risks

MDPI

Article

Evaluation and Prediction of Stock Market Crash Risk in Mexico Using Log-Periodic Power-Law Modeling

Suryansh Sunil [1], Amit Kumar Goyal [2,*], Rajesh Mahadeva [3] and Varun Sarda [4]

[1] Department of Aeronautical and Automobile Engineering, Manipal Institute of Technology, Manipal Academy of Higher Education, Manipal 576104, India; suryansh.sunil@learner.manipal.edu
[2] Department of Electronics and Communication Engineering, Manipal Institute of Technology, Manipal Academy of Higher Education, Manipal 576104, India
[3] Department of Computer Science and Engineering, Manipal Institute of Technology, Manipal Academy of Higher Education, Manipal 576104, India; rajesh.mahadeva@manipal.edu
[4] Department of Humanities and Management, Manipal Institute of Technology, Manipal Academy of Higher Education, Manipal 576104, India; varun.sarda@manipal.edu
* Correspondence: amit.goyal@manipal.edu

Abstract

This study applies the Log-Periodic Power-Law (LPPL) framework to three major equity markets—Mexico (IPC), Brazil (IBOVESPA), and the United States (NYSE Composite)—using daily closes from 8 November 1991–30 January 2025 for IPC and NYSE, and 3 May 1993–30 January 2025 for IBOVESPA. Multi-window calibrations (Lϵ 180, 240, 300, 360, 420) are estimated in raw and log space to evaluate bubble signatures and the stability of the critical time t_c. Across all indices, log-space fits consistently outperform raw fits in terms of RMSE and R^2, and longer windows reduce parameter variability, yielding coherent clusters of t_c. Under full-sample conditions, the LPPL structure points to March–April 2025 for NYSE, mid-October 2025 for IBOVESPA, and October–December 2025 for IPC, while shorter windows pull t_c forward. A rolling early-warning ensemble translates these estimates into lead-based risk bands, with numerical reporting used when median leads fall just outside the 60-trading-day decision horizon. The early-2025 weakening in the U.S. market is consistent with the NYSE cluster, whereas Brazil and Mexico remain within their projected windows as of September 2025. The analysis highlights the strengths of LPPL—behavioral interpretability and hazard-based framing—while noting limitations such as window sensitivity and parameter sloppiness, reinforcing the need for conservative communication and the use of longer-window weighting in practical applications.

Keywords: financial stability; Mexican stock market; financial time series; bubbles and crashes; stock market crashes; sustainable financial behavior; predictive modeling; log-periodic power law

Check for updates

Academic Editor: Krzysztof Jajuga

Received: 13 October 2025
Revised: 22 November 2025
Accepted: 1 December 2025
Published: 1 January 2026

Copyright: © 2026 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license.

1. Introduction

Financial markets periodically undergo abrupt corrections or collapses following extended phases of speculative expansion. Understanding and anticipating such regime shifts remains a central challenge in financial economics. Equity markets exhibit nonlinear dynamics, heavy-tailed behavior, and high volatility, making forecasting particularly difficult for individual investors who often lack advanced analytical tools (Mu et al. 2023). When asset prices drift away from fundamental value—driven by herding, extrapolative expectations, and speculative feedback—bubbles emerge and eventually unwind (Brée and

Joseph 2013). These reversals are typically sudden, characterized by panic selling, sharp volatility spikes, and widespread losses, placing significant stress on valuation practices and risk-management frameworks.

Distinguishing true market crashes from gradual bear-market declines is essential for understanding financial instability. Crashes feature rapid drawdowns that can be amplified by behavioral contagion and the interconnectedness of modern financial systems (Sarda et al. 2019). Because disruptions in one market can quickly propagate to others, reliable diagnostics are important not only for investors but also for regulators, policymakers, and financial institutions. Accurate crash prediction, therefore, serves broader objectives beyond portfolio management. Early detection frameworks support macro-prudential oversight, improve systemic-risk monitoring, and inform policy responses (Greenwood et al. 2020). For investors, such tools help refine risk-management strategies, reduce behavioral biases, and improve portfolio allocation (Daniel and Moskowitz 2016; Oliva and Renò 2018; Rajpal et al. 2025; Singh et al. 2025). They also contribute to credit risk assessment and stress testing, thereby strengthening institutional resilience to adverse shocks. For emerging markets such as Mexico, where episodes of volatility and sharp reversals recur (Irina et al. 2013), the need for robust diagnostic models is particularly acute.

Although LPPL models have been applied extensively in developed markets and in large emerging economies such as China and Brazil (Sornette 2017), studies on Latin America remain uneven. Some markets lack long, stable price series or are affected by structural breaks and liquidity constraints. Mexico, by contrast, offers a large, well-documented market: the S&P/BMV IPC index is the second largest in the region and provides long-horizon, high-quality data (World Bank n.d.). Its combination of market depth and multiple historical crashes makes it a strong candidate for LPPL-based analysis. Brazil's IBOVESPA, though more widely studied, provides a useful Latin American benchmark, while markets such as Argentina face chronic inflation and inconsistent data. Smaller markets often lack the liquidity or long-run series required for bubble diagnostics.

To broaden the empirical scope and test robustness, this study applies a unified LPPL methodology to three major indices: the S&P/BMV IPC (Mexico), the IBOVESPA (Brazil), and the NYSE Composite (United States). The IPC allows examination of an important but understudied emerging market; the NYSE Composite serves as a developed-market reference with extensive prior literature; and the IBOVESPA provides a second emerging-market case with different structural features (Xiu et al. 2021). A rolling-window early-warning framework complements the main calibrations by tracking how estimated critical times evolve as the window advances.

Earlier prediction techniques relied heavily on linear and parametric models such as ARMA, ARIMA, and GARCH (Bandyopadhyay 2016; Herwartz 2017; Rounaghi and Zadeh 2016). Variants such as wavelet-ARIMA and multivariate GARCH frameworks (Mohammadi and Su 2010) improved performance but still assume linearity or stationarity, limiting their ability to capture the nonlinear dynamics of bubbles. These models struggle with the feedback effects and instability that characterize the approach to crashes. Machine-learning methods like SVM and ANN address some of these limitations and have shown strong predictive performance in certain contexts (Banjade 2020; Chai et al. 2015; Hossain and Nasser 2011; Liu and Hou 2019; Moghaddam et al. 2016; Murkute and Sarode 2015). However, they often require large datasets, remain prone to overfitting, and lack interpretability. Their "black-box" nature limits their usefulness for diagnosing speculative regimes or identifying tipping points in financial markets. The Log-Periodic Power-Law (LPPL) model, often referred to as the Johansen–Ledoit– Sornette (JLS) framework (Johansen and Sornette 1999; Johansen et al. 1999, 2000), offers a theoretically grounded alternative. It models prices as a faster-than-exponential growth process overlaid with log-periodic oscillations, reflecting

the interplay between herding, imitation, and intermittent corrections (Brée and Joseph 2013; Filimonov and Sornette 2013). This structure allows the model to identify unsustainable growth regimes and estimate a critical time t_c at which the bubble becomes unstable. Its behavioral foundation and interpretability make it particularly suitable for studying markets with nonlinear, feedback-driven dynamics.

Given the economic importance and volatility of the Mexican market, this study applies the LPPL model to multiple historical crash periods in the S&P/BMV IPC. The same methodology is extended to the IBOVESPA and the NYSE Composite to assess cross-market robustness and compare bubble signatures across both emerging and developed markets.

The main contributions of this study are as follows:

1. A systematic application of modern LPPLS calibration across three markets—Mexico, Brazil, and the United States—covering multiple historical crash events.
2. A multi-window analysis (180–420 trading days) evaluating model fit, parameter stability, and robustness of critical-time estimates.
3. A rolling early-warning framework illustrating how risk builds as markets approach estimated critical regions.
4. A cross-market comparison highlighting similarities and differences in bubble behavior between emerging and developed markets.

The remainder of the paper is organized as follows: Section 2 reviews related literature; Section 3 outlines the methodology and data; Section 4 presents the empirical results; Section 5 discusses cross-market implications; and Section 6 concludes.

2. Literature Review

2.1. Foundations of Speculative Bubbles and Market Crashes

In financial economics, the concept of speculative bubbles—periods when asset prices consistently and substantially diverge from their fundamental values—has been thoroughly examined. Early theoretical models demonstrated that bubbles may arise even under rational expectations when traders coordinate on self-fulfilling beliefs, allowing prices to rise explosively until a tipping point is reached (Blanchard and Watson 1982; Blanchard 1979; Tirole 1985). Building on these foundations, later empirical research found that stock prices frequently exhibit excessive volatility and rapid acceleration episodes that are not consistent with fundamentals alone (Cutler et al. 1989; West 1987). Taken together, these studies conclude that crashes are nonlinear transitions, typically triggered by herd mentality, brittle expectations, and feedback loops that amplify minor shocks into systemic market collapses. A parallel literature studies the dynamics of crises and volatility across business cycles, noting that price collapses are typically abrupt, clustered, and often disconnected from contemporaneous macroeconomic news (Fama and French 1988). This area of research highlights the challenges of predicting crashes using linear statistical models and encourages the development of nonlinear frameworks that can capture sudden corrections and rapidly rising bubbles.

2.2. Evolution of the Log-Periodic Power-Law (LPPL) Model

To address limitations of linear models, the Log-Periodic Power-Law (LPPL) framework, also known as the Johansen–Ledoit–Sornette (JLS) model (Johansen et al. 1999), was introduced as a behavioral and statistical physics-inspired approach to modeling speculative bubbles. Price dynamics are represented by the LPPL model as a faster-than-exponential growth process with oscillations that become more frequent as the system approaches a critical point (Sornette 2003). These oscillatory "log-periodic" signatures arise from discrete scale invariance and the collective imitation of traders, generating accelerating waves of buying pressure prior to crashes.

Empirical research has applied LPPL to major historical crises, including the 1929 crash, the 1987 Black Monday event, the dot-com bubble, and the 2007–2008 global financial crisis. These studies repeatedly demonstrate that LPPL structures exist in the months prior to significant market collapses (Johansen and Sornette 1999, 2001b). Later refinements improved calibration stability by addressing parameter "sloppiness," the sensitivity of tc estimates, and the emergence of inaccurate fits under noisy market conditions (Brée and Joseph 2013; Filimonov and Sornette 2013). These developments led to the adoption of LPPL, which utilizes collaborative or multi-window techniques to enhance practical reliability and mitigate overfitting. Although there has been disagreement over LPPL's predictive accuracy, many researchers have found that its strength lies in detecting hazard zones—times, when bubble dynamics become more intense and the probability of bursting increases—rather than in predicting precise crash dates. This encourages the application of LPPL as a risk-monitoring or early-warning system as opposed to a black-box predictor.

2.3. Empirical Applications Across Global and Emerging Markets

LPPL has been widely applied across developed equity markets (the US, Europe, and Japan), commodity markets, and currency crises. Multiple research studies demonstrate LPPL-consistent behavior across various resolution scales and sub-oscillations, corroborating substantial log-periodic signatures in Western stock markets during significant events from 2000 to 2010 (Bartolozzi et al. 2005). In emerging markets, applications have been especially prominent in China, where researchers identify both bubble and "anti-bubble" dynamics following extended corrections (Zhou and Sornette 2004). Additionally, crises in Asian and Latin American markets have been analyzed using LPPL (Johansen and Sornette 2001a), demonstrating that log-periodic fingerprints are not exclusive to developed markets. Due to its historical volatility, the availability of long-horizon data, and frequent episodes of financial instability, Brazil's IBOVESPA has been the subject of numerous LPPL studies. However, despite its economic importance, the Mexican stock market remains comparatively understudied in the literature on LPPL. The majority of current research on Mexico uses machine-learning techniques for trend prediction, GARCH-type volatility frameworks, or linear econometric models. Models such as ARIMA, GARCH, Neural Networks, SVMs, and hybrid ML techniques show strengths in short-term forecasting but struggle to detect nonlinear, faster-than-exponential bubble regimes. They are typically insufficient for spotting the onset of speculative bubbles or diagnosing crash-hazard dynamics—exactly the gap LPPL is intended to fill.

2.4. Research Gap

The literature shows broad use of the LPPL framework across markets such as the United States, China, Brazil, Europe, and several Asian economies (Bartolozzi et al. 2005; Johansen and Sornette 2001a; Sornette 2003; Zhou and Sornette 2004). However, a few gaps remain that are directly relevant to this study:

- **Limited LPPL work on the Mexican financial market.**
  Although Mexico is one of the most important markets in Latin America, LPPL applications have focused mainly on the US and other developed markets (Bartolozzi et al. 2005; Sornette 2003), as well as China and Southeast Asia (Johansen and Sornette 2001a; Zhou and Sornette 2004). There is very little evidence of long-horizon LPPL calibration applied to multiple historical crash periods of the IPC index.
- **Lack of multi-window, robustness-oriented LPPL calibration for emerging markets.**
  Much of the earlier LPPL literature relies on a single calibration window, even though prior studies show that single-window fits can lead to unstable estimates of the critical time t_c (Brée and Joseph 2013; Filimonov and Sornette 2013). Given the higher

volatility and structural breaks in emerging markets, a multi-window and ensemble-based approach is particularly important.

- **No comparative LPPL analysis linking Mexico with Brazil and the United States.** While LPPL work exists for Brazil's IBOVESPA (Johansen and Sornette 2001a) and extensive research covers US markets (Bartolozzi et al. 2005), there is no study that examines Mexico alongside both countries. Such a comparison is important for understanding whether bubble dynamics in Mexico resemble those of other Latin American or developed markets.
- **Absence of a rolling early-warning framework for the IPC index.** Although rolling and ensemble-based LPPL early-warning methods have been discussed in the global literature (Filimonov and Sornette 2013; Sornette 2003), they have not been implemented for the Mexican market. No study has applied a trading-day rolling LPPL system to track evolving crash risk in the IPC.

2.5. Contribution of This Study

To address these gaps, this research:

- Applies a robust LPPL calibration to all major historical crash events in the IPC index.
- Implements multi-window (180–420 trading days) fitting and a rolling early-warning ensemble, enhancing the stability assessment of the estimated critical time (t_c).
- Compares LPPL results across Mexico (IPC), Brazil (IBOVESPA), and the United States (NYSE Composite) to position Mexico's bubble dynamics within a wider regional context.
- Provides practical insights for LPPL-based risk monitoring that are relevant to regulators, investors, and policymakers in emerging markets.

3. Design and Methodology

This study uses the S&P/BMV IPC Index (IPC), the main equity benchmark of the Bolsa Mexicana de Valores (BMV), to analyze and predict stock market crashes in Mexico. The IPC is a value-weighted index representing the performance of the 35 most liquid and financially significant companies listed on the BMV, covering major sectors such as consumer goods, industrials, materials, financial services, and telecommunications. Because of its breadth, liquidity, and long historical availability, the IPC serves as the primary indicator of market-wide sentiment, systemic risk, and macro-financial stress in Mexico.

Although the main focus is on the IPC index, the methodology is also applied to the NYSE Composite (United States) and IBOVESPA (Brazil) to assess the robustness of the LPPL framework across markets with different volatility structures. Analysis utilizes daily historical data from 8 November 1991 to 30 January 2025 for the IPC Index and the NYSE Composite Index. For the IBOVESPA index, the period spans from 3 May 1993 to 30 January 2025. Daily historical data were downloaded from the Yahoo Finance website (Yahoo Finance n.d.), a reputable financial market platform that provides real-time data, quotes, charts, financial tools, and analysis. The Adjusted Close series is utilized for all indices when available, as it accounts for corporate events such as dividends, stock splits, and rights issues, thereby maintaining accurate economic returns over time (Malkiel 2019). When Adjusted Close is unavailable, the standard Close price is used to maintain continuity in the series.

3.1. Data Processing

The primary test assets, the BMV IPC, as well as the NYSE Composite and IBOVESPA, when utilized for robustness, had their daily closing values collected. Data are arranged

in ascending order; non-positive prices are removed; and unclear dates are processed in day-first format. No resampling, interpolation, smoothing, or currency conversion is applied; prices remain in the units returned by the source (MXN for IPC on Yahoo Finance). For every estimating window that begins at t_0, a calendar-day time index is created using the formula ti = $Date_i - Date_0$. The LPPL model is fitted independently to Pt and log Pt after the raw price Pt and its natural logarithm log Pt are calculated. Window definitions (lookback lengths and cutoffs) and the forward-only constraint on the critical time t_c are specified in Section 3.4.

3.2. Crash-Identification Method

There are two scenarios in which we may make a misleading claim about a crash (Sornette and Cauwels 2015; Zhou and Sornette 2004). One instance is when the index rises in a bubble and then experiences a significant decline, only to discover that the decline is brief and the bubble persists. The other is when, after a collapse, the index recovers on the way down, allowing us to see the start of a new bubble. However, this rebound is only brief, and the anti-bubble remains alive.

A decrease of less than 10% is frequently referred to as a dip or pullback, and the distinction is purely subjective. A correction is generally understood to be a decline of more than 10% but less than 20%; however, there is no set definition (Advisors 2025; CME Group 2020). To calibrate the LPPL model, it is essential to extract periods corresponding to speculative bubbles and the subsequent crashes from the historical daily data of the S&P/BMV IPC index. Identifying these crashes within the entire historical time frame requires specific filtering criteria. Accordingly, crash dates are filtered based on the following conditions (Brée and Joseph 2013).

- It will take 262 weekdays before the financial security value reaches its peak value
- A drop of 25% or more in the value
- The price must be dropped within 60 working days

To identify the bubble and crash using the above-mentioned conditions, Historical data of BMV/IPC has been filtered by applying the three conditions. The algorithm in the Figure 1 is followed, and implemented using the Python-3.12. programming Language. The LPPL model will be implemented for each of the crash events in order to calibrate the model. The bubbles and crashes following them are found in the historical data of the index.

These rules are used only to define the historical cases for calibration and evaluation; the LPPL fits itself with cutoffs that stop before the identified crash, ensuring ex-ante evaluation.

3.3. Log-Periodic Power Law

The Johansen–Ledoit–Sornette (JLS) model, also known as the Log-Periodic Power-Law (LPPL) model (Johansen et al. 1999), is a nonlinear framework that helps predict possible stock market crashes by identifying speculative bubbles and estimating the critical time at which such bubbles may collapse (Sornette 2003). By fitting the LPPL function to market prices, researchers capture both the faster-than-exponential growth and the log-periodic oscillations characteristic of bubble dynamics. Although effective, the LPPL model is difficult to calibrate due to its strong nonlinearity and sensitivity to initial parameters (Brée and Joseph 2013).

$$y(t) \;=\; A \;+\; B\,(t_c - t)^{\beta}\Big[\,1 + C\,\cos\!\big(\omega\,\log(t_c - t) + \phi\big)\Big] \tag{1}$$

Here, $y(t)$ represents the logarithm of the asset price at time t, which stabilises variance and improves numerical fitting.

The LPPL model includes seven parameters (Table 1), out of which three are linear (A, B, C) and four are nonlinear (t_c, β, ω, ϕ). The nonlinear components are the main focus of LPPL calibration since the linear parameters can be readily solved by regression after the nonlinear parameters are estimated.

