from the peak)
- **Route** of the chapter
  - bubbles in history and the statistical signatures of a bubble
  - rational bubbles and explosive roots: the right-tailed unit-root tests of Phillips, Shi and Yu (link with Chapter 3)
  - the LPPL model of Johansen, Ledoit and Sornette: super-exponential growth, log-periodic oscillations, the critical time
  - estimation, confidence from many windows and the evaluation of crash predictions

Self-study guide

- This chapter is for **self-study**
  - there is no seminar
  - each section ends with a recap
  - read a section, then redo its worked example on paper before you look at the solution
  - after each chart, read the interpretation slide and compare it with your own reading of the chart
- Run the lecture notebook: every chart and number of the slides is produced there
  - change one setting (the window, the filter, the threshold) and see how the conclusion changes
- Check yourself with the self-assessment at the end and with the quiz of Chapter 13 on the course site
- Prerequisites: log returns and log prices (Chapter 1), the Dickey–Fuller test (Chapter 3), OLS regression

Learning outcomes

- Describe a bubble by its fundamental value and by its statistical signatures
- Explain why a rational bubble must grow at an explosive rate and test for explosive roots with SADF and GSADF
- Write the LPPL equation and interpret each of its seven parameters
- Estimate LPPL with the two-step method of Filimonov and Sornette and check the filter conditions
- Build a confidence indicator from many windows and evaluate its alarms honestly, with false alarms and base rates

Reading and tools

- Textbook background: Huang and Petukhina (2022) (unit roots and nonlinear least squares, the tools of this chapter)
  - the LPPL model: Johansen, Ledoit and Sornette (2000); Sornette (2017); calibration: Filimonov and Sornette (2013)
  - explosive roots: Phillips, Wu and Yu (2011); Phillips, Shi and Yu (2015a); survey: Phillips and Shi (2020)
  - history of bubbles: Kindleberger and Aliber (2005); Shiller (2015); Garber (1990)
- Python Quantlets of this chapter: Quantlets/Ch_13
  - the SADF/GSADF statistics and the LPPL calibration are written out in the code (no special package)
- Lecture notebook: open in Google Colab
- Data: daily closes of the S&P 500, Nasdaq 100, BET, Shanghai Composite and Bitcoin (EODHD)

Speculative bubbles, 1636–2021



Flora's wagon of fools, c. 1637

Painting: Hendrik Gerritsz Pot (c. 1637); public domain; Wikimedia Commons

- **Tulip mania**, Netherlands, 1636–1637
  - contracts on rare bulbs traded at extreme prices, then collapsed in February 1637
  - Garber (1990): part of the story is a myth; rare bulbs were expensive for good reasons, common bulbs show the mania
- **South Sea** (1720), **Wall Street** (1929), **Japan** (1989), **dot-com** (2000), **housing** (2007), **crypto** (2017, 2021)
- The common pattern Kindleberger and Aliber (2005): a new story, easy credit, rising prices that attract new buyers, then panic

1929 and 1987



Crowd outside the New York Stock Exchange, 29 October 1929

Photo: US government (1929); public domain; Wikimedia Commons



Amsterdam Stock Exchange, 21 October 1987

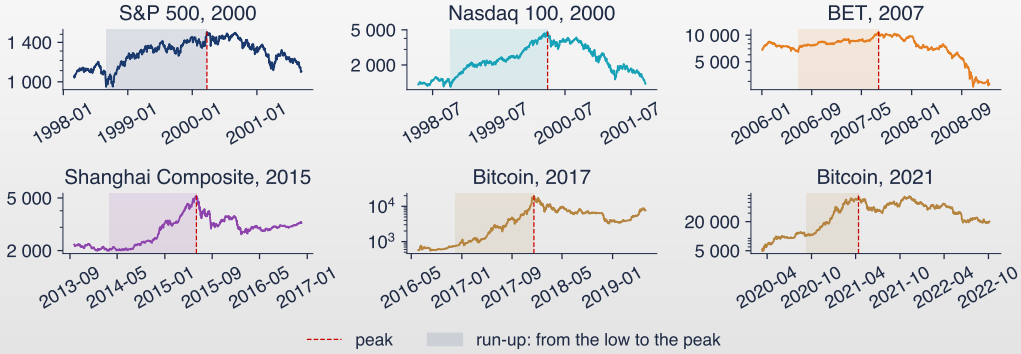
Photo: Bart Molendijk / Anefo, Nationaal Archief (1987); CC0; Wikimedia Commons

- **October 1929**
  - the end of the boom of the 1920s and the start of the Great Depression
- **19 October 1987** (Black Monday)
  - the largest one-day fall of the Dow Jones index, without any major news that day
  - Sornette, Johansen and Bouchaud (1996) found accelerating, log-periodic oscillations in the years before 1987; this observation started the LPPL literature
  - the daily data of this course start in 1990: we study 1987 through the literature and later bubbles with our own data
- A crash without news suggests an **endogenous** cause: the market itself became fragile

Fundamental value and bubble

- **Fundamental value** F_t
  - the present value of the expected future cash flows (dividends D), discounted at the rate r
  - $F_t = \sum_{k=1}^{\infty} \frac{E_t[D_{t+k}]}{(1+r)^k}$
  - D_{t+k}: the dividend paid k periods from now; $E_t[\cdot]$: the expectation given the information available at t; r: the discount rate per period
- **Bubble component**
  - $B_t = P_t - F_t$, the part of the price P_t not explained by fundamentals
  - F_t is not observed: we can only test the **implications** of a bubble in the price path
- Two views
  - efficient markets: prices reflect information; "bubbles" are visible only with hindsight Fama (2014)
  - behavioural finance: herding and extrapolation push prices away from value Shiller (2015); large run-ups do raise the probability of a crash Greenwood, Shleifer and You (2019)

Six run-ups and crashes in the course data



- Daily closes, log scale, from six months before the low of the run-up to eighteen months after the peak; shaded: the run-up from the low to the peak

TSA_ch13_episodes

Interpreting the six episodes

Episode	Low	Peak	Rise	Growth p.a.	Fall in 1 year
S&P 500, 2000	31 August 1998	24 March 2000	60%	30%	27%
Nasdaq 100, 2000	8 October 1998	27 March 2000	317%	97%	66%
BET, 2007	27 June 2006	24 July 2007	68%	48%	50%
Shanghai Composite, 2015	20 March 2014	12 June 2015	159%	77%	49%
Bitcoin, 2017	1 December 2016	16 December 2017	2476%	312%	83%
Bitcoin, 2021	8 September 2020	13 April 2021	527%	309%	53%

- Growth p.a.: $\ln(P_{\mathbf{peak}}/P_{\mathbf{low}})$ divided by the length of the run-up in years; fall: the lowest close in the year after the peak, relative to the peak
- Bitcoin multiplied its price 26 times in about a year; the S&P 500 rose much less: not every peak is a bubble of the same strength

Recap: Bubbles in history

- A bubble is a price above the fundamental value F_t; F_t is not observed
- Six episodes in our data: run-ups of 60% to 2476%, followed by falls of 27% to 83% within a year
- Signatures: explosive and super-exponential growth, accelerating oscillations, then an end