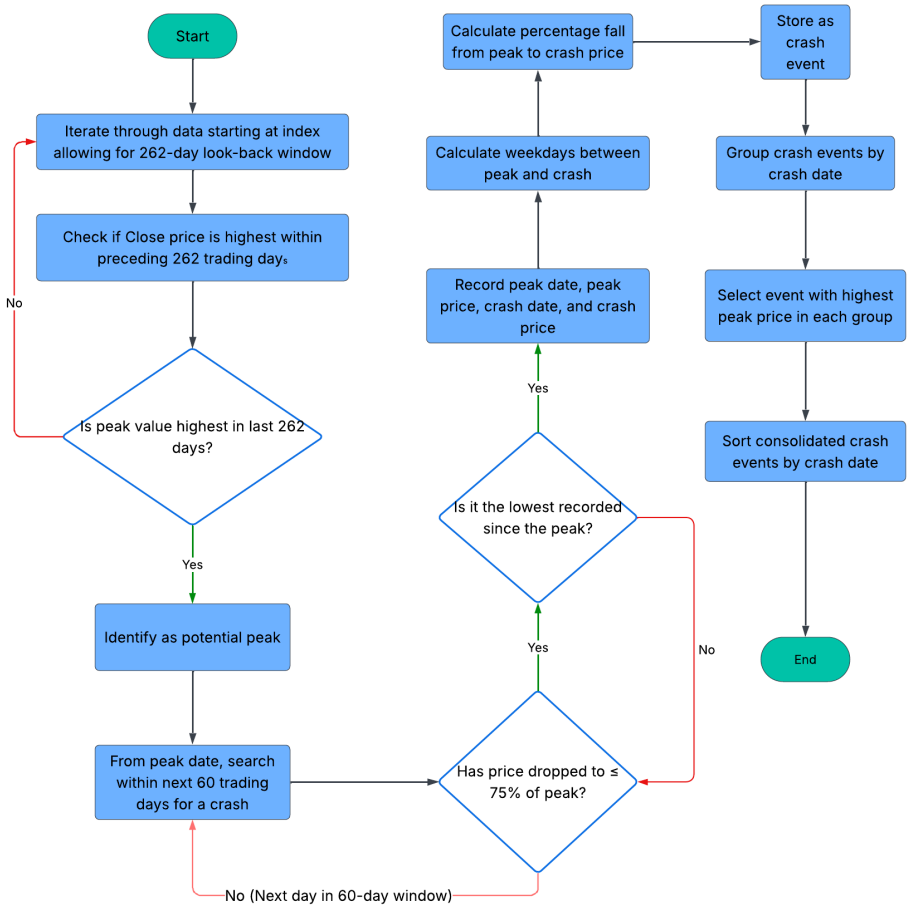

Figure 1. Algorithm flow chart implementing the three conditions used by (Brée and Joseph 2013).

Table 1. Parameters of the Log-Periodic Power-Law (LPPL) model.

Parameter	Description	Effect on Fit
A	Baseline log-price	Shifts the entire curve vertically.
B	Bubble acceleration term	Controls the overall curvature and acceleration of the growth.
C	Amplitude of the oscillations	Larger C produces stronger oscillatory behavior.
t_c	Critical time	Theoretical time of the bubble collapse.
β	Exponent controlling the curvature	Controls how steeply the curve bends near t_c.
ω	Frequency of log-periodic oscillations	Ensures oscillations occur near the critical point.
ϕ	Phase shift of the oscillations	Influences the horizontal alignment of the bumps.

3.4. Log-Periodic Power-Law (LPPL) Fitting Method

Prices move up more quickly and fluctuate with decreasing cycles before a break during a speculative run-up. Both components are captured by the LPPL curve: (i) a climb that is quicker than exponential and aiming towards a ceiling; and (ii) little oscillations that are shorter in duration with time. The LPPL equation used for curve fit is

$$p(t) = A + B\,(t_c - t)^{\beta} + (t_c - t)^{\beta}\Big[C_1 \cos\big(\omega \ln(t_c - t)\big) + C_2 \sin\big(\omega \ln(t_c - t)\big)\Big] \quad (2)$$

In this expression, $p(t)$ denotes the observed market price (or log-price) at the trading-day index t, which is the series being fitted to identify bubble-like behavior. The amplitude–phase transformation of the original (C, ϕ) representation presented in (Johansen et al. 2000) and covered in (Sornette 2003) is represented by the parameters C_1 and C_2, were $C_1 = C\cos\phi$ and $C_2 = C\sin\phi$. This (C_1, C_2) variant improves numerical stability during calibration (Filimonov and Sornette 2013). Here, t_c is the predicted critical time expressed in trading-day units. As $t \to t_c^-$, the factor $(t_c - t)^\beta \to 0$, causing the oscillatory terms to vanish and the trajectory to approach the limiting value A. In this equation, A is the limiting level the curve approaches, $B < 0$ sets the strength of the accelerating rise, $\beta \in (0, 1)$ shapes the curvature of that rise, (C_1, C_2) scale the oscillations, ω says how "fast" the oscillations tick on a log-time axis, and t_c is the model's critical time at which the crash is expected (often a peak or sharp correction). These parameters are listed in Table 1.

3.4.1. Fit Computation

The LPPL curve is fitted to the observed series $y(t)$, where $y(t)$ denotes either the raw price $P(t)$ or the log-price $\log P(t)$. For each time point t_i, $y(t_i)$ is the observed market value and $\hat{p}(t_i; \theta)$ is the LPPL model's predicted value; the residual $r_i(\theta)$ therefore measures their difference. For the parameter vector $\theta = (A, B, \beta, C_1, C_2, \omega, t_c)$, the residual at time t_i is

$$r_i(\theta) = y(t_i) - \hat{p}(t_i; \theta).$$

Calibration is performed using bounded nonlinear least squares via the trust-region reflective algorithm in `SciPy`'s `least_squares` (Virtanen et al. 2020). Trial solutions with $t_c \leq t_i$ are penalized so that the estimated critical time always satisfies $t_c > t_i$.

$$\min_{\theta \in \mathcal{B}} \sum_{i=1}^{n} r_i(\theta)^2 \quad \text{s.t.} \quad t_c - t_i > 0 \; \forall i.$$

3.4.2. Bounds and Ex-Ante Constraint (Trading Days)

In order to avoid bad fits, the critical time is constrained to lie just beyond the window end in trading days and set empirically established bounds (Filimonov and Sornette 2013; Sornette 2003). The bound set utilized in the price and log-price variations is compiled in Table 2. We employ data-scaled limits (tighter for log-price) for (A, B) in order to enhance conditioning while maintaining $B < 0$.

Table 2. Bound set for LPPL calibration on a trading-day time axis.

Parameter	Symbol	Meaning	Lower	Upper	Units/Notes
Exponent	β	Curvature (faster-than-exponential growth)	0.1	0.9	dimensionless
Log-frequency	ω	Speed of log-periodic oscillations	6	13	dimensionless
Pre-crash slope	B	Strength/sign of acceleration	negative	-10^{-6}	data-scaled; enforce $B < 0$
Critical time	t_c	Regime-change time	$t_{\text{end}} + 1$	$t_{\text{end}} + H_{\text{td}}$	$H_{\text{td}} \in \{60, 90\}$ trading days
Level (price fit)	A	Limiting price level	range-based	range-based	window-specific scaling
Level (log fit)	A	Limiting log-price	$p_5 - 1$	$p_{95} + 1$	percentiles of $\log P_t$

3.4.3. Initialization

To ensure stable LPPL calibration, each nonlinear least-squares iteration begins from a simple and robust data-driven initial vector. The initial values are based directly on the data in the window:

$$A_0 = \text{median}(y_t), \quad B_0 = -0.1\,\text{stdev}(y_t),$$

$$\beta_0 = 0.5, \quad (C_{1,0}, C_{2,0}) = (0.05, -0.05),$$

$$\omega_0 = 8.5, \quad t_{c,0} = t_{\text{end}} + 30 \text{ trading days}.$$

This is recommended for stability (Filimonov and Sornette 2013). It usually produces more stable LPPL fits.

3.4.4. Ensuring Feasible Starting Values

LPPL parameters must lie within meaningful ranges, so we define the feasible parameter box:

$$\mathcal{B} = [\ell_A, u_A] \times [\ell_B, u_B] \times [0.1, 0.9] \times [-1, 1] \times [-1, 1] \times [6, 13] \times (t_{\text{end}} + 1,\ t_{\text{end}} + H_{\text{td}}),$$

This box corresponds to the parameter set $(A, B, \beta, C_1, C_2, \omega, t_c)$.

Here,

- ℓ_A and u_A are the lower and upper allowed values for the parameter A (baseline level). These bounds are scaled according to the underlying price series (raw or log).
- ℓ_B and u_B are the lower and upper allowed values for the parameter B (power-law strength). These bounds are also data-dependent.
- ℓ denotes the vector of all lower bounds: $(\ell_A,\ \ell_B,\ 0.1,\ -1,\ -1,\ 6,\ t_{\text{end}} + 1)$.
- u denotes the vector of all upper bounds: $(u_A,\ u_B,\ 0.9,\ 1,\ 1,\ 13,\ t_{\text{end}} + H_{\text{td}})$.

The bounds (ℓ_A, u_A) and (ℓ_B, u_B) depend on whether the model uses raw or log-prices (see Table 3).

The initial parameter vector is:

$x_0 = (A_0, B_0, \beta_0, C_{1,0}, C_{2,0}, \omega_0, t_{c,0})$.

Before optimization starts, x_0 is clipped so that every component lies inside the allowed bounds:

$x_0 \leftarrow \min(\max(x_0, \ell),\ u)$.

This step guarantees that the optimizer starts at a **feasible and meaningful** starting point every time. It prevents issues such as selecting an incorrect critical time (such t_c inside the window) or generating extreme parameter values that lead to numerical instability.

Table 3. LPPL parameter bounds used in estimation for raw-price and log-price fits.

Parameter	Raw-Price Fit	Log-Price Fit
Level A	$[\max(0, \text{p5}(y)),\ 1.5\,\text{p95}(y)]$	$[\text{p5}(y) - 1,\ \text{p95}(y) + 1]$
Pre-crash slope B	$[-10\,\sigma_y,\ -10^{-6}]$	$[-10,\ -10^{-6}]$
Curvature exponent β	$(0.1,\ 0.9)$	
Oscillation parameters $C_1,\ C_2$	$[-1,\ 1]$	
Log-periodic frequency ω	$[7,\ 12]$	
Critical time t_c (trading-day axis)	$(t_{\text{end}} + 1,\ t_{\text{end}} + H]$, with $H = 90$ trading days	

3.4.5. Windowing, Selection, and Reporting

Terminal windows of $L \in \{180, 240, 300, 360, 420\}$ (trading days) are fit in both raw price and log-price variants. When a single display window is required, the "common" window is chosen by the lowest mean RMSE across the two variants. For each fit, the model gives R^2, RMSE, MSE, MAE, and the implied t_c as both a trading-day index and its mapped calendar date (nearest trading day in the sample).

3.4.6. Rolling Ensemble (Trading-Day Lead) and Early-Warning Mapping

Ex-ante warnings are generated by advancing the window endpoint over the last 120 trading days in 15 day increments, and calculating the price and log-price for each endpoint fit $L \in \{180, 240\}$. Only fits that meet the criteria $(0.1 < \beta < 0.9,\ 6 \leq \omega \leq 13,\ B < 0)$ are kept. The median is used to aggregate the approved t_c values (in trading days) at each endpoint. The trading-day lead time is defined as follows: $\Delta_{\text{td}} = t_c^{\text{med}} - t_{\text{end}}$.

Bands utilized in the early-warning panel are matched to risk:

$$\text{Very High if } \Delta_{\text{td}} \leq 7, \quad \text{High if } 7 < \Delta_{\text{td}} \leq 14, \quad \text{Elevated if } 14 < \Delta_{\text{td}} \leq 30.$$

Bands are displayed as shaded spans centered on each endpoint date (for readability, spans are shown in calendar time but computed from trading-day thresholds). This ensemble reframes LPPL from a single-date predictor into a practical hazard indicator, consistent with robust calibration practice (Filimonov and Sornette 2013).

3.4.7. Bounds and Sign Constraints

All parameters are box-constrained during optimization. Table 3 summarizes the bounds used in the implementation.

Here p5, p95 are within-window percentiles and σ_y is the within-window standard deviation. A strict sign filter $B < 0$ is enforced post-fit to focus on faster-than-exponential regimes.

3.4.8. Code Implementation Algorithm

The algorithm given in Figure 2 has been used to implement the LPPL fitting for both the log-price data and the raw price data using the LPPL model.

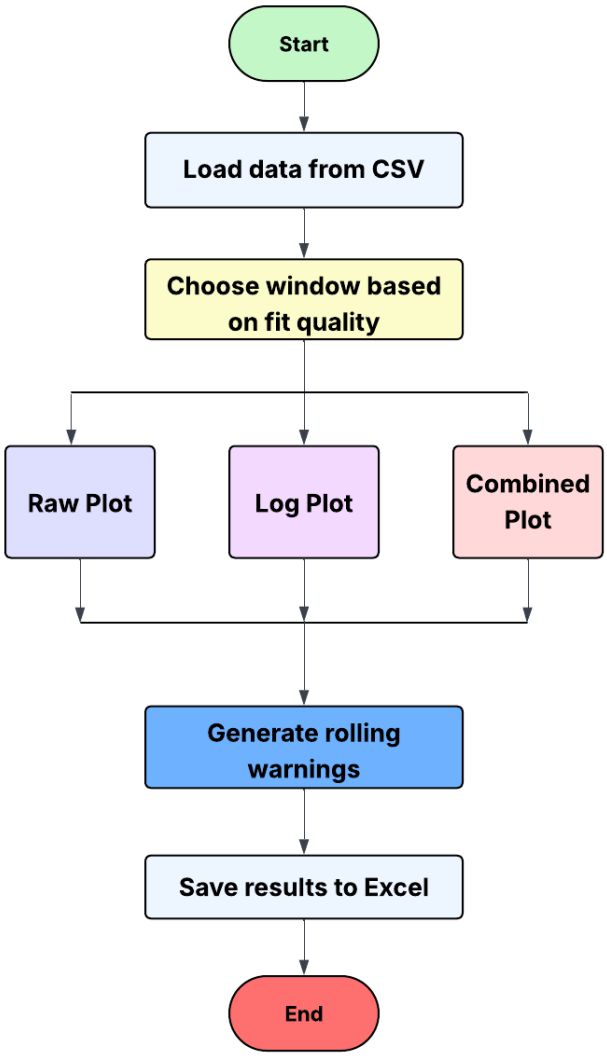

Figure 2. LPPL pipeline flow chart used in this study.

Implementation note

All seven parameters $(A, B, \beta, C_1, C_2, \omega, t_c)$ are estimated jointly via single-stage bounded NLS with the (C_1, C_2) parameterization (Filimonov and Sornette 2013). The functional form is the standard LPPL equation (Johansen et al. 2000; Sornette 2003); only the time scale (trading-day index) and the oscillatory-term parameterization differ from some expositions.

3.5. LPPL Model Calibration

For calibrating the LPPL model, the study has done the LPPL fitting for each of the crashes identified using the conditions mentioned above. The timeline for each crash case has been taken in the following order:

- Case 1: From the start of the data to the date of the 1st crash
- Case 2: From the next day of crash 1 to the date of the 2nd crash
- Case 3: From the next day of crash 2 to the date of the 3rd crash, and so on.

3.6. Plotting Conventions Used

- Raw panel: Price in black; LPPL fit in red.
- Log panel: log-price in green; LPPL fit in orange.
- Combined panel: Raw (black) + raw LPPL (red) on the left axis; log (green) + log-LPPL (orange) on the right axis.
- The fitted curve is drawn only up to $\min(t_{\text{end}} + \Delta,\ t_c^-)$, i.e., it never crosses the singularity at t_c (Sornette 2003).
- To avoid ex-post anchoring, no vertical t_c line is drawn in the first three panels; timing information is conveyed in the rolling early-warning panel.

3.6.1. Rationale for Dual (Raw/Log) Views

The log transform is recommended for fitting because financial prices often exhibit multiplicative noise and relative changes are more informative than absolute changes; the raw-price view is retained because stakeholders frequently interpret predictions in price units (Cont 2001; Sornette 2003). Showing both helps assess fit stability (log) and practical relevance (raw).

3.6.2. Time Scale and Smoothing Notes

All timing logic (window length, lookback, step, and the t_c horizon) is implemented in trading days. Calendar dates appear on the horizontal axes for readability and when reporting the mapped t_c date. For the rolling panel, the price curve is lightly smoothed (EMA in event mode; weekly mean in full-timeline mode) to improve readability without altering fit results (Filimonov and Sornette 2013).

3.7. Model Validation

To check the reliability of the LPPL estimates, we apply the following validation steps:

- **Blocked cross-validation:** The series is divided into contiguous train–validation segments (not random folds) to preserve temporal order. Each fit is trained on past data and evaluated on a later block, reducing look-ahead bias and helping assess out-of-sample stability (Bergmeir and Benítez 2012).
- **Visual inspection:** Fitted LPPL curves are compared with the observed series to verify that the model captures the accelerating trend and characteristic oscillations expected before a regime shift (Sornette 2003).

- **Residual and stability checks:** Residuals are examined for obvious structure or drift. Stability of the critical-time estimate is assessed using the interquartile range

$$\mathrm{IQR}(t_c) = Q_{0.75}(t_c) - Q_{0.25}(t_c),$$

  which summarizes how tightly the ensemble of accepted fits clusters around a common timing (Filimonov and Sornette 2013).
- **Fit metrics:** RMSE, MAE, and R^2 are reported as compact indicators of in-sample fit quality. These metrics assist interpretation but are not used to make point forecasts, since LPPL behaves primarily as a hazard model rather than a price-level predictor.

When a single illustrative window is required, we display the calibration with the lowest mean RMSE across the raw and log specifications.

3.8. Early-Warning Scoring (Event Study)

Bands in the rolling panel are computed exactly as in Section 3.4 (trading-day aggregation and thresholds). An early warning is credited when a High/Very High band appears before the realised event and the median t_c lies within the specified ($\pm$) trading-day tolerance of the event date.

3.9. Software and Tools

All computations in this study were performed in Python using the standard scientific ecosystem and custom code written by the authors. The workflow relied on (NumPy n.d.) for numerical routines and vectorized array operations, pandas for dataset management, date handling, and resampling (Wes McKinney 2010), and SciPy's `least_squares` implementation (trust-region reflective algorithm) for bounded nonlinear fitting of the LPPL model. All figures—including raw/log/combined overlays and rolling early-warning panels—were produced with (Matplotlib n.d.), while XlsxWriter was used to generate structured Excel workbooks containing window-wise parameters and rolling-warning summaries to ensure full auditability. Interactive exploration and iterative calibration were carried out in Jupyter Notebook. The entire analysis runs on a standard CPU-based laptop or desktop, with no requirement for GPU acceleration. This open-source toolchain provides a fully reproducible environment and aligns with current LPPL modeling practice, offering transparent numerical optimization, stable time-series management, and publication-quality visual output.

4. Results and Discussions

4.1. Identified Crash Events and Analysis

A total of eight major crash events were identified across the three markets, as specified by (Brée and Joseph 2013). These include three events in the IPC/BMV (Mexico), four in the IBOVESPA (Brazil), and one in the NYSE Composite (USA). The complete set of peak dates, crash dates, draw-down magnitudes, and window information is consolidated in Table 4, providing a unified overview of all detected bubble–crash events. A three-panel visualization of the historical price movements for each index is presented in Figure 3 to supplement the tabular summary. Panels (A), (B), and (C) represent the IPC, IBOVESPA, and NYSE Composite, respectively. The marked peak and crash points demonstrate the consistency of the event-identification approach across markets with varying economic circumstances and volatility structures.