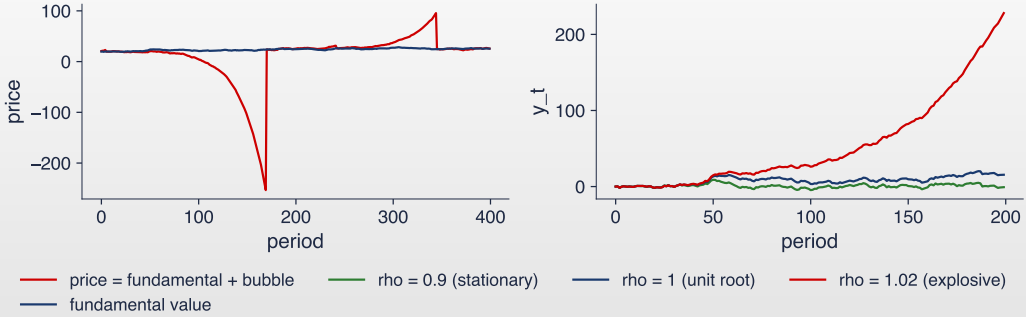
Why a rational bubble must explode

- No-arbitrage pricing: $P_t = \dfrac{E_t[P_{t+\mathbf{1}} + D_{t+\mathbf{1}}]}{1+r}$; the fundamental value F_t is one solution
  - in words: today's price is the discounted expected value of tomorrow's price plus tomorrow's dividend
  - every $P_t = F_t + B_t$ with $E_t[B_{t+\mathbf{1}}] = (1+r)\,B_t$ is also a solution
- The bubble must grow, in expectation, at the rate r: $E_t[B_{t+k}] = (1+r)^k B_t$
  - an investor holds an asset above its value only if the overvaluation is expected to grow
  - $1 + r > 1$: the bubble component is an **explosive** process
- Consequence: if dividends are $I(1)$ (Chapter 3), the price is $I(1)$ without a bubble and explosive with one

Worked example: a bubble that can burst

- Blanchard and Watson (1982): each period the bubble survives with probability π and then grows to $\frac{1+r}{\pi} B_t$; otherwise it bursts ($B_{t+1} \approx 0$)
  - check: $E_t[B_{t+1}] = \pi \cdot \frac{1+r}{\pi} B_t + (1-\pi) \cdot 0 = (1+r) B_t$
- With $r = 2\%$ and $\pi = 0.98$ per period:
  - while it survives the bubble grows by $1.02/0.98 - 1 = 4.08\%$ per period, more than r: the extra growth pays for the risk of a burst
  - expected lifetime $1/(1-\pi) = 50$ periods; it survives more than 34 periods with probability $1/2$ ($0.98^k = 0.5$)
- The longer the bubble lasts, the faster it must grow and the larger the fall when it bursts

A simulated rational bubble and three AR(1) roots



- Left: fundamental value (random walk) plus a Blanchard–Watson bubble ($r = 2\%$, $\pi = 0.98$), 400 periods; right: $y_t = \rho y_{t-1} + \varepsilon_t$ with the same shocks and $\rho = 0.9$, 1, 1.02

TSA_ch13_rational_bubble

Interpreting the simulation

- The bubble grows, bursts, restarts: in this path it reaches 71 and bursts 2 times after a large size
  - the price inherits the explosive episodes; between them it follows the fundamental value
  - the first bubble of this path is negative and drives the price below zero: free disposal rules this out for a real asset, so a rational bubble must be nonnegative
- AR(1) roots (Chapter 3): $\rho = 0.9$ reverts to the mean, $\rho = 1$ wanders (unit root), $\rho = 1.02$ explodes
  - with $\rho = 1.02$ a shock doubles in about 35 periods and is multiplied by 7.24 after 100
- Idea of the tests: an explosive root ($\rho > 1$) in the log price during a run-up is the statistical trace of a rational bubble

Right-tailed unit-root tests

- The regression of Chapter 3 on a window of the log price $y_t = \ln P_t$: $\Delta y_t = a + \delta\, y_{t-\mathbf{1}} + \varepsilon_t$, with $\delta = \rho - 1$
  - Chapter 3 (Dickey and Fuller (1979)): $H_\mathbf{0}$: $\delta = 0$ against $H_\mathbf{1}$: $\delta < 0$ (stationary), the **left** tail
  - bubbles: $H_\mathbf{0}$: $\delta = 0$ (unit root) against $H_\mathbf{1}$: $\delta > 0$ (explosive), the **right** tail: reject for large t-statistics
- Problem: a bubble occupies only a part of the sample, and the crash after it looks like mean reversion
  - one ADF on the whole sample has almost no power: we must compute the test on many **windows**
- Notation
  - $\Delta y_t = y_t - y_{t-\mathbf{1}}$; a: the intercept; ε_t: the error; ρ: the AR(1) coefficient of y_t
  - $r_\mathbf{1}, r_\mathbf{2} \in [0,1]$: fractions of the sample; $\mathrm{ADF}_{r_\mathbf{1}}^{r_\mathbf{2}}$: the t-statistic of $\hat{\delta}$ on the window from $r_\mathbf{1} T$ to $r_\mathbf{2} T$; $r_\mathbf{0}$: the smallest window

SADF, GSADF and BSADF (1/2)

- sup: the largest value of the ADF statistic over all the windows considered
  - a single large value, in any window, is enough to signal an explosive episode
- **SADF** Phillips, Wu and Yu (2011)
  - start fixed at 0, the end r_2 moves forward
  - SADF $= \sup_{r_2 \in [r_0, 1]} \mathrm{ADF}_0^{r_2}$
  - good for one bubble; after a first crash it loses power
- **GSADF** Phillips, Shi and Yu (2015a)
  - both the start r_1 and the end r_2 move
  - GSADF $= \sup_{r_2 \in [r_0, 1]} \sup_{r_1 \in [0, r_2 - r_0]} \mathrm{ADF}_{r_1}^{r_2}$
  - detects several bubbles in one sample; smallest window $r_0 = 0.01 + 1.8/\sqrt{T}$

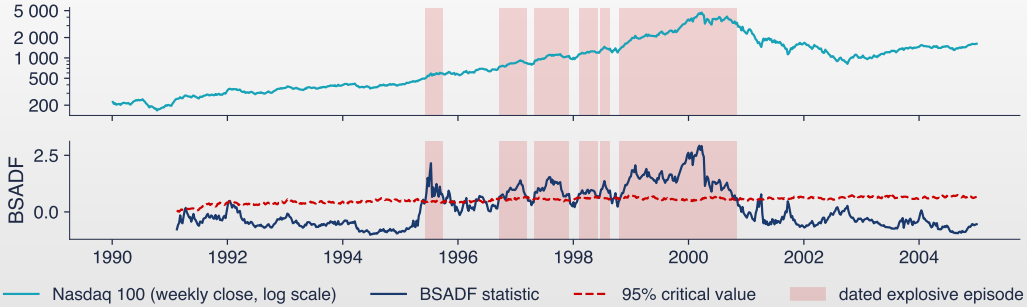
SADF, GSADF and BSADF (2/2)

- **BSADF** (date-stamping)
  - one statistic for each date r_2, using only data up to r_2
  - $\text{BSADF}(r_2) = \sup_{r_1} \text{ADF}_{r_1}^{r_2}$: the largest statistic over all starts r_1 of windows ending at r_2
  - an episode starts when BSADF crosses its critical value and lasts at least $\log T$ observations Phillips, Shi and Yu (2015b) (log: the natural logarithm)
- The distributions are not standard: critical values by Monte Carlo
  - simulated under a random walk with a weak drift, $y_t = T^{-1} + y_{t-1} + \varepsilon_t$
  - the drift T^{-1} is very small: a unit root, not an explosive root, under H_0
- Reading: an ADF, SADF or GSADF statistic above its 95% critical value rejects the unit root in favour of an explosive root

Worked example: one window of the Nasdaq 100

- Weekly log closes of the Nasdaq 100, $T = 782$ weeks; $r_0 = 0.074$, i.e. a smallest window of 58 weeks
  - end of the window: 24 March 2000, the week of the peak; among all start points the largest statistic comes from the start 17 January 1992
- OLS on this window ($n = 427$ weekly changes): $\widehat{\Delta y_t} = -0.0336 + 0.0060\, y_{t-1}$
  - $t = \hat{\delta}/\mathrm{SE}(\hat{\delta}) = 0.006030/0.002072 = 2.91$
  - so BSADF at the peak is 2.91, above its 95% critical value of 0.54
- Reading: $\hat{\delta} > 0$ means $\hat{\rho} = 1 + \hat{\delta} > 1$: the higher the price, the larger the next weekly rise, the mark of an explosive phase

Explosive episodes of the Nasdaq 100, 1990–2004



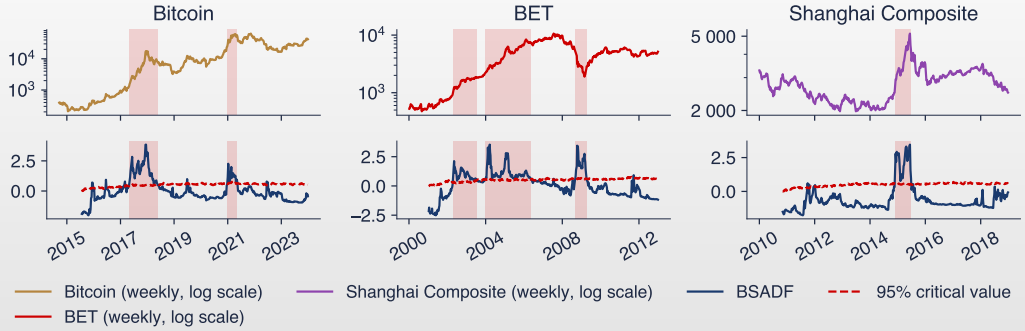
- Top: weekly close, log scale; bottom: BSADF and its 95% critical values (Monte Carlo, 1000 random walks); shaded: dated episodes (at least 7 weeks above the critical value)

TSA_ch13_explosive_roots

Interpreting the Nasdaq 100 tests

- Whole-sample ADF -1.27 (95% critical value -0.13): no evidence
  - SADF 2.72 (critical value 1.48) and GSADF 2.92 (critical value 2.28): explosive behaviour
  - the crash of 2000–2002 hides the bubble from a single test on the whole sample
- BSADF dates 6 episodes; the longest runs from 23 October 1998 to 3 November 2000 (107 weeks)
  - short explosive bursts already from June 1995: the 1990s boom came in waves, as Phillips, Wu and Yu (2011) found
- The end date is late: the statistic stays high for months after the peak; the test dates exuberance, it does not announce the crash

Bitcoin, BET and Shanghai: dated episodes



⊡ Weekly log closes; top: price with the dated episodes; bottom: BSADF and its 95% critical values; Bitcoin 2014–2023, BET 2000–2012, Shanghai Composite 2010–2018

TSA_ch13_explosive_roots

Interpreting the three markets

- GSADF: Bitcoin 3.85, BET 3.53, Shanghai 3.46; above their 95% critical values (2.18, 2.29, 2.15)
  - Bitcoin: April 2017 – May 2018 and January 2021 – May 2021; Shanghai: November 2014 – June 2015
- BET: the explosive phase is December 2003 – May 2006, not the run-up to July 2007, which was strong but not explosive; an earlier episode is dated April 2002 – July 2003
  - the BET flag of September 2008 – April 2009 is the collapse: an **accelerating fall** also gives $\hat{\delta} > 0$; always read the dates against the price chart
- The Bitcoin 2017 episode lasts until May 2018, again well after the peak

Limits of the explosive-root tests

- **Fundamentals**
  - an explosive price is a bubble only if the fundamentals are not explosive
  - better: test the price–dividend ratio (Phillips, Shi and Yu (2015a) for the S&P 500); for Bitcoin there are no dividends at all
- **Volatility**
  - changing volatility distorts the size of the test
  - a wild bootstrap gives robust critical values Phillips and Shi (2020)
- **Timing**
  - BSADF uses only past data, so it can run in real time
  - but it signals a bubble that is already under way and says nothing about when it ends
  - the LPPL model of the next section tries to answer exactly this: when?

Recap: Rational bubbles and explosive roots

- A rational bubble grows in expectation at the rate r: $E_t[B_{t+1}] = (1 + r)B_t$, an explosive process
- Right-tailed ADF: H_1: $\delta > 0$; SADF and GSADF take the supremum over windows; BSADF dates the episodes
- Nasdaq 100, Bitcoin, BET and Shanghai all show explosive episodes; the tests date the run-up, not the crash

Crashes as critical points



Didier Sornette (born 1957)

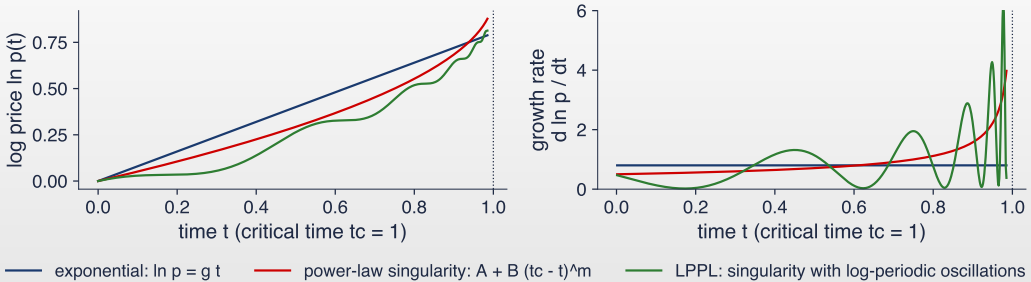
Photo: Didier Sornette (2012); CC BY-SA 3.0 DE; Wikimedia Commons

- Physicist, ETH Zürich; with Anders Johansen and Olivier Ledoit he proposed the **JLS model** Johansen, Ledoit and Sornette (2000)
- Idea: traders imitate their neighbours; imitation creates positive feedback
  - rising prices attract buyers, who push prices higher: growth accelerates
  - the market becomes more and more fragile, like a physical system near a **critical point** (water near boiling)
- The **critical time** t_c
  - the most probable moment of the end of the bubble
  - a crash is likely near t_c but not certain: the bubble may also end with a slow decline

Super-exponential growth

- **Exponential** growth
  - constant growth rate, $\ln P(t) = a + g\,t$
  - $d \ln P/dt = g$
  - a: the initial log price; g: the constant growth rate; $d \ln P/dt$: the growth rate of the price at time t
  - a straight line on a log scale
- **Super-exponential** growth
  - the growth rate increases with time
  - on a log scale the curve bends upward
  - power-law singularity: $\ln P(t) = A + B\,(t_c - t)^m$, with $B < 0$ and $0 < m < 1$
  - t_c: the critical time; $t_c - t$: the time left until it; A: the log price reached at t_c; B: the size of the growth; m: the exponent
  - growth rate $d \ln P/dt = -Bm\,(t_c - t)^{m-\mathbf{1}} \to \infty$ as $t \to t_c$, while $\ln P(t_c) = A$ stays finite
- Such growth cannot last: it must end at or before t_c, a **finite-time singularity**

Exponential and super-exponential growth



- Left: log price; right: its growth rate $d \ln P/dt$; exponential ($g = 0.8$), power-law singularity ($A = 1$, $B = -1$, $m = 0.5$, $t_c = 1$) and the same with log-periodic oscillations

TSA_ch13_lppl_model

Interpreting the growth paths

- All three start with similar growth rates: early in a bubble the paths are hard to tell apart
- The exponential rate stays at 0.8; the singular rate rises from 0.50 to 0.70 halfway and to 3.98 just before t_c
- The oscillations ride on the accelerating trend and become faster near t_c: the next slides explain why

From a crash hazard to the LPPL equation (1/2)

- JLS: the relative price change before the crash has three parts
  - $dP/P = \mu(t)\,dt + \sigma\,dW - \kappa\,dj$
  - $\mu(t)\,dt$: the expected growth over a short interval dt; $\mu(t)$: the drift, which may change in time
  - $\sigma\,dW$: the ordinary noise; W: a Brownian motion (a continuous random walk); σ: the volatility
  - $\kappa\,dj$: the crash; j jumps from 0 to 1 at the crash, which removes a fraction κ of the price
- **Hazard rate** $h(t)$
  - the probability per unit of time that the crash happens now, given that it has not happened yet