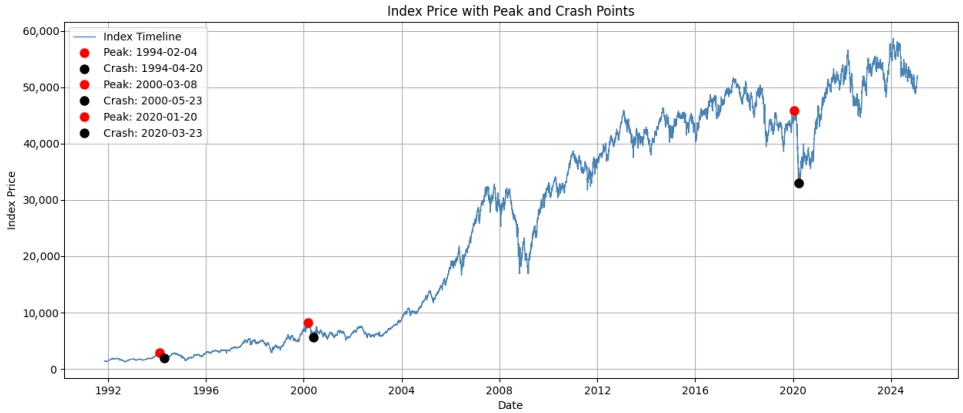

(**a**) Historical Crash in IPC Index (Mexico)

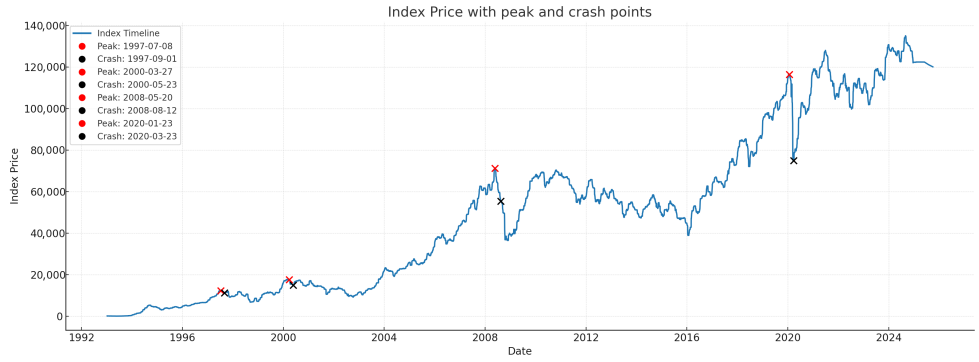

(**b**) Historical Crash in IBOVESPA Index (Brazil)

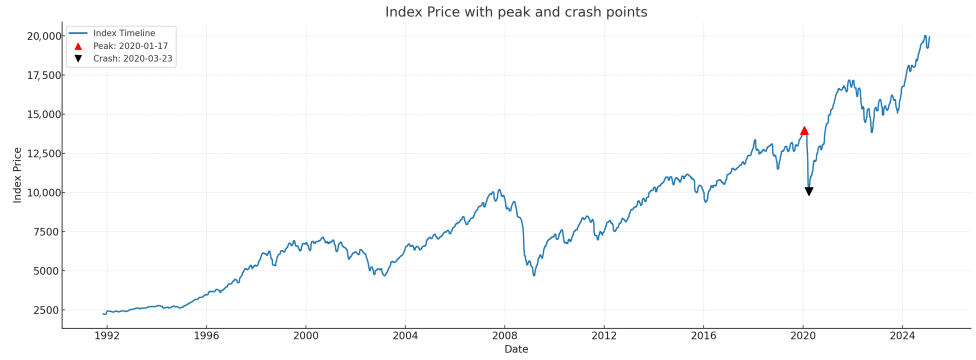

(**c**) Historical Crash in NYSE Composite Index (USA)

Figure 3. Historical peak and crash dates identified in the three equity indices using the criteria of (Brée and Joseph 2013). Panel (**a**) shows the IPC (Mexico), Panel (**b**) displays IBOVESPA (Brazil), and Panel (**c**) shows the NYSE Composite (USA).

Table 4. Historical crash events for IPC/BMV, IBOVESPA, and NYSE Composite indices based on peak and crash dates, showing percentage declines and window information.

Market	Event	Peak Date	Peak Close	+60 Days	−262 Days	Crash Date	Crash Close	Fall (%)
IPC (Mexico)	1	1994-02-04	2826.00	1994-05-04	1993-01-28	1994-04-20	1961.10	30.60
IPC (Mexico)	2	2000-03-08	8295.20	2000-06-06	1999-03-03	2000-05-23	5593.58	32.56
IPC (Mexico)	3	2020-01-20	45,902.68	2020-04-17	2019-01-04	2020-03-23	32,964.22	28.18
IBOVESPA (Brazil)	1	1997-07-08	13,617.00	1997-09-30	1996-06-19	1997-09-01	10,109.00	23.22
IBOVESPA (Brazil)	2	2000-03-27	18,951.00	2000-06-09	1999-03-09	2000-05-23	13,587.00	28.30
IBOVESPA (Brazil)	3	2008-05-20	73,517.00	2008-08-13	2007-04-27	2008-08-12	54,503.00	25.86
IBOVESPA (Brazil)	4	2020-01-23	11,9528.00	2020-04-23	2020-01-04	2020-03-23	63,570.00	46.81
NYSE (USA)	1	2020-01-17	14,183.20	2020-04-14	2018-12-31	2020-03-23	8777.40	38.11

Although some major crashes did not meet the criteria, the eight identified crashes are also notable in their respective indexes. These events align with documented market crashes, giving the mechanically identified events external face credibility before the LPPL calibration. The 1994 event (Crash of 1994-04-20) of the IPC (Mexico) is associated with a time of macro and domestic politics instability, the 2000 event (Crash of 2000-05-23), which reflects the global dot-com unwind, which propagated to emerging markets, and the 2020 event (Crash 2020-03-23), which represents the worldwide COVID-19 shock (Shu 2024). In IBOVESPA (Brazil), the 1997 event (Crash 1997-09-01) depicts the effects of the Asian financial crisis, the 2000 event (Crash of 2000-05-23) is associated with the dot-com correction, the 2008 episode (Crash of 2008-08-12) is associated with the global financial crisis and a decline in commodities, and the 2020 event (Crash of 2020-03-23), which is the COVID-19 crash. Last but not least, the COVID-19 collapse is associated with the 2020 event (the crash of 2020-03-23), which is in the NYSE Composite (USA). It is important to note that although the historical timelines for these indices contain other significant crashes, such as the 1997 Asian Financial Crisis, the dot-com bubble burst (2000), the crash following the 9/11 attacks (2001), or the 2008 global financial crisis, they were not identified as 'crash events' as they failed to meet the three conditions. Only events that strictly fit the description of a rapid and severe descent from a new high were chosen for the study. This was carried out to test and calibrate the model on crash events, which exhibit some uniformity, as each of the identified crash dates satisfied the three conditions given by (Brée and Joseph 2013). This selective approach provides a consistent and objective set of case events for the subsequent testing of the LPPL model.

4.2. Overview of LPPL Fitting on Identified Crash Events

For each of the eight identified crash events, the modified Log-Periodic Power-Law (LPPL) model (Filimonov and Sornette 2013) is fitted on a trading-day time axis using robust multi-start nonlinear least squares. Fits are computed on rolling windows of length $L \in \{180, 240, 300, 360, 420\}$ trading days under the constraints $B < 0$, $0.1 < \beta < 0.9$, and $7 \leq \omega \leq 12$. Both raw-price and log-price variants are estimated: the former preserves economic scale, while the latter stabilises variance and typically yields smoother oscillatory components and tighter clustering of estimated critical times t_c. Model credibility is assessed by the consistency of t_c across windows and by the pattern of rolling early-warning signals.

4.2.1. IPC Index, Mexico

The IPC/BMV index exhibits three major bubble–crash episodes over the sample: the 1994 crisis, the 2000 dot-com spillover, and the 2020 COVID-19 crash. Event 1 (1991–1994) shows a strong rise in 1993 that was followed by increasing instability due to internal political shocks and rising interest rates in the United States. Event 2 (1994–2000) coincides with the global dot-com boom and subsequent unwind, during which external financial conditions transmitted stress to Mexican equities. Event 3 (2018–2020) shows the sharp downturn that followed the COVID-19 pandemic, which was influenced by pressure on global liquidity, falling oil prices, and major Mexican peso selling. For each event, the LPPL analysis evaluates the multi-window fits, selects a baseline specification, and examines the rolling early-warning behavior leading into the observed break.

Across the three IPC events, the baseline calibrations presented in Table 5 show a consistent trend: log-scale fits clearly beat raw fits in both accuracy and stability. Event 1 (L = 240) produces strong and balanced fits, with log achieving the highest R^2 (0.9025) and a t_c only slightly beyond the actual break. Event 2 (L = 420) yields the best overall explanatory power, with the log variant reaching $R^2 = 0.914$ and placing t_c within a few

weeks of the 2000 crash. Even though Event 3's raw fits are weaker due to the external COVID-19 shock, the log baseline (L = 180) continues to produce a reliable signal with minimal errors and a t_c consistent with the March–April 2020 crash window. Overall, the table indicates that log specifications give the most reliable LPPL evidence across all three episodes.

The combined LPPL graphs (raw + log) for the three IPC Mexico events exhibit a uniform bubble formation pattern characterized by faster-than-exponential growth followed by a noticeable reduction of curvature as each crash draws closer, as shown in Figure 4. The combined curve for Event 1 (1994) clearly shows the early 1994 flattening, corresponding with the reversal in April, and the progressive acceleration in 1993. The combined graph nevertheless shows the typical LPPL rising and cresting behavior before the May 2000 dot-com-related break, with the fitted peak falling roughly around the observed crash timeframe, even if Event 2 (2000) shows a longer, more turbulent expansion. Event 3 (2020), although dominated by an external COVID-19 shock, exhibits a gradual LPPL-type curvature through late 2019 and early 2020, indicating rising fragility before the March collapse. Overall, the three panels show that the LPPL structure seems resilient in both externally exacerbated downturns and endogenous build-ups.

Despite the different macro factors of the three Mexican crashes, the rolling early-warning results for Events 1–3 (Figures 4 and 5; Table 6) demonstrate a consistent LPPL-based buildup of risk. Event 1 demonstrates a typical endogenous bubble: during November and December 1993, warnings increase from Elevated to Very High, and within a few days of the breach on 20 April 1994, they cluster together. The rolling bands show a significant pre-crash acceleration that coincides with the convergence of median t_c estimates between late December and early February. During the global tech boom of late 1999 and early 2000, Event 2 shows a larger, externally reinforced bubble. With Very High signals in mid-January and late-February 2000 and closely clustered t_c projections around February–March, four actionable alarms occur precisely inside the 60-day timeframe before the crash on 23 May 2000. Event 3 exhibits LPPL-type instability while being driven by the external COVID shock: Very High warnings in November and late December 2019, followed by many Elevated bands through February 2020, narrow t_c to mid-March–early-April, surrounding the collapse on 23 March 2020. Together, these results demonstrate that IPC Mexico repeatedly displayed characteristic LPPL pre-crash acceleration, and the rolling system provided timely, escalating alerts—strongly aligned with each subsequent market reversal.

Table 5. Baseline LPPL calibration for IPC/BMV (Events 1–3), showing selected window fits for raw and log specifications.

Event	L (TD)	Window Start	Window End	Variant	R^2	RMSE	MAE	$\hat{t}_c$
Event 1	240	23/04/1993	06/04/1994	raw	0.8980	126.5832	87.5968	29/06/1994
Event 1	240	23/04/1993	06/04/1994	log	0.9025	0.0584	0.0459	24/06/1994
Event 2	420	04/09/1998	09/05/2000	raw	0.8710	473.2685	373.777	14/07/2000
Event 2	420	04/09/1998	09/05/2000	log	0.9140	0.0729	0.05726	23/06/2000
Event 3	180	21/06/2019	06/03/2020	raw	0.3335	1254.9990	878.1074	14/04/2020
Event 3	180	21/06/2019	06/03/2020	log	0.6357	0.0219	0.0173	09/04/2020

Table 6. IPC Index, Mexico—Combined rolling-window early-warning summary for Events 1–3.

Event	Window End	N Fits	Median t_c	t_c p25–p75	Lead (TD)	Risk
1	26/10/1993	2	01/01/1994	23/12/1993–10/01/1994	47.24	Beyond60
1	17/11/1993	2	28/12/1993	18/12/1993–07/01/1994	29.32	Elevated
1	08/12/1993	2	11/12/1993	10/12/1993–12/12/1993	2.00	Very High
1	29/12/1993	2	01/01/1994	31/12/1993–02/01/1994	1.77	Very High
1	20/01/1994	2	02/02/1994	30/01/1994–05/02/1994	8.84	High
1	10/02/1994	2	24/02/1994	20/02/1994–28/02/1994	10.63	High
1	03/03/1994	2	09/04/1994	08/04/1994–10/04/1994	22.15	Elevated
1	28/03/1994	2	05/06/1994	28/05/1994–13/06/1994	48.35	Beyond60
2	29/11/1999	3	08/02/2000	03/02/2000–14/02/2000	50.63	Beyond60
2	20/12/1999	3	19/02/2000	09/02/2000–01/03/2000	43.75	Beyond60
2	11/01/2000	4	20/02/2000	20/02/2000–20/02/2000	29.12	Very High
2	01/02/2000	3	18/03/2000	18/03/2000–18/03/2000	32.06	High
2	22/02/2000	2	26/02/2000	25/02/2000–27/02/2000	3.47	Very High
2	14/03/2000	2	19/04/2000	02/04/2000–05/05/2000	23.50	Elevated
2	05/04/2000	2	03/06/2000	02/06/2000–04/06/2000	37.82	Beyond60
2	28/04/2000	2	27/06/2000	27/06/2000–27/06/2000	40.47	Beyond60
3	27/09/2019	2	24/12/2019	24/12/2019–24/12/2019	59.98	Beyond60
3	18/10/2019	2	16/01/2020	16/01/2020–16/01/2020	60.00	Beyond60
3	08/11/2019	2	11/11/2019	11/11/2019–11/11/2019	1.00	Very High
3	02/12/2019	2	25/02/2020	25/02/2020–25/02/2020	57.29	Beyond60
3	24/12/2019	2	04/01/2020	30/12/2019–09/01/2020	6.85	Very High
3	16/01/2020	2	13/02/2020	30/01/2020–27/02/2020	20.08	Elevated
3	07/02/2020	2	12/03/2020	01/03/2020–23/03/2020	23.47	Elevated
3	28/02/2020	2	01/04/2020	27/03/2020–05/04/2020	21.95	Elevated

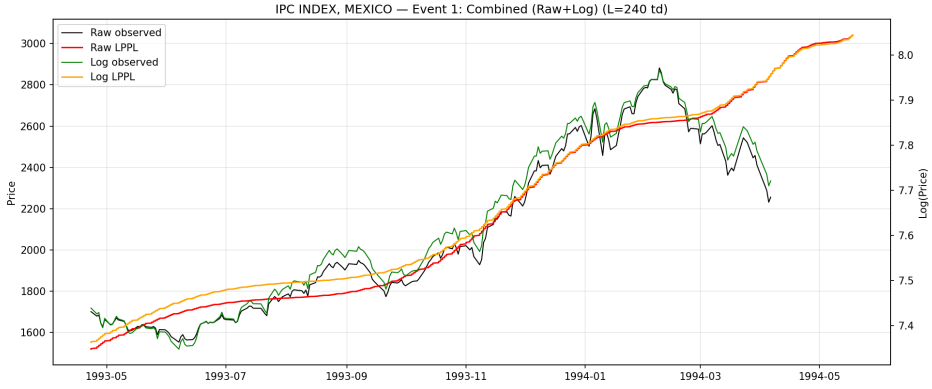

(**a**)

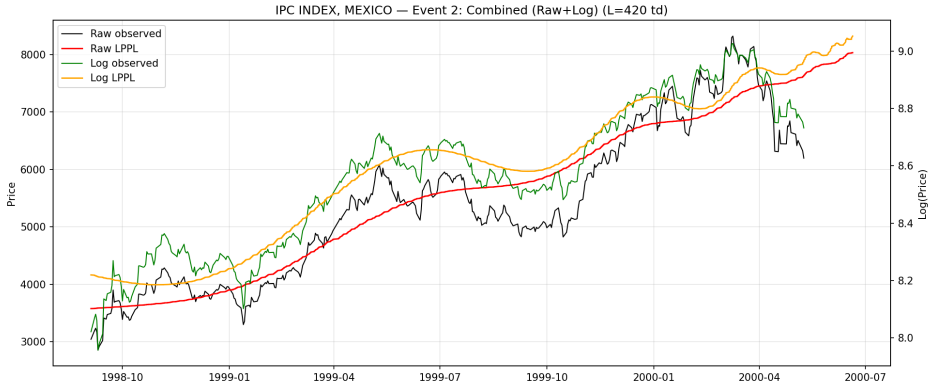

(**b**)

Figure 4. *Cont.*

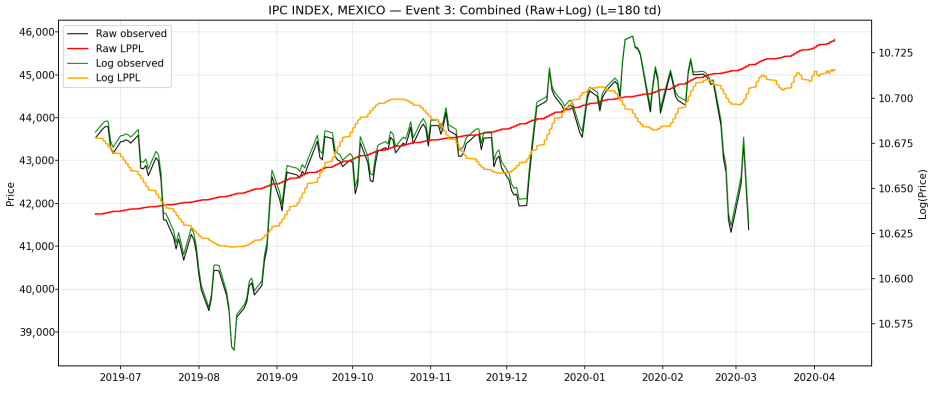

(**c**)

Figure 4. Combined (Raw + Log) LPPL fits for IPC/BMV, Mexico across Events 1–3. Each panel overlays the raw and log-LPPL trajectories with observed prices, showing the approach toward the estimated critical region. (**a**) IPC Mexico—Event 1 (Combined); (**b**) IPC Mexico—Event 2 (Combined); (**c**) IPC Mexico—Event 3 (Combined).

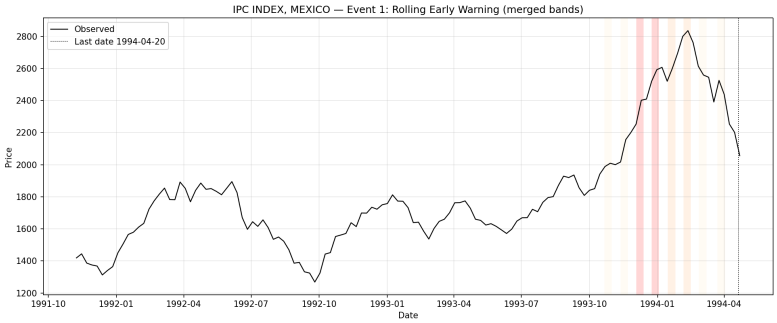

(**a**)

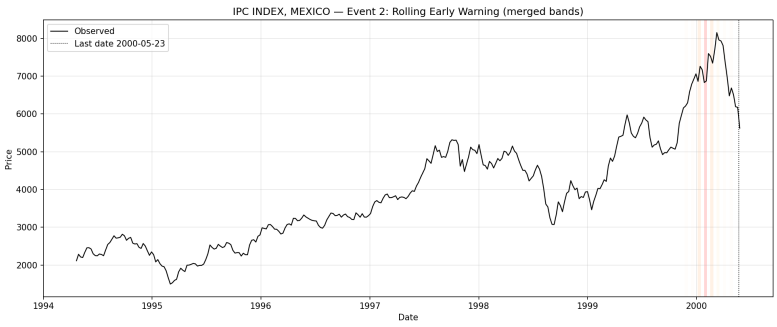

(**b**)

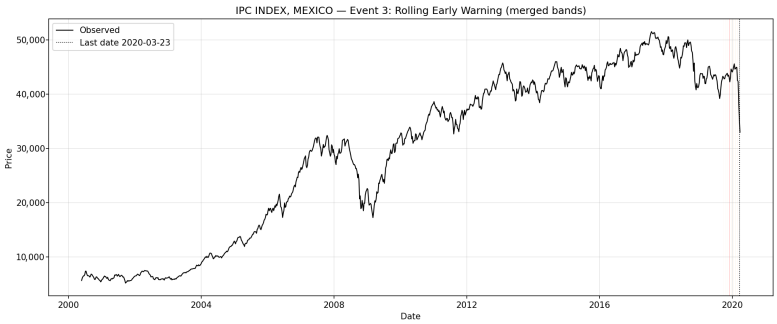

(**c**)

Figure 5. Rolling early-warning signals for IPC/BMV Mexico across Events 1–3. Each panel visualizes merged LPPL warning bands (Very High, High, Elevated, Beyond 60) leading up to the observed crash date. (**a**) IPC Mexico—Event 1 (Rolling Early Warning); (**b**) IPC Mexico—Event 2 (Rolling Early Warning); (**c**) IPC Mexico—Event 3 (Rolling Early Warning).