From a crash hazard to the LPPL equation (2/2)

- No arbitrage: $E[dP] = 0 \Rightarrow \mu(t) = \kappa\, h(t)$
  - $E[dj] = h(t)\, dt$, so the expected change is zero only if the drift pays for the expected crash loss
  - the higher the hazard, the faster the price must grow: risk and growth increase together
- Imitation near a critical point gives a hazard that rises towards t_c with oscillations
  - $h(t) \approx \alpha(t_c - t)^{m-\mathbf{1}}[1 + \beta\cos(\omega \ln(t_c - t) - \phi')]$
  - $\alpha > 0$: the level of the hazard; $(t_c - t)^{m-\mathbf{1}}$ grows without bound as $t \to t_c$, since $m < 1$
  - β: the relative size of the oscillations; ω: their log-frequency; ϕ': their phase
- Integrating $\mu(t) = \kappa h(t)$ over time gives the log price on the next slide

The LPPL equation

- **Log-periodic power law** (LPPL), for $t < t_c$
  - $\ln P(t) = A + B\,(t_c - t)^m + C\,(t_c - t)^m \cos\big(\omega \ln(t_c - t) - \phi\big)$

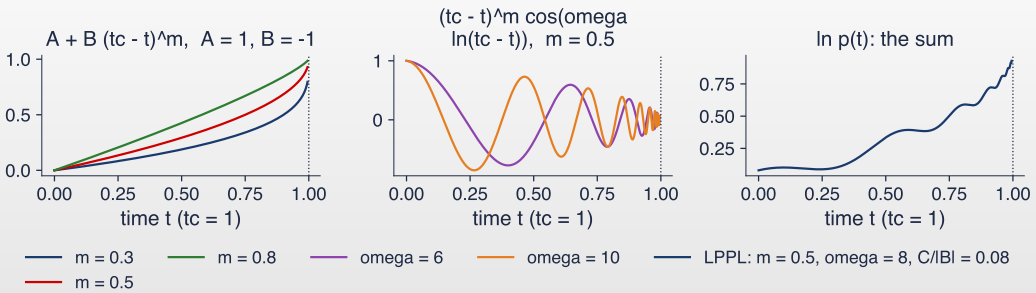
Parameter	Meaning	Usual range
t_c	critical time: most probable end of the bubble	after the last observation
A	log price at t_c (if the bubble reached t_c)	$A > 0$
B	size of the power-law growth	$B < 0$ (rising price)
m	exponent: how fast the growth accelerates	$0 < m < 1$
C	relative size of the oscillations	$\lvert C\rvert < \lvert B\rvert$
ω	angular log-frequency of the oscillations	$2 \leq \omega \leq 25$
ϕ	phase of the oscillations	$[0, 2\pi)$

- Seven parameters: three of them (t_c, m, ω) enter nonlinearly; ϕ disappears in the estimation (Section 5)

Log-periodic oscillations

- $\cos(\omega \ln(t_c - t) - \phi)$ is periodic in $\ln(t_c - t)$, not in t: one full cycle each time $\ln(t_c - t)$ changes by $2\pi/\omega$
  - successive peaks t_n satisfy $\dfrac{t_c - t_n}{t_c - t_{n+1}} = \lambda = e^{2\pi/\omega}$: the time left to t_c shrinks by the same factor at each cycle
- **Discrete scale invariance**
  - the pattern looks the same when time to t_c is rescaled by λ (as in a hierarchy of traders, groups, institutions)
- Worked example: with $\omega = 8$, $\lambda = e^{2\pi/8} = 2.19$
  - if a correction happens 200 days before t_c, the next one comes about $200/2.19 \approx 91$ days before t_c, then about 42, then 19
  - $\omega = 6.28$ gives $\lambda = 2.72$; $\omega = 10$ gives $\lambda = 1.87$; the empirical literature reports λ near 2 Sornette (2017)

The pieces of the LPPL equation



- Left: $A + B(t_c - t)^m$ for $m = 0.3, 0.5, 0.8$; middle: the oscillation $(t_c - t)^m \cos(\omega \ln(t_c - t))$ for $\omega = 6$ and 10; right: the full LPPL path ($m = 0.5$, $\omega = 8$, $C/|B| = 0.08$), which fails the damping condition: $m|B|/(\omega|C|) = 0.78 < 1$; $t_c = 1$

TSA_ch13_lppl_model

Interpreting the components

- A smaller m concentrates the growth close to t_c: the curve is flat for a long time, then shoots up
  - m close to 1 is almost a straight line: hardly a bubble; m close to 0 is a jump at t_c
- The oscillation term shrinks with $(t_c - t)^m$ and speeds up: the cycles get shorter by the factor λ
  - a larger ω gives more cycles before t_c
- Real run-ups (Section 5) show this mixture: an accelerating trend with corrections that come more and more often

Worked example: an LPPL value by hand

- Parameters: $A = 8$, $B = -1.2$, $C = 0.08$, $m = 0.5$, $\omega = 8$, $\phi = 0$, $t_c = 2$ (years); compute $\ln P(t)$
  - $t = 1$: $t_c - t = 1$, $(t_c - t)^m = 1$, $\ln 1 = 0$, $\cos 0 = 1$: $\ln P = 8 - 1.2 + 0.08 = 6.880$, $P = 973$
  - $t = 1.75$: $(0.25)^{\mathbf{0.5}} = 0.500$, $\ln 0.25 = -1.386$, $\cos(8 \cdot (-1.386)) = 0.095$:
    $\ln P = 8 - 1.2 \cdot 0.500 + 0.08 \cdot 0.500 \cdot 0.095 = 7.404$, $P = 1642$
  - $t = 1.99$: $(0.01)^{\mathbf{0.5}} = 0.100$, $\cos(8 \cdot (-4.605)) = 0.654$: $\ln P = 7.885$, $P = 2658$
- The price rises from 973 to 2658 as t goes from 1 to 1.99 and would reach $e^A = 2981$ at t_c; most of the rise comes in the last months

Parameter constraints: the filter

- A fit is **qualified** only if it describes a credible bubble (Shu and Zhu (2020), following Sornette et al. (2015)):
  - $B < 0$: rising price; $0.01 \leq m \leq 0.99$: accelerating growth with a finite price at t_c
  - $2 \leq \omega \leq 25$: neither too slow nor too fast oscillations; at least 2.5 oscillations in the window, $\frac{\omega}{\pi} \ln \frac{t_c - t_1}{t_c - t_2} \geq 2.5$
  - $t_2 \leq t_c \leq t_2 + 0.2\,(t_2 - t_1)$: the end is close, relative to the length of the window $[t_1, t_2]$
  - damping $\dfrac{m|B|}{\omega|C|} \geq 1$: the hazard rate stays positive
- Conditions on the residuals
  - every fitted price within 15% of the observed price; the Lomb periodogram Lomb (1976) confirms the log-periodic cycle (level 10%)
  - Lomb periodogram: a periodogram for unevenly spaced points, here of the residuals as a function of $\ln(t_c - t)$
  - the residuals are stationary: the Dickey–Fuller and Phillips–Perron tests of Chapter 3 reject a unit root (level 10%)

Recap: The LPPL model

- Imitation creates a hazard rate that grows towards t_c; no arbitrage turns it into super-exponential growth
- $\ln P(t) = A + B(t_c - t)^m + C(t_c - t)^m \cos(\omega \ln(t_c - t) - \phi)$, seven parameters
- Log-periodic: cycles shrink by $\lambda = e^{2\pi/\omega}$; a fit counts only if it passes the filter

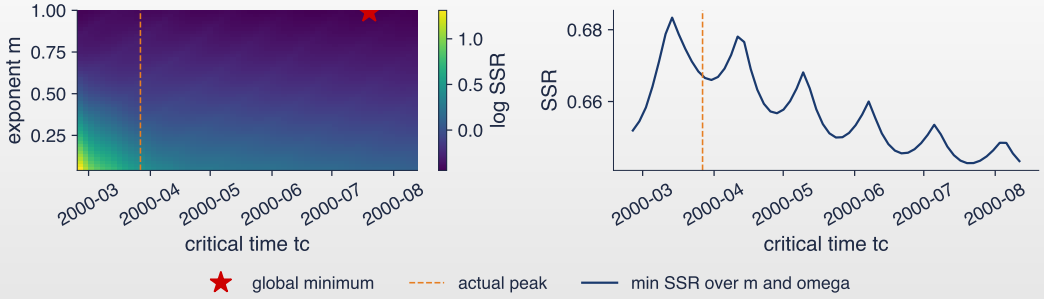
The two-step method of Filimonov and Sornette (1/2)

- Least squares over seven parameters has many local minima; Filimonov and Sornette (2013) reduce the search to three parameters
  - $\min \mathrm{SSR} = \sum_{t=t_1}^{t_2} [\ln P_t - \mathrm{LPPL}(t)]^2$; SSR: the sum of squared residuals over the window $[t_1, t_2]$
  - $\mathrm{LPPL}(t)$: the right-hand side of the LPPL equation
- Expand the cosine, with $\tau = t_c - t$
  - $C \cos(\omega \ln \tau - \phi) = C_1 \cos(\omega \ln \tau) + C_2 \sin(\omega \ln \tau)$, where $C_1 = C \cos \phi$, $C_2 = C \sin \phi$
- The LPPL equation becomes linear in A, B, C_1, C_2
  - $\ln P(t) = A + B\, f_t + C_1\, g_t + C_2\, h_t$
  - $f_t = \tau^m$, $g_t = \tau^m \cos(\omega \ln \tau)$, $h_t = \tau^m \sin(\omega \ln \tau)$: known numbers once t_c, m, ω are fixed

The two-step method of Filimonov and Sornette (2/2)

- **Step 1** (linear)
  - for given (t_c, m, ω), regress $\ln P_t$ on $(1, f_t, g_t, h_t)$ by OLS
  - $(\hat{A}, \hat{B}, \hat{C}_1, \hat{C}_2) = (X'X)^{-1}X'y$; X: the matrix with columns $1, f_t, g_t, h_t$; y: the vector of $\ln P_t$
  - this gives the concentrated cost $\mathrm{SSR}(t_c, m, \omega)$, a function of three parameters only
- **Step 2** (nonlinear)
  - minimise $\mathrm{SSR}(t_c, m, \omega)$ over the search space
  - $t_c \in [t_2, t_2 + (t_2 - t_1)/3]$, $m \in [0, 1]$, $\omega \in [1, 50]$
  - in the code: a grid of $8 \times 8 \times 16$ points, then the Nelder–Mead simplex from the best grid point
- At the end: $C = \sqrt{C_1^2 + C_2^2}$ and $\phi = \mathrm{atan2}(C_2, C_1)$
  - $\mathrm{atan2}(C_2, C_1)$: the angle of the point (C_1, C_2), in $(-\pi, \pi]$

The cost landscape



- Nasdaq 100, window from the low of October 1998 to 30 days before the peak
  - left: $\log \mathrm{SSR}$ over (t_c, m), minimised over $\omega \in [2, 25]$ (135 360 grid points, with the linear step solved at each point)
  - right: the minimum over m and ω as a function of t_c

TSA_ch13_estimation

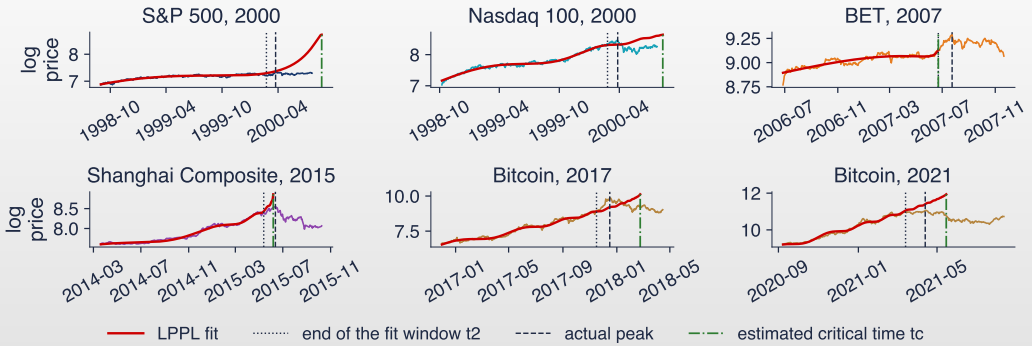
Interpreting the cost landscape

- The profile over t_c has 5 local minima: an optimiser started at a different t_c can end in a different valley
  - the waves come from the oscillation term: shifting t_c shifts the phase of the cycles
- The landscape is flat: the worst and the best t_c differ by only 6.3% in SSR
  - the global minimum is at m near 1 and $t_c = 19$ July 2000, far after the actual peak: the data barely identify t_c
- Lesson: one LPPL fit is never enough; we need many windows and the filter (Section 6)

Worked example: the Shanghai fit at the end of the window

- Shanghai Composite, window 20 March 2014 – 13 May 2015 (281 days); estimates: $t_c = 2015.428$ (6 June 2015), $m = 0.392$, $\omega = 5.50$
  - step 1 gives $\hat{A} = 8.893$, $\hat{B} = -1.270$, $\hat{C}_1 = 0.020$, $\hat{C}_2 = 0.074$, so $\hat{C} = 0.076$
- At $t_2 = 2015.361$ (years): $\tau = t_c - t_2 = 0.067$, $\tau^m = 0.347$, $\ln \tau = -2.702$
  - $\cos(\omega \ln \tau) = -0.656$, $\sin(\omega \ln \tau) = -0.754$
  - $\widehat{\ln P} = 8.893 + (-0.440) + (-0.024) = 8.429$; observed $\ln P = 8.384$
- In levels: fitted 4579 against observed 4376, an error of 4.7%; the three terms are the level A, the power-law growth and the oscillation

LPPL fits 30 days before the peak



- Window: from the low before the run-up to 30 calendar days before the peak; the fitted path is extended to the estimated t_c; data after the window are shown only for comparison

TSA_ch13_estimation

Interpreting the six fits

Episode	t_2	$\hat{t}_c$	Days to peak	$\hat{m}$	$\hat{\omega}$	Param.	All
S&P 500, 2000	23 February 2000	21 August 2000	+150	1.00	1.0	no	no
Nasdaq 100, 2000	25 February 2000	11 August 2000	+137	1.00	5.9	no	no
BET, 2007	22 June 2007	22 June 2007	−32	0.69	1.0	no	no
Shanghai Composite, 2015	13 May 2015	6 June 2015	−6	0.39	5.5	yes	yes
Bitcoin, 2017	16 November 2017	22 February 2018	+68	0.71	14.1	no	no
Bitcoin, 2021	14 March 2021	15 May 2021	+32	0.82	17.4	no	no

- Days to peak: $\hat{t}_c$ minus the actual peak; Param.: the fit passes the parameter conditions; All: it also passes the residual conditions
- Only Shanghai passes the whole filter, with $\hat{t}_c$ within a week of the peak; S&P 500 and Nasdaq 100 hit the boundary $m = 1$: no super-exponential growth in these windows
- Bitcoin fails the 15% error condition (its daily swings are too large for one smooth path) and the t_c condition ($\hat{t}_c$ too far after t_2)