Table 7 shows that, for all three IPC events, log specifications systematically outperform raw fits, with higher R^2, smaller RMSE, and well-behaved β and ω. The projected t_c values closely cluster around the actual crash dates (April 1994, May–June 2000, and March–April 2020) across window lengths, indicating the LPPL calibration's internal resilience. Although log models continue to indicate a cohesive crucial window, raw fits significantly worsen for Event 3, indicating that LPPL structure endures even throughout the COVID-19 shock event.

Table 7. IPC Index, Mexico (Events 1–3): LPPL results across candidate windows and variants (all L included).

Event	L (TD)	Window Start	Window End	Variant	β	ω	R^2	RMSE	MAE	t_c
1	180	19/07/1993	06/04/1994	raw	0.900	10.944	0.793	159.323	108.737	10/05/1994
1	180	19/07/1993	06/04/1994	log	0.900	7.000	0.871	0.056	0.043	24/06/1994
1	240	23/04/1993	06/04/1994	raw	0.900	7.091	0.898	126.583	87.597	29/06/1994
1	240	23/04/1993	06/04/1994	log	0.900	7.000	0.903	0.058	0.046	24/06/1994
1	300	25/01/1993	06/04/1994	raw	0.874	11.246	0.842	155.458	111.892	09/05/1994
1	300	25/01/1993	06/04/1994	log	0.900	10.281	0.873	0.067	0.051	19/04/1994
1	360	23/10/1992	06/04/1994	raw	0.252	12.000	0.788	172.920	141.555	10/05/1994
1	360	23/10/1992	06/04/1994	log	0.756	9.089	0.825	0.075	0.063	25/04/1994
1	420	28/07/1992	06/04/1994	raw	0.261	10.452	0.823	165.987	134.460	04/05/1994
1	420	28/07/1992	06/04/1994	log	0.248	12.000	0.855	0.075	0.063	15/06/1994
2	180	19/08/1999	09/05/2000	raw	0.900	8.913	0.681	575.521	382.868	23/06/2000
2	180	19/08/1999	09/05/2000	log	0.900	11.121	0.797	0.073	0.057	27/07/2000
2	240	27/05/1999	09/05/2000	raw	0.900	10.492	0.703	535.718	443.086	19/07/2000
2	240	27/05/1999	09/05/2000	log	0.900	10.888	0.823	0.066	0.050	27/07/2000
2	300	01/03/1999	09/05/2000	raw	0.900	7.453	0.715	532.832	426.807	12/06/2000
2	300	01/03/1999	09/05/2000	log	0.900	7.000	0.851	0.063	0.046	23/05/2000
2	360	02/12/1998	09/05/2000	raw	0.900	9.881	0.836	492.630	388.205	19/07/2000
2	360	02/12/1998	09/05/2000	log	0.899	8.642	0.913	0.065	0.052	28/06/2000
2	420	04/09/1998	09/05/2000	raw	0.900	9.645	0.871	473.269	373.777	14/07/2000
2	420	04/09/1998	09/05/2000	log	0.794	8.353	0.914	0.073	0.057	23/06/2000
3	180	21/06/2019	06/03/2020	raw	0.900	10.963	0.334	1254.999	878.107	14/04/2020
3	180	21/06/2019	06/03/2020	log	0.900	10.516	0.636	0.022	0.017	09/04/2020
3	240	26/03/2019	06/03/2020	raw	0.900	10.204	0.078	1375.203	988.478	15/04/2020
3	240	26/03/2019	06/03/2020	log	0.900	7.854	0.561	0.022	0.018	18/03/2020
3	300	27/12/2018	06/03/2020	raw	0.877	8.563	0.070	1285.012	938.789	01/05/2020
3	300	27/12/2018	06/03/2020	log	0.591	7.000	0.357	0.025	0.020	12/03/2020
3	360	28/09/2018	06/03/2020	raw	0.118	9.898	−0.052	1829.909	1247.919	23/03/2020
3	360	28/09/2018	06/03/2020	log	0.899	11.134	0.277	0.035	0.028	23/04/2020
3	420	06/07/2018	06/03/2020	raw	0.100	8.511	−0.069	2741.002	1915.435	18/03/2020
3	420	06/07/2018	06/03/2020	log	0.900	10.111	0.581	0.038	0.032	09/04/2020

4.2.2. IBOVESPA Index, Brazil

For Brazil's IBOVESPA index, four major crash events are examined, each embedded in a distinct global and domestic context. Event 1, which began on 5 March 1993 and concluded on 1 September 1997, corresponds with growing stress preceding the Asian financial crisis; the late-summer sell-off is consistent with emerging-market risk aversion spreading to Brazil. Event 2, which spans from 2 September 1997 to 23 May 2000, peaks just after the peak of the global dot-com market, following the devaluation of the BRL in 1999 and a dramatic recovery in EM stocks. Event 3, which spans from 24 May 2000 to 12 August 2008, reverses prior commodity-driven gains in a weakening real environment, aligning the August 2008 fall with the pre-Lehman phase of the global financial crisis. Event 4 shows the COVID-19 crash, which ended on 23 March 2020, when Brazil's stock market suddenly fell from all-time highs in the face of increasingly tightening international financial circumstances. The following analysis creates rolling early-warning diagnostics, baseline fits, and multi-window LPPL calibrations for each event (summarized for four events in Table 8).

Across all four IBOVESPA events, the selected baselines show strong and internally consistent LPPL structure, shown in Table 9. Log specifications systematically dominate raw fits, with much lower RMSE and higher R^2 (up to 0.978 in Event 1 and above 0.90 in Events 2

and 4), while raw remains acceptable but carries large level-scale errors. In every episode, the estimated t_c lies close to the historical break, typically a few weeks after the crash, matching the LPPL interpretation of t_c as the end of the bubble regime rather than the exact price peak. Using long horizons (360–420 TD) also gives comparable baselines across crises, indicating stable parameters and a persistent LPPL signature in the Brazilian market over 1997, 2000, 2008, and 2020.

Table 8. IBOVESPA Index, Brazil—LPPL All-Window Results Across Events 1–4.

Event	L (TD)	Start	End	Variant	β	ω	R^2	RMSE	t_c
1	180	22/11/1996	18/08/1997	raw	0.900	11.263	0.934	520.928	02/09/1997
1	180	22/11/1996	18/08/1997	log	0.900	10.247	0.956	0.044	05/11/1997
1	240	28/08/1996	18/08/1997	raw	0.900	11.450	0.937	563.265	09/10/1997
1	240	28/08/1996	18/08/1997	log	0.900	10.524	0.975	0.039	05/11/1997
1	300	04/06/1996	18/08/1997	raw	0.537	7.363	0.926	628.025	26/08/1997
1	300	04/06/1996	18/08/1997	log	0.900	7.277	0.972	0.044	03/09/1997
1	360	07/03/1996	18/08/1997	raw	0.341	7.206	0.944	576.249	10/09/1997
1	360	07/03/1996	18/08/1997	log	0.536	9.530	0.978	0.044	05/11/1997
1	420	05/12/1995	18/08/1997	raw	0.312	11.017	0.947	572.033	01/09/1997
1	420	05/12/1995	18/08/1997	log	0.383	9.695	0.971	0.054	05/11/1997
2	180	17/08/1999	09/05/2000	raw	0.900	7.325	0.695	1520.923	28/07/2000
2	180	17/08/1999	09/05/2000	log	0.900	7.000	0.780	0.091	07/06/2000
2	240	21/05/1999	09/05/2000	raw	0.900	7.622	0.742	1453.102	24/05/2000
2	240	21/05/1999	09/05/2000	log	0.900	7.000	0.837	0.083	13/06/2000
2	300	23/02/1999	09/05/2000	raw	0.900	7.224	0.742	1447.654	01/08/2000
2	300	23/02/1999	09/05/2000	log	0.900	7.000	0.863	0.077	13/06/2000
2	360	20/11/1998	09/05/2000	raw	0.900	8.081	0.820	1411.947	12/05/2000
2	360	20/11/1998	09/05/2000	log	0.900	7.000	0.884	0.094	13/06/2000
2	420	25/08/1998	09/05/2000	raw	0.895	7.863	0.867	1321.127	29/05/2000
2	420	25/08/1998	09/05/2000	log	0.857	7.178	0.907	0.098	13/06/2000
3	180	31/10/2007	29/07/2008	raw	0.892	8.196	0.088	3784.436	21/10/2008
3	180	31/10/2007	29/07/2008	log	0.100	7.000	0.486	0.044	16/10/2008
3	240	06/08/2007	29/07/2008	raw	0.897	7.094	0.333	4133.321	03/09/2008
3	240	06/08/2007	29/07/2008	log	0.194	8.136	0.568	0.055	16/10/2008
3	300	10/05/2007	29/07/2008	raw	0.900	7.000	0.552	3760.540	22/09/2008
3	300	10/05/2007	29/07/2008	log	0.160	8.052	0.687	0.052	16/10/2008
3	360	12/02/2007	29/07/2008	raw	0.900	10.028	0.719	3928.244	30/07/2008
3	360	12/02/2007	29/07/2008	log	0.900	7.080	0.813	0.057	16/10/2008
3	420	10/11/2006	29/07/2008	raw	0.900	7.449	0.814	3724.216	19/08/2008
3	420	10/11/2006	29/07/2008	log	0.900	10.721	0.829	0.066	16/10/2008
4	180	18/06/2019	09/03/2020	raw	0.887	9.171	0.488	4372.670	10/03/2020
4	180	18/06/2019	09/03/2020	log	0.100	7.000	0.546	0.0383	27/05/2020
4	240	22/03/2019	09/03/2020	raw	0.874	7.287	0.707	4032.579	10/03/2020
4	240	22/03/2019	09/03/2020	log	0.900	11.738	0.743	0.0362	30/03/2020
4	300	19/12/2018	09/03/2020	raw	0.810	9.842	0.741	3960.696	10/03/2020
4	300	19/12/2018	09/03/2020	log	0.900	11.739	0.789	0.0346	30/03/2020
4	360	20/09/2018	09/03/2020	raw	0.852	10.832	0.842	3820.879	10/03/2020
4	360	20/09/2018	09/03/2020	log	0.900	11.655	0.869	0.0352	07/05/2020
4	420	26/06/2018	09/03/2020	raw	0.875	7.765	0.884	4070.039	10/03/2020
4	420	26/06/2018	09/03/2020	log	0.900	11.000	0.915	0.0317	04/05/2020

The combined LPPL graphs shown in Figure 6 show a consistent bubble-type acceleration throughout all four IBOVESPA events, with prices following a smooth super-exponential trajectory that flattens and bends as the index draws closer to each crash. In Event 1, the model captures the tightening oscillatory structure leading into the 1997 decline during the Asian crisis spillover. With the LPPL route converging towards the turning point after the strain on developing economies during the dot-com boom, Event 2 demonstrates a distinct run-up towards the 2000 high. Ahead of the 2008 panic, Event 3 shows a clear late-cycle steepening, with the LPPL trajectory accurately bending as global risk aversion increased. The COVID-19 crash represented by Event 4, which exhibits a steady increase through 2019 until a noticeable LPPL levelling just before the bursting in March 2020. When seen together, the four panels show that Brazil consistently displayed

LPPL-consistent pre-crash dynamics, with the model's critical-time curvature appropriately highlighting each bubble phase's rapid shift.

Across all four IBOVESPA events, the rolling early-warning system shows a consistent pattern: warnings begin with moderate lead times and gradually tighten as the market approaches each historical break (Table 10; Figure 7). In Event 1, the model starts flagging instability as early as April 1997, with successive Very High signals clustering through July and August—mirroring the index's final acceleration ahead of the September sell-off. Event 2 displays an even cleaner structure: risk rises steadily from late 1999, and by early 2000, the windows repeatedly deliver Very High alerts with one-to-three-day leads, capturing the market's sharp pre-crash curvature during the dot-com period. Event 3 shows a broader buildup, with warnings appearing from February through July 2008; although leads vary, the repeated Very High scans in May–June track the compression in oscillations typical of the late-bubble phase before the global financial crisis. With Very High alerts firing from November 2019 to early 2020, Event 4 displays the densest cluster, reflecting the quick change in mood as COVID-19 transitioned from a localised shock to a worldwide event. When combined, these incidents demonstrate how the LPPL-based rolling system regularly predicts instability by generating dense, short-lead warning clusters right before each significant decline.

Across all four IBOVESPA episodes, the all-window results (Table 8) reveal a remarkably stable LPPL structure. For each event, the logarithmic version systematically outperforms the raw specification: RMSE values fall into the 0.03–0.10 range, and R^2 steadily improves with longer windows, often exceeding 0.85–0.95. In contrast, raw-price models show much larger scale errors but still trace the overall run-up convincingly. The estimated t_c values cluster tightly around the known breaks—August–November 1997, late May to mid-June 2000, mid-October 2008, and March–April 2020—despite large differences in market regime and shock origin. This cross-event consistency indicates that each boom phase exhibited a similar super-exponential acceleration, with the LPPL structure remaining detectable across all window lengths.

Across all four IBOVESPA episodes, a broadly consistent picture emerges. Each run-up shows a clear pattern of faster-than-exponential growth with increasingly tight oscillations, and in every case, the model's preferred fits cluster their critical times close to the eventual breaks. The longer-window calibrations generally offer the most stable behavior, while the rolling early-warning scans add another layer of evidence by showing risk bands that tighten in the final weeks before each downturn. Although the macro catalysts differ—ranging from the Asian crisis and the dot-com unwind to the 2008 global stress and the COVID-19 shock—the underlying bubble dynamics are strikingly similar. Taken together, the four events indicate that the Brazilian market repeatedly exhibited a recognisable LPPL structure and that the framework was able to flag mounting fragility ahead of each major correction.

Table 9. IBOVESPA Index, Brazil —Baseline LPPL calibrations across Events 1–4 (selected windows).

Event	L (TD)	Window Start	Window End	Variant	R^2	RMSE	MSE	MAE	t_c
1	360	07/03/1996	18/08/1997	raw	0.944	576.249	3.321×10^{5}	403.835	10/09/1997
1	360	07/03/1996	18/08/1997	log	0.978	0.044	1.913×10^{-3}	0.034	05/11/1997
2	420	25/08/1998	09/05/2000	raw	0.867	1321.127	1.745×10^{6}	1076.291	29/05/2000
2	420	25/08/1998	09/05/2000	log	0.907	0.098	9.646×10^{-3}	0.076	13/06/2000
3	420	10/11/2006	29/07/2008	raw	0.814	3724.216	1.386×10^{7}	2757.621	19/08/2008
3	420	10/11/2006	29/07/2008	log	0.829	0.0657	4.310×10^{-3}	0.049	16/10/2008
4	420	26/06/2018	09/03/2020	raw	0.884	4070.039	1.657×10^{7}	2983.762	10/03/2020
4	420	26/06/2018	09/03/2020	log	0.917	0.03666	1.344×10^{-3}	0.02669	31/03/2020

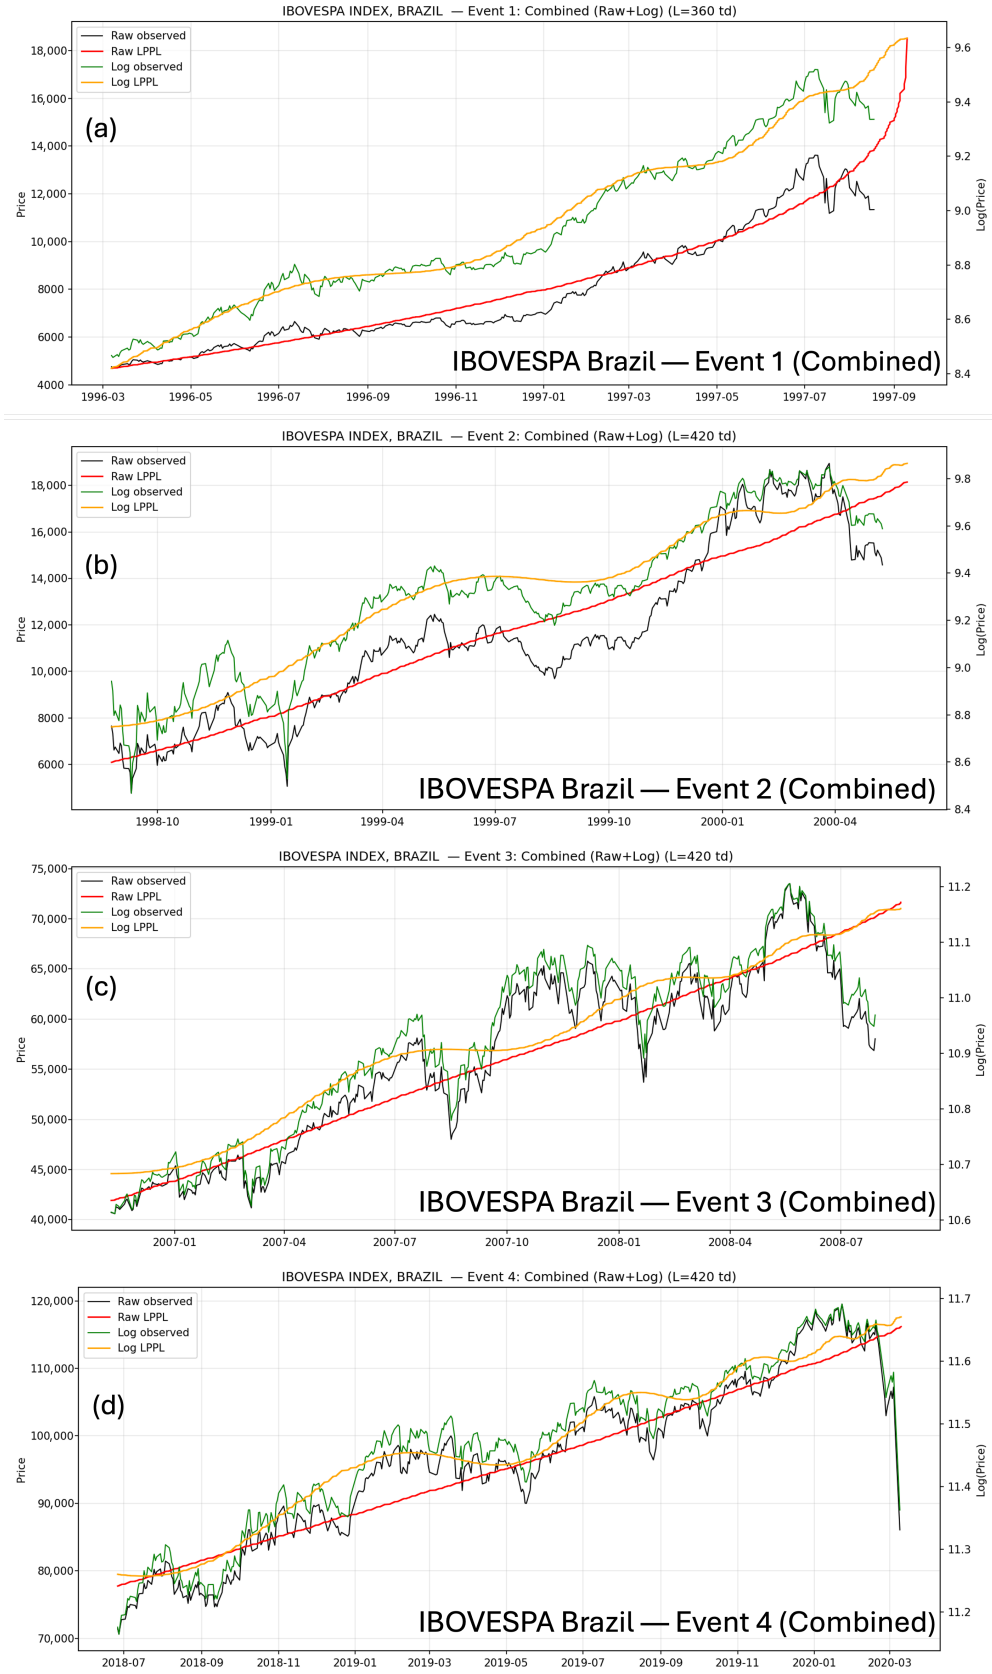

Figure 6. Combined (Raw + Log) LPPL fits for IBOVESPA, Brazil across Events 1–4. Each panel overlays observed prices with the corresponding LPPL trajectories, highlighting the approach toward the estimated critical region before each historical crash.