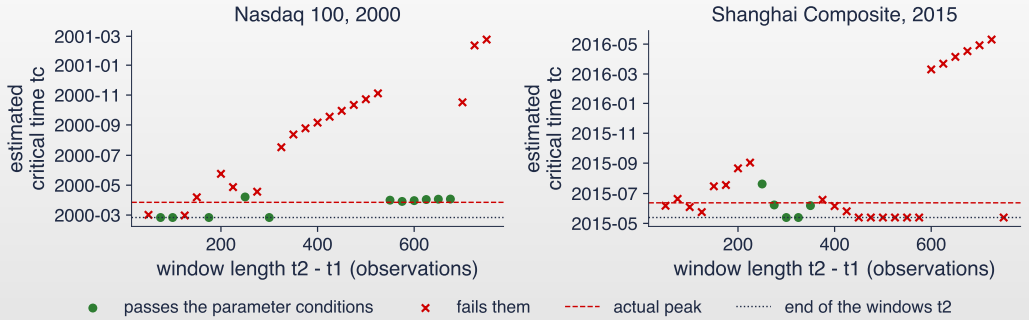
Recap: Estimation

- Four linear parameters by OLS, three nonlinear ones (t_c, m, ω) by search: the method of Filimonov and Sornette (2013)
- The cost surface is rugged and flat in t_c: the critical time is weakly identified
- One window, one fit: only one of six episodes gives a fully qualified fit 30 days before the peak

The LPPLS confidence indicator

- So far t_1 was the low before the run-up, a date known only **after** the bubble; in real time the start is unknown
  - remedy: for each end date t_2 (today), fit LPPL on 29 windows $[t_2 - L, t_2]$, $L = 750, 725, \ldots, 50$ trading days Shu and Zhu (2020)
- Indicator = share of the windows ending at t_2 whose fit is qualified Sornette et al. (2015); 0: no window sees a bubble, 1: all do
  - LPPLS (LPPL singularity) is the name used for this indicator; the equation is the same
- Two versions in this chapter
  - **parameter conditions**: B, m, ω, t_c, oscillations, damping
  - **all conditions**: also the 15% error, the Lomb test and stationary residuals (Shu and Zhu (2020))
- Computed only with data up to t_2: a real-time indicator; an alarm when it exceeds a threshold fixed in advance (here 0.2)

Estimated critical time across estimation windows



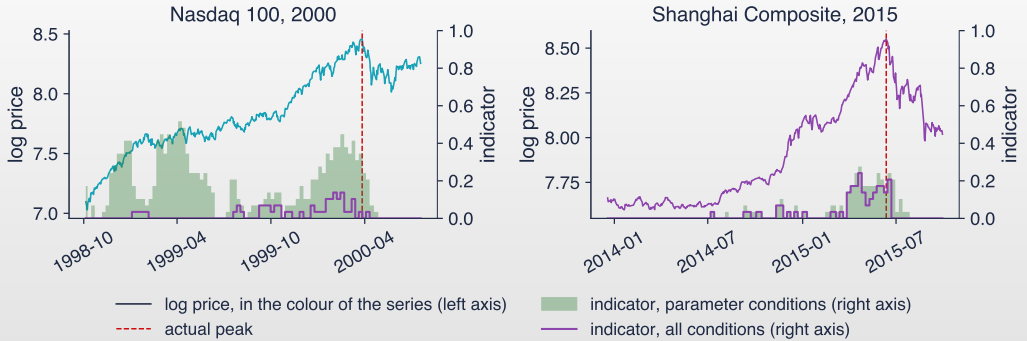
- Estimated t_c for each of the 29 windows ending 30 days before the peak; green dots: the fit passes the parameter conditions; red crosses: it fails them

TSA_ch13_confidence

Interpreting the windows

- Nasdaq 100 (end 25 February 2000): 11 of 29 windows pass the parameter conditions, 1 the whole filter
  - qualified $\hat{t}_c$: median 30 March 2000, 10%–90% range 25 February 2000 – 2 April 2000; actual peak 27 March 2000
- Shanghai (end 13 May 2015): 5 of 29 pass the parameter conditions; median $\hat{t}_c$ 6 June 2015, actual peak 12 June 2015
  - the condition on m fails most often: only 41% of the Shanghai windows have $0.01 \leq m \leq 0.99$
- Short windows put t_c right after t_2; long windows spread it out: the window length is a hidden modelling choice

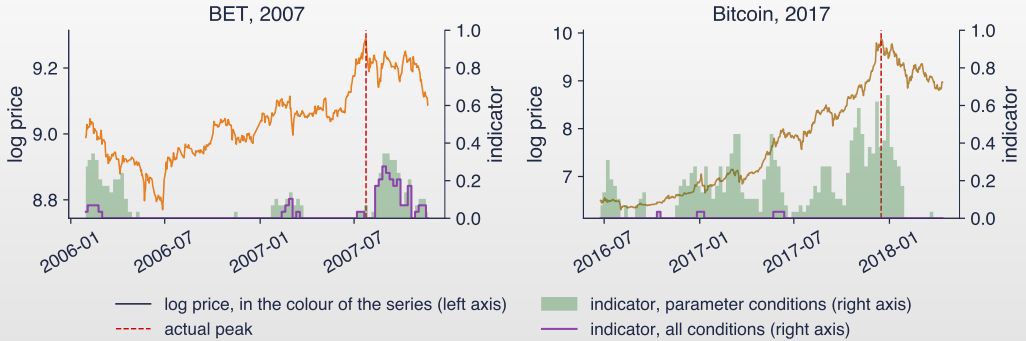
The indicator around four peaks (1/2)



- Log price (left axis) and the indicator (right axis) every 5 trading days, from 18 months before to 4 months after the peak
  - green area: parameter conditions; purple line: all conditions

TSA_ch13_confidence

The indicator around four peaks (2/2)



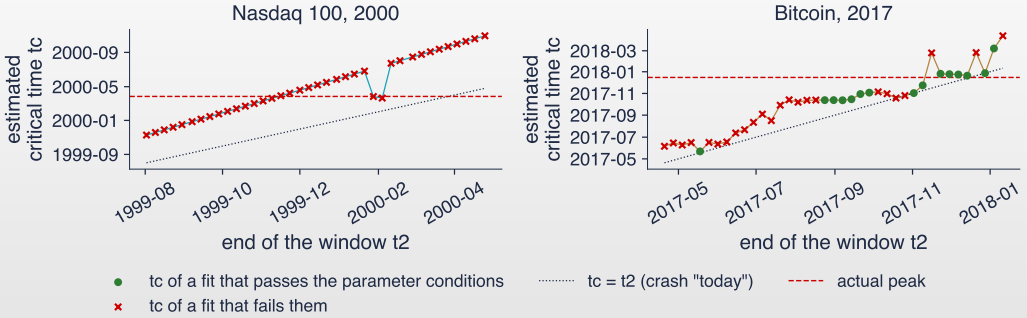
- The same layout and colours as in (1/2), for the BET (2007) and Bitcoin (2017); Bitcoin: one value every 7 days

TSA_ch13_confidence

Interpreting the indicator

- Nasdaq 100: the parameter version reaches 0.52 on 7 April 1999, a year before the peak, and rises again in the last months; the full version stays below 0.14
  - Shanghai: both versions rise in the two months before the peak (maximum on 22 April 2015: 0.24 with the full filter, 0.28 with the parameter conditions only)
- BET: the strongest full-filter signal, 0.28, comes on 28 August 2007, after the peak: the fits then describe the last part of the run-up, not the future
  - Bitcoin 2017: the parameter version peaks at 0.66 on 29 December 2017; the full version is almost always 0 (the 15% error condition)
- The indicator does light up in bubbles, but early, late and with gaps; the choice of filter changes the picture

The critical time moves with the data



- Window start fixed at the low before the run-up; the end t_2 moves from 8 months before to 1 month after the peak (weekly); each point: the $\hat{t}_c$ of the fit ending at t_2

TSA_ch13_confidence

Interpreting the moving critical time

- Nasdaq 100: over 38 end dates $\hat{t}_c$ ranges from 9 November 1999 to 30 October 2000 (356 days); none of these fits passes the parameter conditions
  - Bitcoin 2017: range 21 May 2017 – 12 April 2018; 38% of the fits pass; median error 65 days
- Bitcoin: 4 of the 39 fits put $\hat{t}_c$ exactly at t_2 and 11 within 20 days of it: the fit then says "the end is now"; Nasdaq 100: $\hat{t}_c$ mostly sits at the upper bound of the search space
  - a forecast that keeps moving with the data is a description of the past, not a prediction
- Picking, after the crash, the one window whose $\hat{t}_c$ was right is **look-ahead bias**

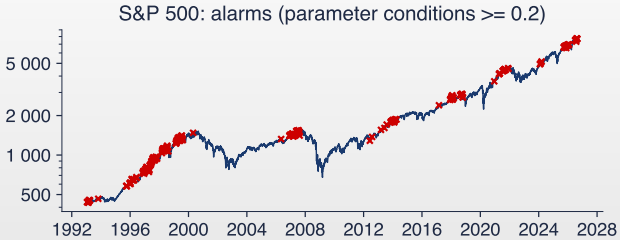
Recap: Confidence from many windows

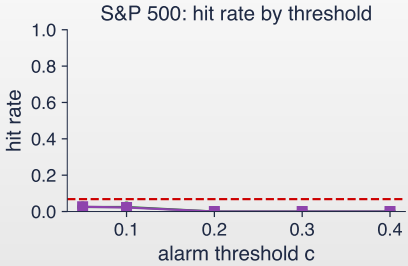
- The start of a bubble is unknown: fit 29 windows for each end date
- Indicator = share of qualified windows; it depends strongly on the filter
- $\hat{t}_c$ moves with the end of the data: report its spread over windows, never one value

Common errors in evaluating crash predictions

- **Look-ahead bias**
  - using information from after t_2
  - choosing t_1 at the low of the run-up, the episodes after the crash, or the filter after seeing the results
- **Selection**
  - studying only the bubbles that burst
  - the run-ups that did not end in a crash are forgotten
- **False alarms and base rates**
  - an alarm is useful only if a crash is **more likely after an alarm** than on an ordinary day
  - the comparison: $P(\text{crash} \mid \text{alarm})$ against $P(\text{crash})$, the unconditional frequency
- A fair test fixes everything in advance (windows, filter, threshold, the definition of a crash) and runs over the whole sample

Alarms over the whole sample (1/2)



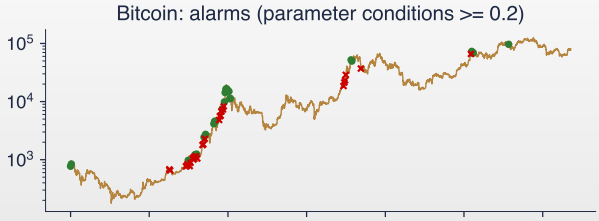


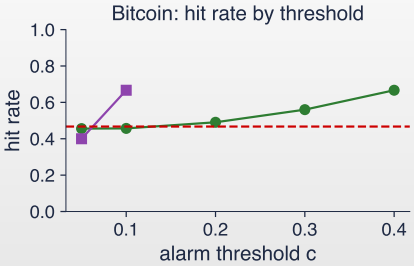
S&P 500 — alarm followed by a fall of 20% or more — false alarm — hit rate, parameter conditions — hit rate, all conditions — unconditional frequency

- S&P 500: indicator every 5 trading days, January 1993 – September 2026
  - a crash: a fall of at least 20% within 182 days after the alarm; right: hit rate by threshold

TSA_ch13_evaluation

Alarms over the whole sample (2/2)





Bitcoin
alarm followed by a fall of 20% or more
false alarm
hit rate, parameter conditions
hit rate, all conditions
unconditional frequency

- Bitcoin: indicator every 7 days, January 2014 – September 2026
  - hit rate: the share of the alarm dates followed by a crash; dashed line: the unconditional frequency

TSA_ch13_evaluation

Interpreting the evaluation

- S&P 500: a fall of 20% within six months follows 7% of all dates, but 0% of the alarm dates (threshold 0.2, parameter conditions)
  - 20 alarm clusters, none followed by such a fall within six months; the alarms come in calm, steady uptrends
  - of the 4 drawdowns of 20% or more since January 1993, 2 had an alarm in the six months before their peak (2007 and 2022), but the falls came later
- Bitcoin: base rate 47%, after an alarm 49%: almost no gain; 12 of its 14 drawdowns of 20% had an alarm before the peak
  - falls of 20% are so frequent in Bitcoin that almost any alarm looks "right"; higher thresholds help a little, on very few dates (56% at 0.3, on 4% of the dates)
- The indicator says "this looks like a bubble regime", not "the crash comes within six months"

The literature on crash prediction

- Successes reported by the authors of the method
  - Jiang et al. (2010): the Chinese bubbles of 2007 and 2009 diagnosed in advance; Sornette et al. (2015): the Shanghai 2015 peak announced before it happened
  - Gerlach, Demos and Sornette (2019): Bitcoin's bubbles of 2012–2018 dissected with the confidence indicator
- Critical evaluations
  - Feigenbaum (2001): the log-periodic term adds little once the estimation error is taken into account
  - Brée and Joseph (2013): on many crashes, the LPPL conditions are rarely all met before the crash; parameters are unstable
- Our own numbers point the same way: the indicator reacts to strong run-ups, but its timing is loose and its alarms are often false

Strengths and limits of LPPL

- **Can**
  - describe a run-up with accelerating growth and shrinking corrections in a few interpretable parameters
  - measure "how bubble-like" a market is today, across many windows (the indicator)
- **Cannot**
  - give the date of a crash: t_c marks the end of the regime
  - guarantee a crash: the regime may end with a slow decline
  - tell whether the price is above the fundamental value: that needs fundamentals
- Use it with the explosive-root tests, with fundamentals (valuation ratios) and with risk measures such as VaR 1% (Chapter 5), not alone

Recap: Evaluation of crash predictions

- Fix windows, filter, threshold and the definition of a crash before looking at the outcomes
- Compare the hit rate after alarms with the base rate; count the false alarms
- LPPL describes bubbles well; as a crash-timing tool its record is weak

Possible contribution of AI

- **Code**
  - a first version of the BSADF recursion, of the Filimonov–Sornette two-step fit, of the indicator over many windows
- **Explanation**
  - a second explanation of the hazard-rate argument or of log-periodicity, with your own numbers
- **Exploration**
  - the indicator on many markets and periods, with different filters, under a protocol fixed in advance
- Example prompt
  - `Write Python code that fits the LPPL model to the daily log price of the Shanghai Composite between two dates with the Filimonov-Sornette method (OLS for A, B, C1, C2; grid search plus Nelder-Mead for tc, m, omega), checks the filter conditions of Shu and Zhu (2020) and plots the fit.`

Checks you must run

- Simulate an LPPL path with known parameters plus noise and check that the code recovers t_c, m and ω
- The sign conventions: $B < 0$ for a rising bubble; the test of explosive roots uses the right tail, not the left tail of Chapter 3
- Critical values: SADF and GSADF need their own simulated critical values, not the Dickey–Fuller table
- No look-ahead: every quantity at date t_2 must use only data up to t_2
- A crash "predicted" by an AI answer after the fact is not a prediction; ask for the full list of alarms and false alarms
- Every cited reference: it must exist; check the DOI

Key takeaways

- A bubble is a price above its fundamental value; a rational bubble must grow explosively
- Right-tailed ADF tests (SADF, GSADF, BSADF) detect and date explosive episodes with data up to each date
- LPPL: super-exponential growth with log-periodic oscillations ending at a critical time t_c
- Estimate in two steps (OLS + search over t_c, m, ω), filter the fits, use many windows
- Judge predictions by hit rates against base rates, with all the false alarms: the timing of crashes remains hard