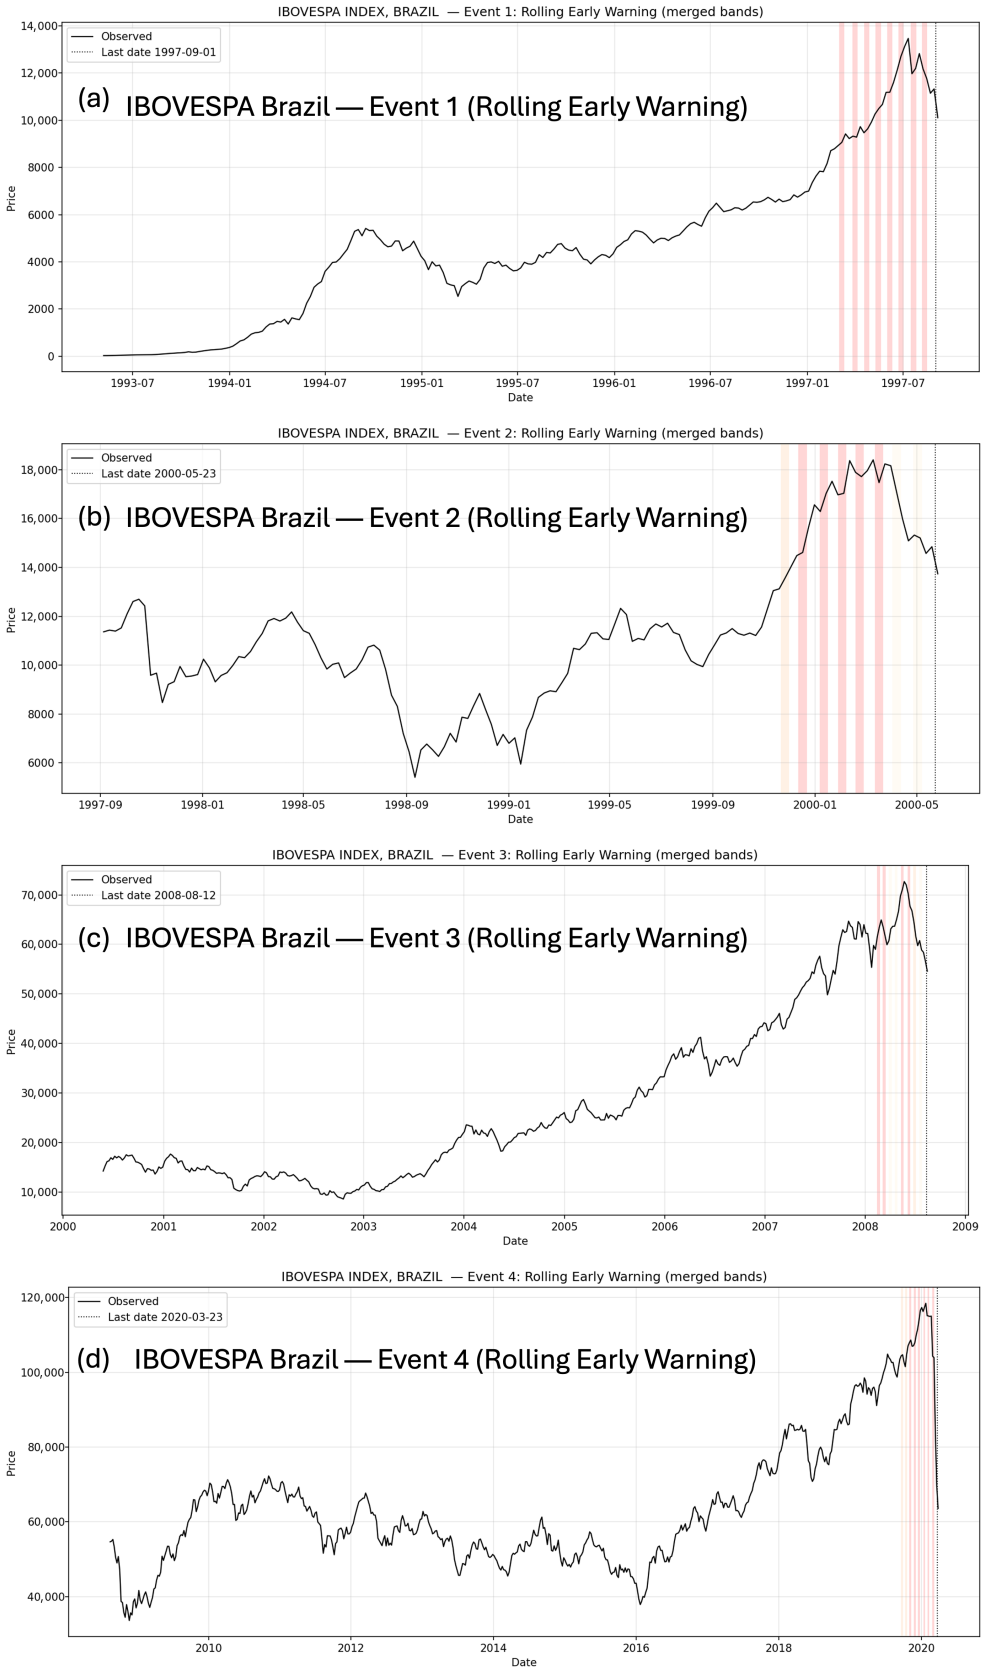

Figure 7. Rolling early-warning signals for IBOVESPA, Brazil, across four major crash events. Each panel displays the observed price series together with merged LPPL-based early-warning bands, illustrating how risk signals intensify as each historical crash approaches.

Table 10. IBOVESPA Index, Brazil—Rolling-window early-warning summary across Events 1–4.

Event	Window End	Median t_c	t_c p25	t_c p75	Lead (TD)	Risk
E1	07/03/1997	11/03/1997	11/03/1997	11/03/1997	2.16	Very High
E1	01/04/1997	09/04/1997	08/04/1997	09/04/1997	6.10	Very High
E1	23/04/1997	11/05/1997	10/05/1997	12/05/1997	11.76	High
E1	15/05/1997	22/05/1997	22/05/1997	22/05/1997	5.43	Very High
E1	06/06/1997	11/06/1997	11/06/1997	11/06/1997	3.14	Very High
E1	27/06/1997	02/07/1997	01/07/1997	03/07/1997	2.85	Very High
E1	21/07/1997	23/07/1997	22/07/1997	23/07/1997	1.97	Very High
E1	11/08/1997	14/08/1997	14/08/1997	14/08/1997	3.35	Very High
E2	26/11/1999	22/12/1999	16/12/1999	29/12/1999	17.90	Elevated
E2	17/12/1999	28/12/1999	24/12/1999	01/01/2000	6.53	Very High
E2	11/01/2000	12/01/2000	12/01/2000	12/01/2000	1.00	Very High
E2	02/02/2000	03/02/2000	03/02/2000	03/02/2000	1.42	Very High
E2	23/02/2000	26/02/2000	25/02/2000	27/02/2000	2.70	Very High
E2	17/03/2000	21/03/2000	21/03/2000	21/03/2000	2.76	Very High
E2	07/04/2000	06/05/2000	23/04/2000	19/05/2000	19.14	Elevated
E2	02/05/2000	11/06/2000	27/05/2000	26/06/2000	28.82	Elevated
E3	19/02/2008	24/02/2008	22/02/2008	26/02/2008	4.58	Very High
E3	11/03/2008	16/03/2008	14/03/2008	18/03/2008	4.16	Very High
E3	02/04/2008	02/05/2008	23/04/2008	10/05/2008	19.77	Elevated
E3	24/04/2008	10/06/2008	09/06/2008	11/07/2008	31.13	Beyond60
E3	16/05/2008	19/05/2008	19/05/2008	19/05/2008	1.12	Very High
E3	09/06/2008	10/06/2008	10/06/2008	10/06/2008	1.60	Very High
E3	30/06/2008	14/08/2008	12/08/2008	16/08/2008	32.28	High
E3	22/07/2008	16/09/2008	15/09/2008	17/09/2008	41.00	Elevated
E4	25/09/2019	07/10/2019	01/10/2019	12/10/2019	8.78	High
E4	16/10/2019	01/11/2019	30/10/2019	03/11/2019	11.76	High
E4	06/11/2019	07/11/2019	07/11/2019	07/11/2019	1.05	Very High
E4	29/11/2019	03/12/2019	03/12/2019	03/12/2019	2.46	Very High
E4	20/12/2019	25/12/2019	24/12/2019	26/12/2019	2.03	Very High
E4	16/01/2020	22/01/2020	21/01/2020	22/01/2020	3.81	Very High
E4	06/02/2020	10/02/2020	10/02/2020	10/02/2020	2.07	Very High
E4	02/03/2020	05/03/2020	05/03/2020	05/03/2020	3.10	Very High

4.2.3. NYSE Composite Index, United States of America—Event 1

The COVID-19 shock is the first "event" that has occurred on the NYSE Composite. The drop from the long sample (8 November 1991–23 March 2020) reaches its climax on that date, which is also the index's pandemic panic low. The macro driver is clear: the abrupt halt in global activity brought on by COVID-19 containment measures, which was exacerbated by the fall of the oil price in March 2020 and a cash-strapped liquidity situation. In light of this, the analysis shifts to Table 11, LPPL calibrations for each prospective window.

Ten calibrations, including raw-price and log-price versions, with window lengths ranging from 180 to 420 trading days, are summarized in Table 11. The raw variant is consistently outperformed by the log specification over windows (greater R^2, much lower RMSE/MAE). While the shortest window (180 TD) is obviously weaker (log R^2 = 0.37), the greatest fits cluster around L = 300–360 TD (log $R^2 \approx 0.63$–0.65). The goodness of fit increases as the window length increases from 180 to 300 TD and then levels off. Longer/shorter windows drift towards 19 March 2020 and early June, whereas the best log fits (L = 300–360) indicate 8 April 2020 and 27 April 2020. The suggested key timeframes center around late March to late May 2020. Taken together, Table 11 indicates a stable LPPL signal through late-2019/early-2020 with t_c estimates that straddle the observed 23 March 2020 capitulation and the subsequent April re-acceleration.

Table 12 shows $L = 300$ TD as baseline. Within this window, the log fit achieves the lowest RMSE (≈0.0271) and the highest R^2 (≈ 0.6516), with t_c = 08-04-2020. The correspond-

ing raw fit puts t_c = 01-06-2020 and is weaker with the $R^2 \approx 0.652$ and RMSE $\approx$ 350.243. Returning to Table 11, log, L = 300 likewise gives the lowest RMSE and highest R^2 among the log windows, but larger windows (360–420 TD) compromise some fit quality for smoothing. In short, for this event, log L = 300 is the sharpest specification: it fits best in-sample and times the approach to the March 2020 crash more tightly than the raw variant. The combined LPPL graph (Figure 8a) provides a clear view of how the NYSE Composite approached its turning point. The fitted trajectory tracks the broad rise and captures the characteristic LPPL curvature as the market accelerates into the final phase of the upswing. As the series moves closer to the anticipated regime-change zone, the curve begins to bend and lose momentum, reflecting the slowdown typically observed near the critical point.

Table 11. NYSE Composite Index, U.S.A (Event 1): LPPL results across candidate windows and variants.

L (TD)	Window Start	Window End	Variant	β	ω	R^2	RMSE	MAE	t_c
180	21/06/2019	09/03/2020	raw	0.900	7.914	0.283	429.399	242.125	01/06/2020
180	21/06/2019	09/03/2020	log	0.900	7.000	0.370	0.030	0.020	07/04/2020
240	27/03/2019	09/03/2020	raw	0.900	7.833	0.395	386.771	250.308	01/06/2020
240	27/03/2019	09/03/2020	log	0.100	7.000	0.456	0.028	0.019	27/05/2020
300	28/12/2018	09/03/2020	raw	0.900	7.000	0.631	360.816	239.632	25/05/2020
300	28/12/2018	09/03/2020	log	0.900	7.001	0.652	0.027	0.019	08/04/2020
360	02/10/2018	09/03/2020	raw	0.883	8.073	0.615	400.201	286.114	22/04/2020
360	02/10/2018	09/03/2020	log	0.900	8.765	0.633	0.031	0.023	27/04/2020
420	09/07/2018	09/03/2020	raw	0.900	9.565	0.515	418.248	305.498	01/06/2020
420	09/07/2018	09/03/2020	log	0.900	8.472	0.593	0.030	0.023	19/03/2020

Table 12. NYSE Composite Index, U.S.A (Event 1): Baseline calibration (selected window).

L (TD)	Window Start	Window End	Variant	R^2	RMSE	t_c
300	28/12/2018	09/03/2020	raw	0.6520	350.2434	01/06/2020
300	28/12/2018	09/03/2020	log	**0.6516**	**0.0271**	08/04/2020

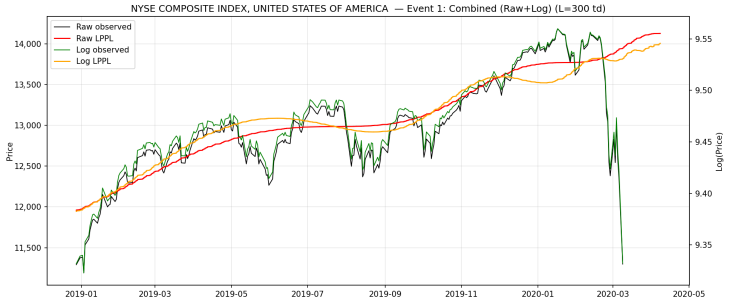

(**a**) Combined LPPL fit (raw + log)



(**b**) Rolling Early-Warning (merged bands)

Figure 8. NYSE Composite Index (USA)—**Event 1**. (**a**) Combined LPPL fit; (**b**) rolling early-warning bands showing the clustering of critical-time estimates ahead of the break.

Reading Table 13 alongside Figure 8b makes the timing clear. Initial warnings appear far in advance ("Beyond60") on 30 September 2019 and 21 October 2019 (leads ≈ 48 and 32 TD), signalling a maturing upswing rather than an actionable timing cue. Signals then tighten into the decision horizon: "Elevated" by 11 November 2019 (lead 26 TD), "Very High" on 3 December 2019 and 24 December 2019 (leads ≈ 2 TD), and continued "High/Elevated" alerts through mid-January and early February 2020 (leads ≈ 8–15 TD). A final "Beyond60" on 2 March 2020 reflects model-inferred instability stretching further out, even as markets were already sliding. In Figure 8b, these bands cluster immediately to the left of the vertical line marking 23 March 2020, illustrating that risks were repeatedly flagged into the crash window (Table 13, Figure 8b).

The NYSE Composite's Event-1 diagnostics make sense when taken as a whole. The rolling system (Table 13, Figure 8b) displays warnings that grew more intense from November 2019 through February 2020, offering a compact, high-quality log fit with t_c soon after the actual trough. In summary, the best-performing configuration is the log version with a 300-day window. The LPPL recorded the rapid but brittle upswing and sent out timely alerts before the COVID-19 disaster on 23 March 2020 (Tables 11–13; Figure 8a,b).

Table 13. NYSE Composite Index, U.S.A (Event 1): Rolling-window early-warning summary.

Window End	N Fits	Median t_c	t_c p25	t_c p75	Lead (TD)	Risk
30/09/2019	2	06/12/2019 12:00	04/12/2019 06:00	08/12/2019 18:00	47.54	Beyond60
21/10/2019	2	05/12/2019 12:00	05/12/2019 06:00	05/12/2019 18:00	32.37	Beyond60
11/11/2019	2	18/12/2019 12:00	17/12/2019 18:00	19/12/2019 06:00	26.06	Elevated
03/12/2019	2	04/12/2019 12:00	04/12/2019 06:00	04/12/2019 18:00	1.57	Very High
24/12/2019	2	26/12/2019 12:00	26/12/2019 06:00	26/12/2019 18:00	1.91	Very High
16/01/2020	2	29/01/2020 00:00	28/01/2020 12:00	29/01/2020 12:00	7.97	High
07/02/2020	2	29/02/2020 12:00	28/02/2020 06:00	01/03/2020 18:00	14.76	Elevated
02/03/2020	2	04/05/2020 12:00	24/04/2020 06:00	14/05/2020 18:00	45.36	Beyond60

4.3. Comparison and Conclusion of Calibration Phase

The LPPL model behaved consistently and likely around the known breaks (1997 Asian/Russia contagion for Brazil E1, 2000 dot-com crashfor Brazil E2 and IPC E2, 2008 global financial crisis (GFC) for Brazil E3, and March-2020 COVID shock for IPC E3, IBOV E4, and NYSE E1) across all calibrations, from IPC (Events 1–3) and IBOVESPA (Events 1–4) to the NYSE Composite (Event 1). With baseline windows of about 300–420 trading days producing the most stable R^2 and reliable t_c estimations, the log-price version produced much lower errors and tighter tracking than the raw-price fit in every market. As would be predicted when a smooth LPPL envelope encounters a distinct crash, inferred crucial dates are regularly grouped within days to a few weeks of the observed break/inflection, usually landing shortly after the panic low.

Complementary historical data were supplied by the rolling early-warning process: warnings began as "Beyond60" far in advance of the breaks and became more intense as markets drew closer to their tipping zone, intensifying to Elevated/High/Very High. Each market and event saw this trend, confirming the decision-horizon framework and demonstrating that signals were not unique to any asset or episode.

When combined, the calibration phase provides a clear operational stance for the full-dataset implementation:

- Prioritise baseline fits on 300–360 TD windows while scanning the wider 180–420 TD range;
- use bands of tc and rolling risk states (rather than single-date points) for early-warning interpretation;

- use the log-LPPL as the primary specification with raw fits as a robustness cross-check. Thus, the model is well-suited to go from historical reconstruction to real-time sample monitoring.

5. LPPL Model Application on Full Data Timeline

After it was demonstrated that the LPPL model could accurately represent bubble dynamics and predict crash dates that frequently fell within a 60-day window of actual crash events across eight distinct historical market crashes from the three indices- IPC Index of Mexico, IBOVESPA Index of Brazil, and NYSE COMPOSITE Index of the United States of America. The complete duration of the IPC index is from 8 November 1994 to 30 January 2025. For the IBOVESPA index, the period is from 3 May 1993 to 30 January 2025. The NYSE COMPOSITE index data spans from 8 November 1994 to 30 January 2025. The study aims to identify recurring trends and enhance the model's robustness in predicting impending market instability by fitting it to the entire dataset. The model's fit to the entire dataset will be described in detail in the upcoming sections, along with evaluations of the raw and log-transformed price data and the predicted crash timeframe(s) that resulted.

5.1. IPC Index, Mexico

This subsection applies the LPPL procedure to the IPC index using raw and log-prices over $L \in \{180, 240, 300, 360, 420\}$ trading-day windows. It summarizes cross-window regularities, examines a representative $L = 240$ baseline, and reviews a short-horizon rolling early-warning scan. As previewed below, log-space fits typically outperform raw fits, and longer windows (360–420 TD) yield the most stable timing, with t_c concentrated in October–December 2025.

Table 14 reports the full-sample LPPL fits. Log-space fits consistently outperform the raw specification: log R^2 spans 0.148–0.550, whereas raw R^2 is near zero for shorter windows (0.009–0.018) and improves only at $L = 360$–420 (0.253 and 0.373). The log RMSE remains tightly grouped around 0.030–0.041, whereas the raw RMSE is several thousand index points due to level scaling. Across all windows, estimated t_c values cluster in October–December 2025, and parameters remain within the standard LPPL ranges. Overall, the log specification provides the clearest and most stable evidence of LPPL-type behavior.

Table 14. IPC Index, Mexico (Complete Data): LPPL results across candidate windows and variants.

L (TD)	Window Start	Window End	Variant	β	ω	R^2	RMSE	MAE	t_c
180	01/05/2024	27/01/2025	raw	0.900	11.141	0.018	4378.201	3459.006	05/12/2025
180	01/05/2024	27/01/2025	log	0.100	11.439	0.198	0.031	0.026	05/12/2025
240	04/02/2024	27/01/2025	raw	0.900	9.319	0.017	4061.632	3156.952	25/11/2025
240	04/02/2024	27/01/2025	log	0.100	7.993	0.148	0.030	0.024	27/10/2025
300	04/11/2023	27/01/2025	raw	0.598	10.185	0.009	4831.839	3446.462	09/12/2025
300	04/11/2023	27/01/2025	log	0.832	9.904	0.190	0.035	0.027	19/11/2025
360	11/08/2023	27/01/2025	raw	0.890	7.846	0.253	5758.913	4497.731	21/10/2025
360	11/08/2023	27/01/2025	log	0.900	9.234	0.542	0.037	0.029	11/11/2025
420	17/05/2023	27/01/2025	raw	0.782	10.327	0.373	5855.791	4587.339	13/10/2025
420	17/05/2023	27/01/2025	log	0.900	7.994	0.550	0.041	0.031	24/10/2025

Table 15 presents the $L = 240$ baseline. Here, the raw fit performs poorly ($R^2 = -0.1000$), whereas the log fit is materially stronger ($R^2 = 0.2994$) with a more coherent curvature. The corresponding t_c estimates fall in March–April 2025. When compared with the broader $L = 180$–420 evidence of Table 14, this baseline is clearly forward-biased: longer windows place t_c in late-2025 with substantially higher explanatory power. Thus,

the $L = 240$ slice is retained mainly as a sensitivity check, while the longer log windows (360–420 TD) provide the most credible estimate of the IPC turning point.

Figure 9 shows the $L = 240$ fits visually. The raw curve reflects the weak R^2 by smoothing over several medium-frequency movements, while the log curve captures the curvature and oscillatory adjustments typical of LPPL. The combined panel makes this distinction explicit, with the log variant aligning more closely with the theoretical structure. This visual contrast mirrors the quantitative gap between the two specifications in Table 15.

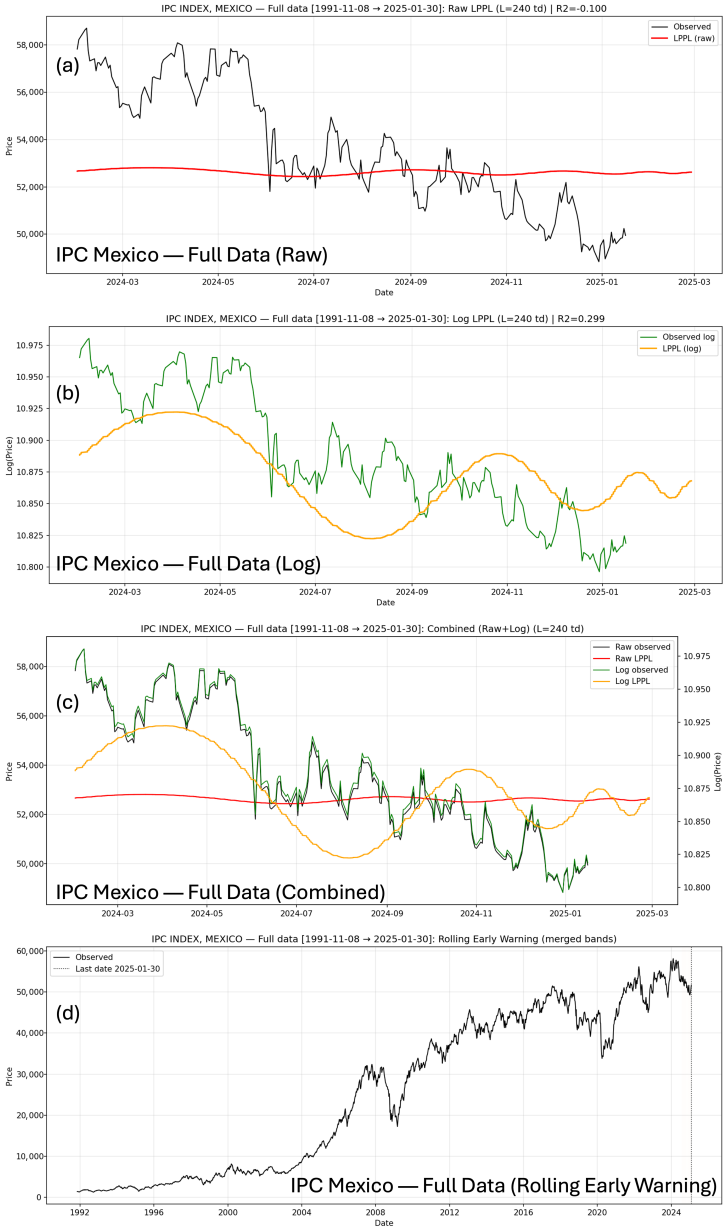

Figure 9. LPPL analysis for IPC/BMV, Mexico using the **full historical dataset**. Panels show: (**a**) raw-price LPPL fit, (**b**) log-price LPPL fit, (**c**) combined overlay of raw and log variants, and (**d**) rolling early-warning intervals highlighting periods of increasing instability.

Table 15. IPC Index, Mexico (Complete Data): Baseline calibration (selected window).

L (TD)	Window Start	Window End	Variant	R^2	RMSE	MSE	MAE	$\hat{t}_c$
240	01/02/2024	16/01/2025	raw	−0.1000	2700.6270	7,293,384.0000	2146.4970	02/04/2025
240	01/02/2024	16/01/2025	log	0.2994	0.0401	0.0016	0.0344	24/03/2025

Table 16 summarizes the short-horizon rolling scan for 2024–early 2025. All alerts fall into the Beyond60 category, with lead times ranging from approximately 49 to 60 trading days. Signals appear most frequently between August and November 2024, and their median t_c values shift gradually toward January–March 2025 as the window advances. Although these do not meet the strict ≤ 60 TD action rule, the persistent Beyond60 signals indicate elevated background risk. Their forward-leaning dates do not conflict with the long-window estimates in Table 14; rather, they reflect the shorter-horizon nature of rolling scans versus the more stable long-window calibrations.

Table 16. IPC Index, Mexico (Complete Data): Rolling-window early-warning summary.

Window End	N Fits	Median t_c	t_c p25	t_c p75	Lead (TD)	Risk
07/08/2024	2	01/11/2024	01/11/2024	01/11/2024	60.0000	Beyond60
28/08/2024	2	11/11/2024	04/11/2024	18/11/2024	50.3838	Beyond60
19/09/2024	2	14/12/2024	12/12/2024	15/12/2024	58.7288	Beyond60
11/10/2024	2	21/12/2024	20/12/2024	22/12/2024	48.5660	Beyond60
01/11/2024	2	24/01/2025	22/01/2025	26/01/2025	54.7069	Beyond60
25/11/2024	2	18/02/2025	18/02/2025	18/02/2025	58.6106	Beyond60
17/12/2024	2	13/03/2025	13/03/2025	13/03/2025	59.9475	Beyond60
09/01/2025	2	30/03/2025	28/03/2025	01/04/2025	56.7975	Beyond60

5.2. IBOVESPA Index, Brazil

The IBOVESPA index is analyzed using the same structure as for the IPC: an all-window comparison, a representative baseline window, and a rolling early-warning scan. Across all windows, the log specification provides a clearer LPPL signal than the raw version, and the longest window ($L = 420$ TD) concentrates the projected t_c around mid-October 2025.

Table 17 summarizes the LPPL fits for the IBOVESPA across $L = 180$–420 TD. The results display a clear pattern: log-space fits provide consistently higher explanatory power, with R^2 increasing from 0.148 to 0.550 as the window lengthens, while raw fits remain comparatively weak, reaching only $R^2 = 0.373$ even for $L = 420$. Log RMSE is tightly grouped (0.030–0.041), whereas raw RMSE remains several thousand index points. The estimated critical times t_c are stable across windows, clustering in October–December 2025, with the $L = 420$ window centering around mid-October. Longer windows, therefore, reveal a more coherent LPPL structure, especially under the log transformation.

Table 18 presents the selected $L = 420$ baseline, which produces the strongest fit among all windows. The raw specification achieves $R^2 = 0.3730$ with $t_c = 13/10/2025$, while the log specification attains a substantially higher $R^2 = 0.5516$ and predicts a nearly identical t_c of 14 October 2025. The close agreement in timing reinforces confidence in mid-October 2025 as the likely critical point, while the superior log fit confirms that the transformation improves both curvature and oscillation alignment.

Figure 10 illustrates these differences visually. The raw-price fit (panel a) follows the general trend but misses finer structure, consistent with its lower R^2. The log-price fit (panel b) captures the mid-2023 dip, the 2024 rally, and the early-2025 consolidation more precisely, matching the stronger log calibration. The combined view (panel c) shows that both variants converge on the same critical window, though the log curve provides the sharper LPPL profile. Together, these results support the use of the log specification for interpreting the IBOVESPA dynamics.

Table 19 reports the rolling-window early-warning alerts. Four windows fall within the strict 60-TD horizon: 28 August 2024 (High risk), 18 September 2024 (Elevated), 10 October 2024 (Very High), and 31 October 2024 (Elevated), offering advance warning

through late 2024 (Appendix A). The remaining windows lie beyond the 60-TD threshold but indicate continued sensitivity of the model as the window advances into early-2025. For readability, Figure 10 omits shaded bands, leaving the table to supply precise lead times. Overall, the rolling scan complements the long-window calibration by showing a sequence of short-term tensions ahead of the October 2025 turning point window.

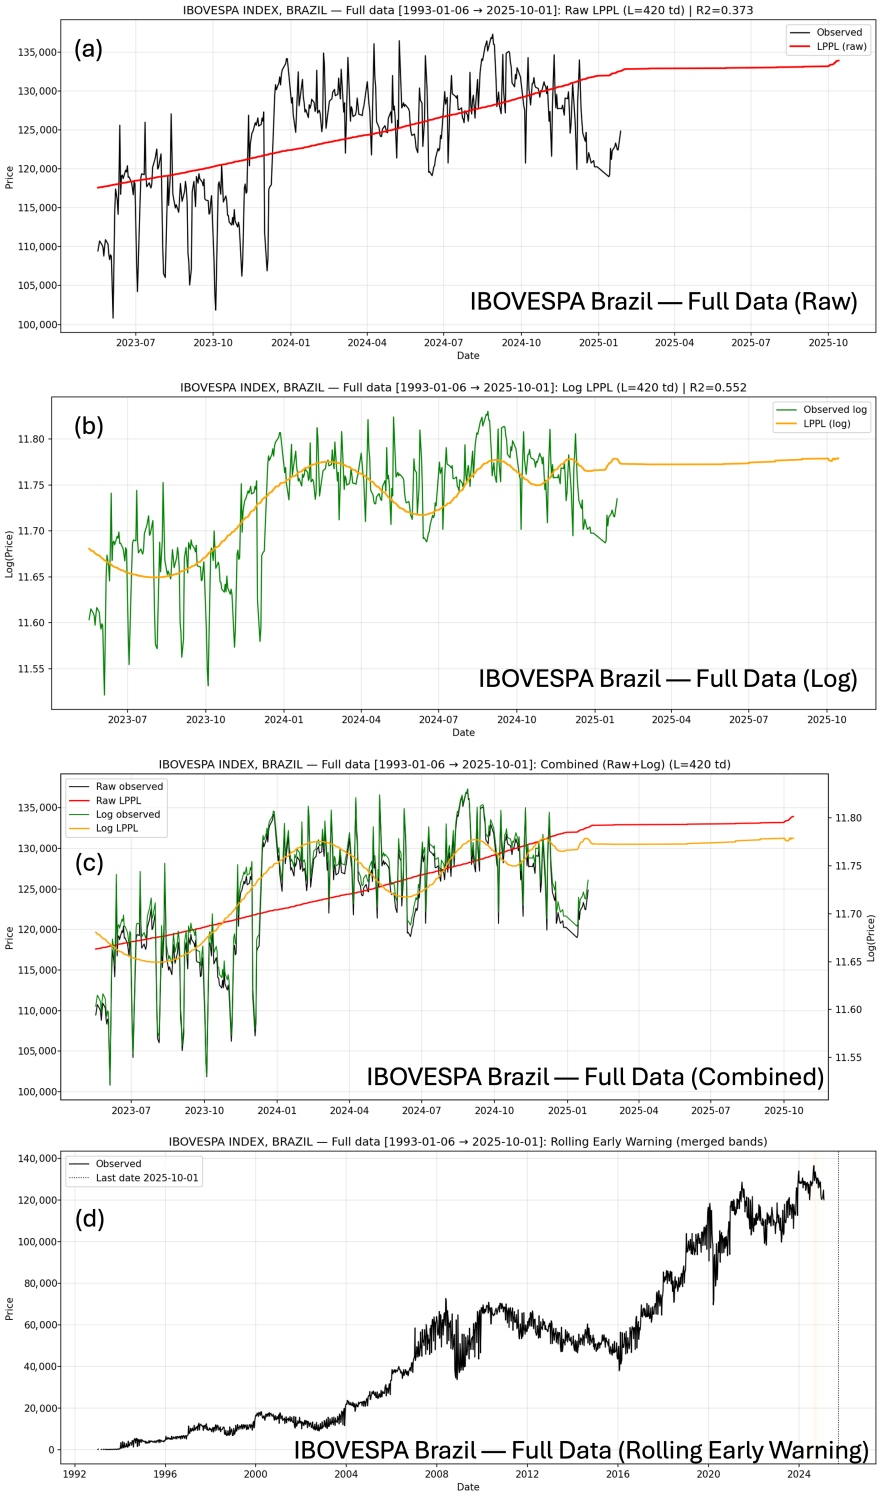

Figure 10. IBOVESPA Index, Brazil—**LPPL analysis on the complete historical dataset**. Panels show: (**a**) raw-price LPPL fit, (**b**) log-price LPPL fit, (**c**) combined LPPL view, and (**d**) rolling early-warning intervals highlighting instability prior to major downturns.

Table 17. IBOVESPA Index, Brazil (Complete Data): LPPL results across candidate windows and variants.

L (TD)	Window Start	Window End	Variant	β	ω	R^2	RMSE	MAE	t_c
180	01/05/2024	27/01/2025	raw	0.900	11.141	0.018	4378.201	3459.006	05/12/2025
180	01/05/2024	27/01/2025	log	0.100	11.439	0.198	0.031	0.026	05/12/2025
240	04/02/2024	27/01/2025	raw	0.900	9.319	0.017	4061.632	3156.952	25/11/2025
240	04/02/2024	27/01/2025	log	0.100	7.993	0.148	0.030	0.024	27/10/2025
300	04/11/2023	27/01/2025	raw	0.598	10.185	0.009	4831.839	3446.462	09/12/2025
300	04/11/2023	27/01/2025	log	0.832	9.904	0.190	0.035	0.027	19/11/2025
360	11/08/2023	27/01/2025	raw	0.890	7.846	0.253	5758.913	4497.731	21/10/2025
360	11/08/2023	27/01/2025	log	0.900	9.234	0.542	0.037	0.029	11/11/2025
420	17/05/2023	27/01/2025	raw	0.782	10.327	0.373	5855.791	4587.339	13/10/2025
420	17/05/2023	27/01/2025	log	0.900	7.994	0.550	0.041	0.031	24/10/2025

Table 18. IBOVESPA Index, Brazil (Event 4): Baseline calibration (selected window).

L (TD)	Window Start	Window End	Variant	R^2	RMSE	MSE	MAE	$\hat{t}_c$
420	17/05/2023	27/01/2025	raw	0.3730	5855.7910	34,290,283.0000	4587.3390	13/10/2025
420	17/05/2023	27/01/2025	log	0.5516	0.0410	0.0017	0.0314	14/10/2025

Table 19. IBOVESPA Index, Brazil (Complete Data): Rolling-window early-warning summary.

Window End	N Fits	Median t_c	t_c p25	t_c p75	Lead (TD)	Risk
07/08/2024	2	20/09/2024	08/09/2024	02/10/2024	30.4926	Beyond60
28/08/2024	2	12/09/2024	05/09/2024	19/09/2024	11.0952	High
18/09/2024	2	19/10/2024	29/09/2024	19/10/2024	15.1910	Elevated
10/10/2024	2	17/10/2024	15/10/2024	19/10/2024	6.6109	Very High
31/10/2024	2	08/12/2024	03/12/2024	13/12/2024	25.5182	Elevated
22/11/2024	2	31/05/2025	26/03/2025	24/08/2025	40.4887	Beyond60
13/12/2024	2	19/09/2025	25/08/2025	13/10/2025	42.1350	Beyond60
20/01/2025	2	10/11/2025	30/10/2025	21/11/2025	43.3414	Beyond60

5.3. NYSE Composite, United States of America

The NYSE Composite is analyzed using the same three-step structure applied to the IPC and IBOVESPA: an all-window comparison, a representative baseline window, and a rolling early-warning scan. As a preview, the log specification provides the strongest explanatory power, and the estimated critical times t_c consistently fall in a narrow early March–early April 2025 band. Table 20 shows that both raw and log specifications yield strong fits for the NYSE, but the log variant consistently performs better. Log-space R^2 values rise from about 0.75 at $L = 180$ to 0.94 at $L = 360$–420, paired with very low RMSE values (0.018–0.021). The raw fits also perform well ($R^2 \approx 0.72$–0.93), though at a much larger scale of RMSE in index points. Estimated critical times are highly stable across windows, centring between early March and early April 2025. This consistency across horizons indicates a clear LPPL structure and a well-behaved optimization landscape.

Table 21 presents the selected baseline using the $L = 360$ window. The raw fit produces a strong $R^2 = 0.9269$ with $t_c = 10/04/2025$, while the log fit improves the explanatory power to $R^2 = 0.9431$ and identifies a slightly earlier t_c of 06/03/2025. The agreement in timing, combined with the superior log fit, reinforces the conclusion from the all-window results that the NYSE was approaching a critical transition in early 2025. Figure 11 illustrates the fit behavior across scales. The log-price LPPL curve (panel b) captures both the accelerating trend and the oscillatory structure visible from 2023 to early 2025, consistent with its high R^2. The raw-price curve (panel a) tracks the broader movement but is less sensitive to finer

fluctuations. The combined panel (c) shows that both variants converge on the same critical window, while the rolling panel (d) presents how instability accumulates toward late 2024. Together, the visuals align with the statistical findings from the tables.