Key formulas

Quantity	Formula
Rational bubble	$P_t = F_t + B_t, \quad E_t[B_{t+1}] = (1+r)B_t$
Right-tailed ADF	$\Delta y_t = a + \delta y_{t-1} + \varepsilon_t, \quad H_1\colon \delta > 0$
GSADF, BSADF	$\sup_{r_2} \sup_{r_1} \mathrm{ADF}_{r_1}^{r_2}, \quad \mathrm{BSADF}(r_2) = \sup_{r_1} \mathrm{ADF}_{r_1}^{r_2}, \quad r_0 = 0.01 + 1.8/\sqrt{T}$
LPPL	$\ln P(t) = A + B(t_c - t)^m + C(t_c - t)^m \cos(\omega \ln(t_c - t) - \phi)$
Linear form	$\ln P(t) = A + Bf_t + C_1 g_t + C_2 h_t, \quad C = \sqrt{C_1^2 + C_2^2}$
Scaling ratio	$\lambda = e^{2\pi/\omega}$
Damping	$m\lvert B\rvert/(\omega\lvert C\rvert) \geq 1$
Indicator	share of qualified windows ending at t_2

Self-assessment (1/2)

- **Question**: in a Blanchard–Watson bubble with $r = 1\%$ and $\pi = 0.95$, by how much does the bubble grow in a period in which it survives?
  - **Answer**: by $1.01/0.95 - 1 \approx 6.3\%$; expected lifetime $1/(1 - 0.95) = 20$ periods
- **Question**: a right-tailed ADF on a window gives $\hat{\delta} = 0.004$ with SE 0.002. Is the window explosive at 5% if the critical value is 1.5?
  - **Answer**: yes: $t = 0.004/0.002 = 2 > 1.5$; we reject the unit root in favour of an explosive root
- **Question**: why do we use the right tail and not the left tail of Chapter 3?
  - **Answer**: the alternative is $\rho > 1$ (explosive), not $\rho < 1$ (stationary); evidence for it is a large positive t

Self-assessment (2/2)

- **Question**: an LPPL fit gives $\omega = 6.28$. By what factor does the time left to t_c shrink between two corrections?
  - **Answer**: $\lambda = e^{2\pi/6.28} \approx e \approx 2.72$
- **Question**: a fit has $m = 1.0$ at the edge of the search space. Is it a qualified bubble fit?
  - **Answer**: no: the filter needs $m \leq 0.99$; $m = 1$ is a linear trend in t, not an accelerating one
- **Question**: an indicator gives alarms before 8 of 10 crashes. Is it a good predictor?
  - **Answer**: not yet: we also need the false alarms and the base rate; if alarms are on almost every day, catching 8 crashes is easy

Project idea

- **Is the BET in a bubble today?** A real-time protocol, fixed before looking at the result
  - BSADF on the weekly BET and BET-TR; the indicator of this chapter on the daily BET
  - the evaluation of Section 7 on the BET since 2000: alarms, false alarms, base rate
- Extensions
  - the price–earnings ratio instead of the price; Bitcoin and Ethereum together
  - compare with the regime-switching models of Chapter 10
- Further reading on the Bucharest Stock Exchange: Pele and Mazurencu-Marinescu (2012)
- Next: Chapter 14, multivariate GARCH models

References (1/3)

- Blanchard, O. J., & Watson, M. W. (1982). Bubbles, rational expectations and financial markets. NBER Working Paper 945. doi:10.3386/w0945
- Brée, D. S., & Joseph, N. L. (2013). Testing for financial crashes using the log periodic power law model. *International Review of Financial Analysis*, 30, 287–297. doi:10.1016/j.irfa.2013.05.005
- Dickey, D. A., & Fuller, W. A. (1979). Distribution of the estimators for autoregressive time series with a unit root. *Journal of the American Statistical Association*, 74(366), 427–431. doi:10.1080/01621459.1979.10482531
- Fama, E. F. (2014). Two pillars of asset pricing. *American Economic Review*, 104(6), 1467–1485. doi:10.1257/aer.104.6.1467
- Feigenbaum, J. A. (2001). A statistical analysis of log-periodic precursors to financial crashes. *Quantitative Finance*, 1(3), 346–360. doi:10.1088/1469-7688/1/3/306
- Filimonov, V., & Sornette, D. (2013). A stable and robust calibration scheme of the log-periodic power law model. *Physica A*, 392(17), 3698–3707. doi:10.1016/j.physa.2013.04.012
- Garber, P. M. (1990). Famous first bubbles. *Journal of Economic Perspectives*, 4(2), 35–54. doi:10.1257/jep.4.2.35
- Gerlach, J.-C., Demos, G., & Sornette, D. (2019). Dissection of Bitcoin's multiscale bubble history from January 2012 to February 2018. *Royal Society Open Science*, 6(7), 180643. doi:10.1098/rsos.180643

References (2/3)

- Greenwood, R., Shleifer, A., & You, Y. (2019). Bubbles for Fama. *Journal of Financial Economics*, 131(1), 20–43. doi:10.1016/j.jfineco.2018.09.002
- Huang, C., & Petukhina, A. (2022). *Applied Time Series Analysis and Forecasting with Python*. Springer. doi:10.1007/978-3-031-13584-2
- 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 & Organization*, 74(3), 149–162. doi:10.1016/j.jebo.2010.02.007
- Johansen, A., Ledoit, O., & Sornette, D. (2000). Crashes as critical points. *International Journal of Theoretical and Applied Finance*, 3(2), 219–255. doi:10.1142/S0219024900000115
- Kindleberger, C. P., & Aliber, R. Z. (2005). *Manias, Panics and Crashes: A History of Financial Crises* (5th ed.). Palgrave Macmillan. doi:10.1057/9780230628045
- Lomb, N. R. (1976). Least-squares frequency analysis of unequally spaced data. *Astrophysics and Space Science*, 39(2), 447–462. doi:10.1007/BF00648343
- Pele, D. T., & Mazurencu-Marinescu, M. (2012). Modelling stock market crashes: The case of Bucharest Stock Exchange. *Procedia – Social and Behavioral Sciences*, 58, 533–542. doi:10.1016/j.sbspro.2012.09.1030
- Phillips, P. C. B., & Shi, S. (2020). Real time monitoring of asset markets: Bubbles and crises. In *Handbook of Statistics: Financial, Macro and Micro Econometrics Using R*, 61–80. Elsevier. doi:10.1016/bs.host.2018.12.002

References (3/3)

- Phillips, P. C. B., Shi, S., & Yu, J. (2015a). Testing for multiple bubbles: Historical episodes of exuberance and collapse in the S&P 500. *International Economic Review*, 56(4), 1043–1078. doi:10.1111/iere.12132
- Phillips, P. C. B., Shi, S., & Yu, J. (2015b). Testing for multiple bubbles: Limit theory of real-time detectors. *International Economic Review*, 56(4), 1079–1134. doi:10.1111/iere.12131
- Phillips, P. C. B., Wu, Y., & Yu, J. (2011). Explosive behavior in the 1990s Nasdaq: When did exuberance escalate asset values? *International Economic Review*, 52(1), 201–226. doi:10.1111/j.1468-2354.2010.00625.x
- Shiller, R. J. (2015). *Irrational Exuberance* (3rd ed.). Princeton University Press. doi:10.1515/9781400865536
- Shu, M., & Zhu, W. (2020). Detection of Chinese stock market bubbles with LPPLS confidence indicator. *Physica A*, 557, 124892. doi:10.1016/j.physa.2020.124892
- Sornette, D. (2017). *Why Stock Markets Crash: Critical Events in Complex Financial Systems*. Princeton University Press. doi:10.23943/princeton/9780691175959.001.0001
- Sornette, D., Demos, G., Zhang, Q., Cauwels, P., & Zhang, Q. (2015). Real-time prediction and post-mortem analysis of the Shanghai 2015 stock market bubble and crash. *Journal of Investment Strategies*, 4(4), 77–95. doi:10.21314/JOIS.2015.063
- Sornette, D., Johansen, A., & Bouchaud, J.-P. (1996). Stock market crashes, precursors and replicas. *Journal de Physique I*, 6(1), 167–175. doi:10.1051/jp1:1996135