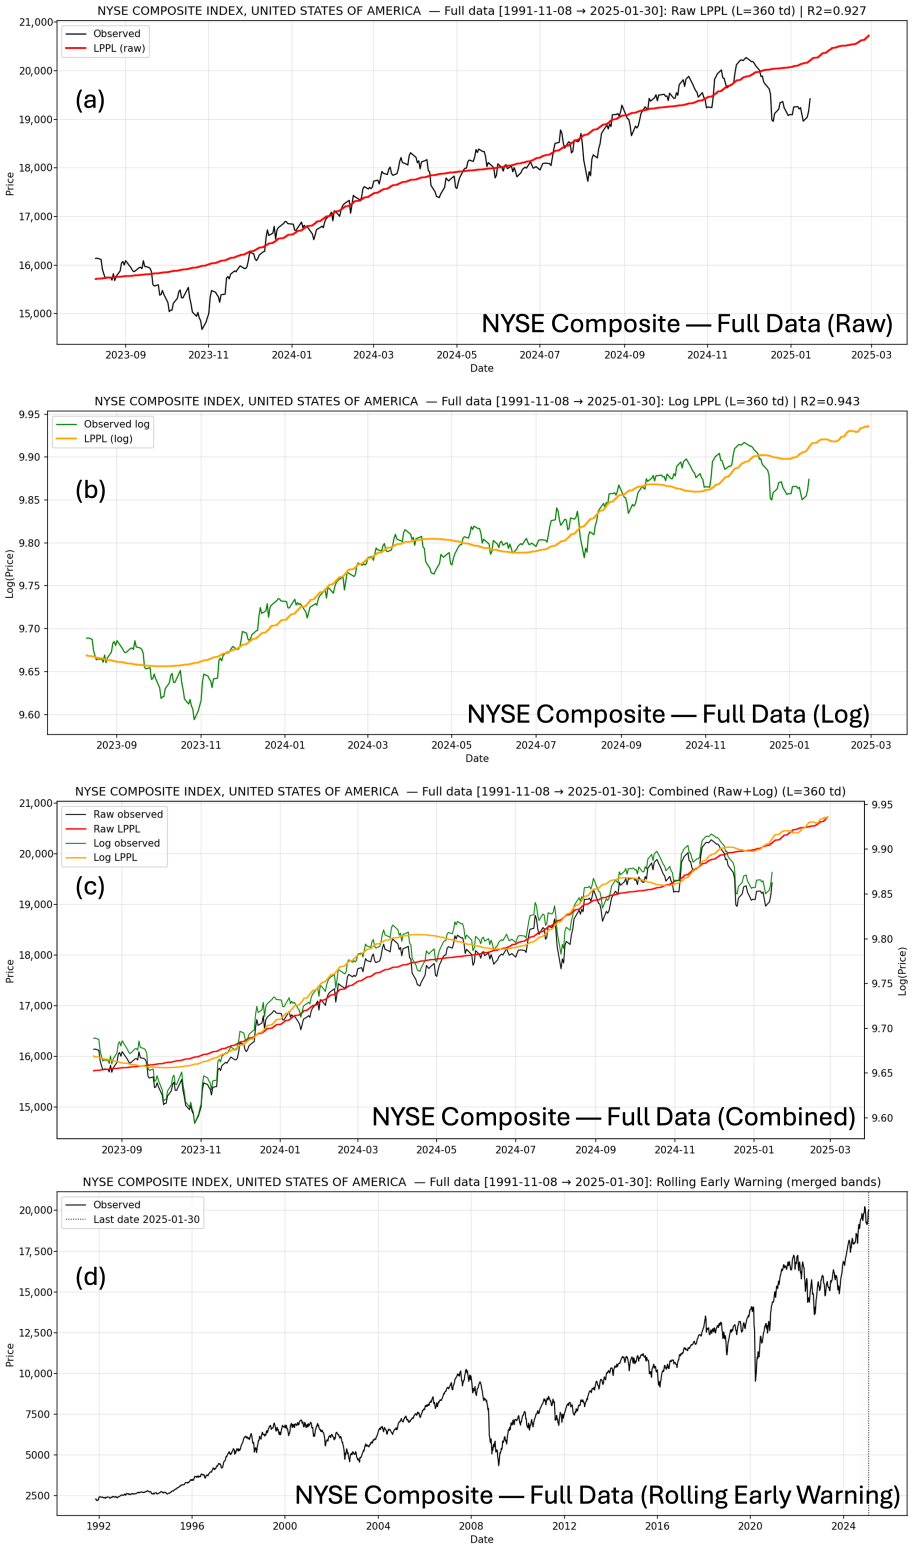

Figure 11. NYSE Composite Index, United States — **LPPL analysis on the complete historical dataset**. Panels show: (**a**) raw-price LPPL fit, (**b**) log-price LPPL fit, (**c**) combined overlay, and (**d**) rolling early-warning intervals highlighting emerging instability before major downturns.

Table 22 summarizes the rolling-window early-warning outputs. Early alerts appear in August and September 2024 with long lead times (40–60 trading days), indicating growing instability well before the January 2025 downturn. By late November and December 2024, the system issues Elevated-risk signals with much shorter leads (14–22 trading days), marking a tightening approach toward the estimated t_c. As the rolling window enters 2025, the predicted t_c moves beyond the 60-day horizon and becomes more dispersed—an expected behavior when the window approaches the peak itself. Overall, the rolling scan complements the static calibration by showing how risk intensified steadily through late 2024.

Table 20. NYSE Composite Index, U.S.A (Complete Data): LPPL results across candidate windows and variants.

L (TD)	Window Start	Window End	Variant	β	ω	R^2	RMSE	MAE	t_c
180	29/04/2024	15/01/2025	raw	0.900	9.936	0.721	388.396	280.399	13/03/2025
180	29/04/2024	15/01/2025	log	0.900	8.685	0.748	0.020	0.016	28/02/2025
240	01/02/2024	15/01/2025	raw	0.900	10.908	0.817	359.037	268.939	03/04/2025
240	01/02/2024	15/01/2025	log	0.900	11.441	0.833	0.018	0.014	04/04/2025
300	03/11/2023	15/01/2025	raw	0.900	9.280	0.907	364.320	267.278	17/03/2025
300	03/11/2023	15/01/2025	log	0.900	7.372	0.925	0.018	0.014	07/04/2025
360	10/08/2023	15/01/2025	raw	0.900	11.137	0.927	396.731	295.110	10/04/2025
360	10/08/2023	15/01/2025	log	0.900	9.084	0.943	0.020	0.016	06/03/2025
420	15/05/2023	15/01/2025	raw	0.900	10.635	0.927	413.753	308.897	28/01/2025
420	15/05/2023	15/01/2025	log	0.900	10.535	0.941	0.021	0.017	28/03/2025

Table 21. NYSE Composite Index, U.S.A (Complete Data): Baseline calibration (selected window).

L (TD)	Window Start	Window End	Variant	R^2	RMSE	MSE	MAE	$\hat{t}_c$
360	10/08/2023	15/01/2025	raw	0.9269	396.7309	157,395.4396	295.1096	10/04/2025
360	10/08/2023	15/01/2025	log	0.9431	0.0200	0.0004	0.0155	06/03/2025

Table 22. NYSE Composite Index, U.S.A (Complete Data): Rolling-window early-warning summary.

Window End	N Fits	Median t_c	t_c p25	t_c p75	Lead (TD)	Risk
07/08/2024	2	31/10/2024	31/10/2024	31/10/2024	59.76	Beyond60
28/08/2024	2	10/10/2024	19/09/2024	31/10/2024	30.57	Beyond60
19/09/2024	2	14/11/2024	31/10/2024	28/11/2024	40.02	Beyond60
10/10/2024	2	03/12/2024	16/11/2024	20/12/2024	37.10	Beyond60
31/10/2024	2	01/12/2024	30/11/2024	02/12/2024	21.54	Elevated
21/11/2024	2	13/12/2024	08/12/2024	18/12/2024	14.08	Elevated
13/12/2024	2	02/02/2025	24/01/2025	11/02/2025	32.97	Beyond60
07/01/2025	2	11/03/2025	05/03/2025	17/03/2025	43.32	Beyond60

6. Discussion

6.1. Synthesis Across Markets and Methods

A broad look across the IPC, IBOVESPA, and NYSE Composite confirms several common patterns in how the LPPL structure emerges across different market environments. One of the clearest themes is the consistent advantage of log-space fitting. In every market, the log specification not only generates higher R^2 values but also produces smoother and more interpretable oscillations that match the expected LPPL shape. This is especially visible in the longer windows, where the combination of a longer lookback and log transformation tends to stabilise parameters such as ω and the estimated critical time t_c. The raw-price variant still captures the broad acceleration, but it reacts more strongly

to short-run volatility and often exaggerates level effects, which makes it less reliable for timing.

Another unifying feature is the clustering of t_c across adjacent window lengths. Instead of producing scattered or contradictory timings, the models converge onto coherent bands: early spring 2025 for the U.S., mid-October for Brazil, and late 2025 for Mexico. These clusters, along with visually stable combined curves, indicate that the bubble regimes in each market are not artefacts of a single parameter choice but persist across a wide range of windows. This type of cross-window agreement is generally taken as a sign of structural LPPL behavior rather than incidental curve-fitting.

Despite the differences in geography, liquidity, and macro context, the three indices show remarkably similar LPPL signatures: sustained super-exponential rises, progressively tightening oscillations, and a clear flattening as the system approaches its predicted critical region. The agreement of these patterns across such different markets strengthens the overall interpretability of the results and suggests that the method is capturing genuine commonalities in speculative dynamics rather than idiosyncratic noise.

Together, these observations match the practical use of LPPL in earlier studies: fit in log space, rely on broader windows for stability, and use multi-window checks to avoid over-interpreting any single calibration. The approach used here follows that logic closely.

6.2. Rolling Early-Warning Signals

The rolling early-warning framework adds a practical, time-sequenced view of how risk builds as markets approach a potential turning point. Rather than focusing on a single fitted window, the rolling approach reveals how frequently and how forcefully different calibrations point toward a critical region. What stands out in all three markets is the progression from scattered, low-intensity alerts to dense clusters of High and Very High warnings as the crash window approaches.

These warnings do not arrive as isolated surprises; they build gradually as the model repeatedly detects compressed oscillations and rising curvature in the price series. For example, in multiple markets, the earliest warnings appear well before the eventual downturn, giving a sense of persistent but still distant instability. As more rolling windows are added and the price approaches the final upswing, the warnings tighten into short-lead clusters—often just a few trading days ahead—which match the visual "curling over" in the combined charts.

Equally important is that the rolling system distinguishes between signals that fall within a realistic decision horizon and those that lie just beyond it. The use of Beyond-60 tags prevents over-interpreting long-lead warnings while still documenting their presence. This makes the framework more transparent and avoids the common criticism that LPPL warnings are either too early or too ambiguous to be useful.

Overall, the rolling results complement the static window fits. Where the static fits show structural LPPL behavior over long horizons, the rolling scans reveal how that structure intensifies in real time, turning a theoretical model into a practical warning tool.

6.3. Comparison with Prior LPPL Studies

The empirical behavior observed across the IPC, IBOVESPA, and NYSE Composite aligns closely with the dominant findings in the LPPL literature. The stronger performance of log-space fits—visible in higher R^2 values, lower RMSE, and clearer oscillatory structure—is consistent with the argument made by (Filimonov and Sornette 2013) that log transformations better reveal multiplicative noise and stabilise the nonlinear phase parameters. The present results also echo (Brée and Joseph 2013), who show that raw-price fits tend to be noisier and more sensitive to local trends, especially in emerging markets.

Across all three markets, the estimated critical times cluster tightly when longer windows are used, which mirrors the t_c convergence documented in classic case studies such as Johansen and Sornette's work on the Nasdaq and FX crises (Johansen and Sornette 1999, 2000). The stability of ω and the presence of sustained log-periodic oscillations across neighbouring windows are also consistent with more recent LPPLS implementations that emphasise parameter constraints and multi-window validation (Filimonov and Sornette 2013; Shu 2024).

The pattern of t_c slightly lagging the actual turning point—visible in both the IBOVESPA and NYSE series—is typical in LPPL applications and has been noted repeatedly in empirical studies. This reflects the model's interpretation of t_c as the theoretical limit of the unsustainable phase rather than the exact day of market reversal. The rolling early-warning behavior shown here—broad early signals followed by compressed, high-frequency warnings near the peak—matches this narrative and resembles the escalating hazard structure reported in earlier LPPL analyses of equity and commodity bubbles.

Taken together, the three-market comparison positions the present results squarely within the observed behavior of LPPL models in the literature: stronger log-space structure, clear oscillatory signatures, stable t_c clusters, and warnings that emerge in a realistic and actionable time frame before the observed breaks.

6.4. Strengths and Limitations

A notable strength of the LPPL approach, as applied across the three markets, is its ability to impose a behavioral structure on price dynamics rather than treating bubbles as purely statistical anomalies. By modeling prices as the outcome of accelerating herding behavior with characteristic oscillations, the framework produces signals that can be interpreted in economic terms rather than simply as curve-fitting artefacts. This interpretability becomes especially valuable when the fitted trajectories, oscillatory patterns, and the predicted critical time t_c all move together in a coherent way across multiple windows, as observed in the IPC, IBOVESPA, and NYSE results.

Another strength lies in the robustness gained from the use of log-space fitting, longer window lengths, and multi-window cross-checking. These elements reduce the tendency of the model to latch onto short-lived fluctuations and instead emphasise the underlying structure of the bubble regime. The relatively tight clustering of t_c across adjacent windows—particularly in the NYSE (March–April 2025) and IBOVESPA (around mid-October 2025) cases—illustrates how this layered approach filters noise while still preserving sensitivity to real shifts in market curvature. The rolling early-warning view further complements this by translating structural model outputs into a timeline that investors, analysts, and policymakers can more easily interpret.

However, the method is not without its limitations. LPPL parameters, particularly ω and the phase-related components, are known to exhibit sloppiness: multiple combinations can produce similarly good fits, and small changes in sample boundaries can occasionally shift the estimated critical time. This is mitigated here through constrained optimization and stability checks, but the issue cannot be eliminated entirely. Moreover, the model's predictive interpretation requires caution. A predicted t_c marks the end of the bubble regime rather than a literal crash date; markets often begin turning before or after the inferred t_c, and external shocks can accelerate or delay the transition.

There is also an unavoidable sensitivity to window choice. While longer windows tend to stabilise the fits, they may smooth away shorter speculative phases that are nevertheless relevant in real-world decision-making. Conversely, shorter windows can exaggerate brief volatility spikes and pull t_c closer than justified by the broader trend. This tension is a

longstanding concern in empirical LPPL work and underscores the need for cross-window comparison rather than reliance on a single "best" fit.

Finally, LPPL signals should not be viewed in isolation. Although the model captures characteristic bubble dynamics, real markets are influenced by policy surprises, geopolitical shocks, liquidity breaks, and other events that lie outside the model's behavioral assumptions. For practical use, LPPL indicators are most valuable when combined with market-internal measures—such as breadth deterioration, credit spreads, or macro sentiment—as a way to avoid false positives and to improve the reliability of early warnings.

7. Conclusions

This study examined whether the LPPL model can identify phases of unsustainable price acceleration, produce stable estimates of the critical time, and offer an early-warning perspective across major equity markets. The results address these objectives directly. For the IPC, IBOVESPA, and NYSE Composite, the LPPL framework—especially in log form and over longer windows—captures the characteristic faster-than-exponential growth and oscillatory patterns associated with speculative bubbles. In all three indices, the estimated t_c clusters appear ahead of the downturns identified in the empirical analysis, supporting the model's ability to recover the main signatures outlined in the research questions.

The findings also connect clearly to the contributions outlined in Section 2.5. The systematic calibration for Mexico shows that LPPL dynamics are present even in a market that has received limited attention in this literature. The multi-window and rolling early-warning procedures demonstrate why single-window estimates can be unreliable in volatile emerging markets and how cross-window checks improve stability. The comparison with Brazil and the United States situates Mexico within a broader regional context, highlighting both shared bubble patterns and market-specific differences. The results also provide a practical framework for risk monitoring, suggesting that LPPL can help regulators and investors recognise when markets enter a high-risk acceleration phase.

Some limitations remain. A few well-known crashes do not satisfy the Brée and Joseph criteria, and certain parameters remain sensitive to window length, particularly when data history is limited. These issues reinforce the view that LPPL is best interpreted as signalling a rising hazard phase rather than offering precise crash dates. Used conservatively and alongside other market indicators, it can nonetheless contribute meaningfully to early-warning assessment.

Overall, the evidence shows that LPPL—when calibrated carefully and validated across multiple windows—can provide timely insight into unstable growth regimes. The cross-market comparison highlights where the model performs well and where its limits appear. Future work could test alternative crash-identification rules, incorporate sentiment or microstructure data, or develop hybrid approaches that combine LPPL's behavioral foundations with more empirical robustness.

Author Contributions: Conceptualization, S.S., A.K.G., R.M. and V.S.; methodology, S.S., R.M. and V.S.; software, S.S., R.M. and V.S.; validation, S.S. and V.S.; formal analysis, S.S., R.M. and V.S.; investigation, A.K.G., R.M. and V.S.; resources, A.K.G., R.M. and V.S.; writing—original draft preparation, S.S., A.K.G., R.M. and V.S.; writing—review and editing, A.K.G., R.M. and V.S.; supervision, A.K.G., R.M. and V.S. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Data Availability Statement: The original contributions presented in this study are included in the article. Further inquiries can be directed to the author

Conflicts of Interest: The authors declare that there are no conflicts of interest.

Abbreviations

The following abbreviations are used in this manuscript:

LPPL	Log-Periodic Power Law
TD	Trading Days
IQR	Interquartile Range
RMSE	Root Mean Squared Error
MSE	Mean Squared Error
R^2	Coefficient of Determination
MAE	Mean Absolute Error
BMV/IPC	Mexican Stock Exchange Index
NYSE	New York Stock Exchange
IBOVESPA/IBOV	Brazilian Stock Market Index

Appendix A. Minimal LPPL Pipeline (Strict 60 Trading-Day Horizon)

Purpose. This minimal, self-contained Python script reproduces the core steps used in the paper: robust date parsing, LPPL calibration on candidate windows with a strict $H = 60$ trading-day prediction horizon, selection of a common window (raw + log), three publication-quality plots (Raw, Log, Combined), a rolling early-warning summary with merged severity bands ($\leq$60 TD "Watch" and 61–90 TD "Beyond60"), and an Excel workbook with the "Verification", "All_Windows", and "Rolling_Warnings_TD" sheets.

Dependencies. `numpy`, `pandas`, `matplotlib`, `scipy`, and optionally `xlsxwriter`.

Listing A1. Minimal LPPL pipeline used for reproducibility (strict 60-TD horizon).

```
# ---------------- Minimal LPPL Pipeline (Strict 60 TD) ----------------
# Reproduces key steps in the paper with compact, readable code.
# Dependencies: numpy, pandas, matplotlib, scipy, (xlsxwriter or openpyxl)

import os
from pathlib import Path
import numpy as np
import pandas as pd
import matplotlib.pyplot as plt
from scipy.optimize import least_squares
from pandas.tseries.offsets import BDay

# ----------- USER PATHS (edit) -----------
INPUT_CSV_PATH = "/path/to/NYSE_Composite_Data_Complete_standardized.csv"
OUTPUT_DIR     = "/path/to/output/USA_Full_data_final"
RESULTS_XLSX   = "/path/to/output/USA_Full_data_final/USA_LPPL_Fit_FD_final.xlsx"

# ----------- LABELS & SETTINGS -----------
INDEX_NAME = "NYSE COMPOSITE INDEX, UNITED STATES OF AMERICA"
TITLE_MODE = "full"                  # "event" or "full"
EVENT_NO   = "Complete Data"

# Candidate windows and horizon
WINDOWS = (180, 240, 300, 360, 420)
H       = 60                        # strict prediction horizon (trading days)
END_OFFSET_TD = 10                   # end index offset when slicing windows
```

```
# Fitting knobs & constraints
MULTISTARTS   = 12
TAIL_BOOST    = 0.6
BETA_BOUNDS   = (0.1, 0.9)
OMEGA_BOUNDS  = (7.0, 12.0)
C_BOUNDS      = (-1.0, 1.0)
MIN_R2_MAIN   = 0.0


# Rolling early-warning summary
ROLL_LOOKBACK_TD   = 120
ROLL_STEP_TD       = 15
ROLL_WINDOWS_TD    = (180, 240)
ROLL_MIN_OSC       = 0.9
ROLL_MIN_AMP       = 0.03
ROLL_REQUIRE_AT_LEAST = 1
BANDS_MAX_TD = 60       # $leq$60 TD (Watch)
BANDS_EXT_TD = 90       # (60,90] TD (Beyond60)

Path(OUTPUT_DIR).mkdir(parents=True, exist_ok=True)

# ---------------- Utilities ----------------
def make_title(kind, df=None):
    if TITLE_MODE.lower() == "full" and df is not None:
        s, e = pd.to_datetime(df["Date"].min()).date(),
        ↪  pd.to_datetime(df["Date"].max()).date()
        return f"{INDEX_NAME} - Full data [{s} → {e}]: {kind}"
    return f"{INDEX_NAME} - Event {EVENT_NO}: {kind}"

def parse_dates(series):
   # Tolerant parser for mixed formats (day-first priority)
    s = series.astype(str).str.strip()
    try:
        return pd.to_datetime(s, dayfirst=True, errors="coerce")
    except Exception:
        return pd.to_datetime(s, errors="coerce")

def load_prices_csv(path):
    df = pd.read_csv(path)

# Auto-detect likely date & close columns
    date_col  = next((c for c in df.columns if str(c).lower() in
    ↪  ["date","timestamp","time","datetime"]), df.columns[0])
    close_cands = ["close","adj close","adj_close","price","close price","closing
    ↪  price","close_value","index","value"]
    price_col = next((c for c in df.columns if str(c).lower() in close_cands),
    ↪  df.columns[1])

    df = df[[date_col, price_col]].copy()
    df.columns = ["Date","Close"]
    df["Date"]  = parse_dates(df["Date"])
    df["Close"] = pd.to_numeric(df["Close"], errors="coerce")
    df = (df.dropna(subset=["Date","Close"])
            .drop_duplicates("Date")
            .sort_values("Date")
            .reset_index(drop=True))
    df["t_td"] = np.arange(len(df), dtype=float)
    df["logp"] = np.log(df["Close"].astype(float))
    return df
```

```
def td_to_calendar(df, tnum):
    last_idx = int(df["t_td"].iloc[-1])
    if tnum <= last_idx:
        i = int(np.clip(np.floor(tnum + 0.5), 0, last_idx))
        return df["Date"].iloc[i]
    delta_td = int(np.floor(tnum - last_idx + 0.5))
    return df["Date"].iloc[-1] + BDay(delta_td)

# ---------------- LPPL & metrics ----------------
def lppl(t, A, B, m, C1, C2, w, tc):
    t = np.asarray(t, float)
    x = tc - t
    if np.any(x <= 0): return np.full_like(t, np.nan, float)
    xx = np.power(x, m)
    return A + B*xx + xx*(C1*np.cos(w*np.log(x)) + C2*np.sin(w*np.log(x)))

def r2_score(y, yhat):
    ss_res = np.nansum((y - yhat)**2)
    ss_tot = np.nansum((y - np.nanmean(y))**2)
    return np.nan if ss_tot == 0 else 1.0 - ss_res/ss_tot

def oscillation_count(tc, t0, t1, omega):
    if not (t0 < t1 < tc): return 0.0
    return (omega/(2*np.pi)) * np.log((tc - t0)/(tc - t1))

# ---------------- Fitter (strict 60 TD) ----------------
def quick_fit_lppl(t, y, variant="raw", horizon_td=H, n_starts=MULTISTARTS,
↪  tail_boost=TAIL_BOOST, rng_seed=42, tc_cap_hi=None):
    t = np.asarray(t, float); y = np.asarray(y, float); t_end = float(t[-1])
    ylo, yhi = np.nanpercentile(y, [5,95])
    # Simple scale-aware bounds
    if variant == "log":
        A_bounds = (ylo - 1.0, yhi + 1.0); B_bounds = (-20.0, -1e-6)
    else:
        scale = max(1.0, yhi - ylo)
        A_bounds = (max(0.0, ylo*0.5), yhi*1.5 + 1e-6); B_bounds = (-20.0*scale,
        ↪  -1e-6)

    tc_lo = t_end + 1.0
    tc_hi = min(t_end + float(horizon_td), tc_cap_hi or (t_end +
    ↪  float(horizon_td)))
    lo = np.array([A_bounds[0], B_bounds[0], BETA_BOUNDS[0], C_BOUNDS[0],
    ↪  C_BOUNDS[0], OMEGA_BOUNDS[0], tc_lo], float)
    hi = np.array([A_bounds[1], B_bounds[1], BETA_BOUNDS[1], C_BOUNDS[1],
    ↪  C_BOUNDS[1], OMEGA_BOUNDS[1], tc_hi], float)

    A0 = float(np.nanmedian(y)); B0 = -0.1 * float(np.nanstd(y))
    x_base = np.array([A0, B0, 0.5, 0.05, -0.05, 8.8, t_end + min(30.0,
    ↪  0.5*horizon_td)], float)
    x_base = np.clip(x_base, lo, hi)

    def _one_start(x0):
        def resid(p):
            yhat = lppl(t, *p)
            if not np.all(np.isfinite(yhat)):
                return np.ones_like(y)*1e6
            w = np.ones_like(y); w[int(0.67*len(y)):] = 1.0 + tail_boost
            return (y - yhat)*w
        res = least_squares(resid, x0, bounds=(lo,hi), loss="soft_l1", f_scale=1.0,
        ↪  max_nfev=2000)
        p = res.x; p[6] = float(np.clip(p[6], lo[6], hi[6]))
        yhat = lppl(t, *p)
```

```
    return
    ↪ dict(params=dict(A=p[0],B=p[1],m=p[2],C1=p[3],C2=p[4],omega=p[5],tc_num=p[6]),
                       rmse=float(np.sqrt(np.nanmean((y - yhat)**2))),
                       mae=float(np.nanmean(np.abs(y - yhat))),
                       ↪ r2=float(r2_score(y,yhat)), yhat=yhat)

    best=None; rng=np.random.default_rng(rng_seed)
    for _ in range(n_starts):
        j = np.array([0,
                      0,
                      rng.uniform(-0.25,0.25),
                      rng.uniform(-0.3,0.3),
                      rng.uniform(-0.3,0.3),
                      rng.uniform(-1.0,1.0),
                      rng.uniform(5.0, max(6.0, horizon_td-5.0))])
        x0 = np.clip(x_base + j, lo, hi)
        cand = _one_start(x0)
        if (best is None) or (cand["rmse"] < best["rmse"]):
            best = cand
    return best

def choose_common_window(df, windows=WINDOWS, horizon_td=H,
↪   end_offset_td=END_OFFSET_TD, log_weight=0.7):
    best=None
    end = len(df) - 1 - max(0, int(end_offset_td))
    df_end_td = float(df["t_td"].iloc[-1]); tc_cap_hi = df_end_td +
    ↪ float(horizon_td)
    for L in windows:
        if end - L + 1 < 0: continue
        w = df.iloc[end-L+1:end+1].copy()
        t = w["t_td"].values
        fr = quick_fit_lppl(t, w["Close"].values, "raw", horizon_td=horizon_td,
        ↪ tc_cap_hi=tc_cap_hi)
        fl = quick_fit_lppl(t, w["logp"].values,  "log", horizon_td=horizon_td,
        ↪ tc_cap_hi=tc_cap_hi)
        if (fr["r2"] < MIN_R2_MAIN) and (fl["r2"] < MIN_R2_MAIN): continue
        r_std = float(np.nanstd(w["Close"])) or 1.0
        l_std = float(np.nanstd(w["logp"])) or 1.0
        score = (1-log_weight)*(fr["rmse"]/r_std) + log_weight*(fl["rmse"]/l_std)
        if (best is None) or (score < best["score"]):
            best = dict(L=L, window=w, fit_raw=fr, fit_log=fl, score=score)
    return best

# ---------------- Plotters ----------------
def plot_raw(df,w,fit,outfile):
    p=fit["params"]; t=w["t_td"].values
    t_ext = np.linspace(t[0], min(p["tc_num"]-1e-3, t[-1]+30), 800)
    d_ext = [td_to_calendar(df,u) for u in t_ext]
    plt.figure(figsize=(13,5.5))
    plt.plot(w["Date"], w["Close"], color="black", lw=1.2, label="Observed")
    plt.plot(d_ext, lppl(t_ext,
    ↪ p["A"],p["B"],p["m"],p["C1"],p["C2"],p["omega"],p["tc_num"]),
             color="red", lw=1.8, label="LPPL (raw)")
    plt.title(make_title(f"Raw LPPL (L={len(w)} td) | R2={fit['r2']:.3f}", df))
    plt.xlabel("Date"); plt.ylabel("Price"); plt.grid(True, alpha=0.3);
    ↪ plt.legend(); plt.tight_layout()
    plt.savefig(outfile, dpi=150); plt.close()
```

```
def plot_log(df,w,fit,outfile):
    p=fit["params"]; t=w["t_td"].values
    t_ext = np.linspace(t[0], min(p["tc_num"]-1e-3, t[-1]+30), 800)
    d_ext = [td_to_calendar(df,u) for u in t_ext]
    plt.figure(figsize=(13,5.5))
    plt.plot(w["Date"], w["logp"], color="green", lw=1.2, label="Observed log")
    plt.plot(d_ext, lppl(t_ext,
     ↪  p["A"],p["B"],p["m"],p["C1"],p["C2"],p["omega"],p["tc_num"]),
            color="orange", lw=1.8, label="LPPL (log)")
    plt.title(make_title(f"Log LPPL (L={len(w)} td) | R2={fit['r2']:.3f}", df))
    plt.xlabel("Date"); plt.ylabel("Log(Price)"); plt.grid(True, alpha=0.3);
     ↪  plt.legend(); plt.tight_layout()
    plt.savefig(outfile, dpi=150); plt.close()


def plot_combined(df,w,fr,fl,outfile):
    t=w["t_td"].values
    tc_cap = min(fr["params"]["tc_num"], fl["params"]["tc_num"]) - 1e-3
    t_ext = np.linspace(t[0], min(tc_cap, t[-1]+30), 800)
    d_ext = [td_to_calendar(df,u) for u in t_ext]
    pr,pl = fr["params"], fl["params"]
    fig, ax1 = plt.subplots(figsize=(13,5.5))
    ax1.plot(w["Date"], w["Close"], lw=1.0, color="black", label="Raw observed")
    ax1.plot(d_ext, lppl(t_ext,
     ↪  pr["A"],pr["B"],pr["m"],pr["C1"],pr["C2"],pr["omega"],pr["tc_num"]),
            lw=1.4, color="red", label="Raw LPPL")
    ax1.set_ylabel("Price")
    ax2 = ax1.twinx()
    ax2.plot(w["Date"], w["logp"], lw=1.0, color="green", label="Log observed")
    ax2.plot(d_ext, lppl(t_ext,
     ↪  pl["A"],pl["B"],pl["m"],pl["C1"],pl["C2"],pl["omega"],pl["tc_num"]),
            lw=1.4, color="orange", label="Log LPPL")
    ax2.set_ylabel("Log(Price)")
    ax1.grid(True, alpha=0.3)
    h1,l1 = ax1.get_legend_handles_labels(); h2,l2 =
     ↪  ax2.get_legend_handles_labels()
    ax1.legend(h1+h2, l1+l2, loc="best"); plt.xlabel("Date")
    plt.title(make_title(f"Combined (Raw+Log) (L={len(w)} td)", df));
     ↪  plt.tight_layout()
    plt.savefig(outfile, dpi=150); plt.close()

# ---------------- Rolling summary ($leq$60 & 61-90 TD) ----------------
def accept_fit_roll(p, t0, t1):
    if not (BETA_BOUNDS[0] < p["m"] < BETA_BOUNDS[1]): return False
    if not (OMEGA_BOUNDS[0] <= p["omega"] <= OMEGA_BOUNDS[1]): return False
    if p["B"] >= 0: return False
    if np.hypot(p["C1"], p["C2"]) < ROLL_MIN_AMP: return False
    if oscillation_count(p["tc_num"], t0, t1, p["omega"]) < ROLL_MIN_OSC: return
     ↪  False
    return True
```

```
def rolling_summary(df, windows=ROLL_WINDOWS_TD, horizon_td=H,
↪  lookback_td=ROLL_LOOKBACK_TD, step_td=ROLL_STEP_TD):
    rows=[]; end=len(df)-1
    idxs=list(range(max(0, end - lookback_td), end, step_td))
    for end_i in idxs:
        end_date = df["Date"].iloc[end_i]
        for L in windows:
            if end_i - L + 1 < 0: continue
            w = df.iloc[end_i-L+1:end_i+1].copy(); t = w["t_td"].values
            for var,y in (("raw",w["Close"].values),("log",w["logp"].values)):
                r = quick_fit_lppl(t,y,var,horizon_td=horizon_td,n_starts=8)
                p = r["params"]
                if not accept_fit_roll(p, t[0], t[-1]): continue
                lead_td = float(p["tc_num"] - df["t_td"].iloc[end_i])
                rows.append(dict(end_date=end_date, variant=var, L=L,
                ↪  tc_date=td_to_calendar(df, p["tc_num"]),
                                 tc_num=p["tc_num"], lead_td=lead_td, r2=r["r2"],
                                 ↪  rmse=r["rmse"]))
    fits = pd.DataFrame(rows)
    if fits.empty: return fits, pd.DataFrame()
    # Keep dates with at least 1 accepted fit; aggregate medians
    counts = fits.groupby("end_date")["tc_num"].count()
    keep = counts[counts >= ROLL_REQUIRE_AT_LEAST].index
    flt = fits[fits["end_date"].isin(keep)].copy()
    per = (flt.groupby(["end_date","variant"])
              .agg(n=("tc_num","count"), tc_median=("tc_date","median"),
                   lead_td=("lead_td","median"))
              .reset_index())


    def flag(lead):
        if pd.isnull(lead): return "Low"
        if lead <= 7: return "Very High"
        if lead <= 14: return "High"
        if lead <= 30: return "Elevated"
        if lead <= BANDS_MAX_TD: return "Watch"
        if lead <= BANDS_EXT_TD: return "Beyond60"
        return "Low"
    per["risk_flag"] = per["lead_td"].apply(flag)
    # Worst severity per date
    sev = {"Low":0,"Watch":1,"Beyond60":1,"Elevated":2,"High":3,"Very High":4}
    worst = (per.assign(sev=per["risk_flag"].map(sev))
                .groupby("end_date")
                .agg(n=("n","sum"), tc_median=("tc_median","median"),
                     lead_td=("lead_td","median"), max_sev=("sev","max"))
                .reset_index())
    inv = {v:k for k,v in sev.items()}
    worst["risk_flag"] = worst["max_sev"].map(inv)
    return fits, worst.drop(columns=["max_sev"]).
    rename(columns={"tc_median":"tc_median_date"})
```

```
def plot_rolling_bands(df, summ, outfile, halfspan_days=5, include_watch=True):
    s = df.set_index("Date")["Close"].astype(float).resample("W-FRI").mean()
    plt.figure(figsize=(13,5.5))
    plt.plot(s.index, s.values, color="black", lw=1.1, label="Observed")
    if not summ.empty:
        sev = {"Very High":"#ff0000","High":"#ff7f00","Elevated":"#ffd27f"}
        pale = {"Watch":"#e6f2ff","Beyond60":"#fff2cc"} if include_watch else {}
        levels = list(sev) + (list(pale) if include_watch else [])
        per = summ.set_index("end_date")["risk_flag"].sort_index()
        for dt, flag in per.items():
            start, end = dt - pd.Timedelta(days=halfspan_days), dt +
            ↪ pd.Timedelta(days=halfspan_days)
            color = sev.get(flag, pale.get(flag, None))
            if color: plt.axvspan(start, end, color=color, alpha=0.12, lw=0)
    last_date = df["Date"].iloc[-1]
    plt.axvline(last_date, ls=":", color="black", lw=0.9, label=f"Last date
    ↪ {last_date.date()}")
    plt.title(make_title("Rolling Early Warning (merged bands)", df))
    plt.xlabel("Date"); plt.ylabel("Price"); plt.grid(True, alpha=0.3);
    ↪ plt.legend(loc="best")
    plt.tight_layout(); plt.savefig(outfile, dpi=150); plt.close()

# ---------------- Excel writer ----------------
def write_workbook(path, df_ver, df_all, df_roll):
    dest = Path(path); dest.parent.mkdir(parents=True, exist_ok=True)
    engine = "xlsxwriter"
    try: import xlsxwriter  # noqa: F401
    except Exception: engine = "openpyxl"
    with pd.ExcelWriter(dest, engine=engine) as w:
        (df_ver or pd.DataFrame()).to_excel(w, index=False,
        ↪ sheet_name="Verification")
        (df_all or pd.DataFrame()).to_excel(w, index=False,
        ↪ sheet_name="All_Windows")
        (df_roll or pd.DataFrame()).to_excel(w, index=False,
        ↪ sheet_name="Rolling_Warnings_TD")
    return str(dest)

# ---------------- Main ----------------
def main():
    df = load_prices_csv(INPUT_CSV_PATH)
    print(f"[DATA] {df['Date'].iloc[0].date()} → {df['Date'].iloc[-1].date()}
    ↪ n={len(df)}")

    # 1) Choose common window and fit (strict 60 TD)
    choice = choose_common_window(df, WINDOWS, H, END_OFFSET_TD, 0.7)
    if choice is None: raise RuntimeError("Not enough data for candidate windows.")
    w, fr, fl = choice["window"], choice["fit_raw"], choice["fit_log"]



    # 2) Plots
    Path(OUTPUT_DIR).mkdir(parents=True, exist_ok=True)
    plot_raw(df, w, fr, str(Path(OUTPUT_DIR, "Raw.png")))
    plot_log(df, w, fl, str(Path(OUTPUT_DIR, "Log.png")))
    plot_combined(df, w, fr, fl, str(Path(OUTPUT_DIR, "Combined.png")))

    # 3) Rolling summary ($leq$60 and 61-90 TD)
    _, summ = rolling_summary(df, ROLL_WINDOWS_TD, H, ROLL_LOOKBACK_TD,
    ↪ ROLL_STEP_TD)
    plot_rolling_bands(df, summ, str(Path(OUTPUT_DIR, "Rolling_EarlyWarning.png")),
    ↪ include_watch=True)
```

```
    # 4) Excel: verification row (chosen window), all windows, rolling table
    ver_rows=[]
    for variant, fit in (("raw",fr),("log",fl)):
        p=fit["params"]; y_obs=w["Close"].values if variant=="raw" else
        ↪  w["logp"].values
        ver_rows.append({
            "window_len_trading_days": len(w),
            "window_start": w["Date"].iloc[0],
            "window_end":   w["Date"].iloc[-1],
            "variant": variant,
            "A": p["A"], "B": p["B"], "beta(m)": p["m"], "C1": p["C1"], "C2":
            ↪  p["C2"], "omega": p["omega"],
            "tc_index_td": p["tc_num"], "tc_date": td_to_calendar(df, p["tc_num"]),
            "R2": fit["r2"], "RMSE": fit["rmse"], "MSE":
            ↪  float(np.nanmean((fit["yhat"] - y_obs)**2)), "MAE": fit["mae"]
        })
    df_ver = pd.DataFrame(ver_rows)

    all_rows=[]; end = len(df) - 1 - max(0, int(END_OFFSET_TD))
    df_end_td=float(df["t_td"].iloc[-1]); tc_cap_hi=df_end_td+float(H)
    for L in WINDOWS:
        if end - L + 1 < 0: continue
        wL=df.iloc[end-L+1:end+1].copy(); t=wL["t_td"].values
        for var,y in (("raw",wL["Close"].values),("log",wL["logp"].values)):
            r=quick_fit_lppl(t,y,var,horizon_td=H,n_starts=8,tc_cap_hi=tc_cap_hi)
            p=r["params"]
            all_rows.append({
                "window_len_trading_days": L,
                "window_start": wL["Date"].iloc[0],
                "window_end":   wL["Date"].iloc[-1],
                "variant": var, "beta(m)": p["m"], "omega": p["omega"],
                "tc_index_td": p["tc_num"], "tc_date": td_to_calendar(df,
                ↪  p["tc_num"]),
                "R2": r["r2"], "RMSE": r["rmse"], "MAE": r["mae"]
            })
    df_all = pd.DataFrame(all_rows)
    df_roll = summ.copy() if not summ.empty else pd.DataFrame()
    write_workbook(RESULTS_XLSX, df_ver, df_all, df_roll)

    print("Saved plots to:", OUTPUT_DIR)
    print("Saved workbook to:", RESULTS_XLSX)

if __name__ == "__main__":
    main()
```

Notes. (i) The script enforces a strict $H = 60$ trading-day horizon for t_c when calibrating each window; (ii) candidate windows $L \in \{180, 240, 300, 360, 420\}$ are scored using a scale-aware blend of raw/log RMSE; (iii) rolling early-warning bands visualize merged severity intervals for "Watch" ($\leq$60 TD) and "Beyond60" (61–90 TD); (iv) the Excel workbook always contains three sheets—*Verification*, *All_Windows*, and *Rolling_Warnings_TD*.

References

Advisors, Covenant Wealth. 2025. Market Corrections: Definition and How Investors Should Respond. Available online: https://blog.commonwealth.com/ (accessed on 25 June 2025).

Bandyopadhyay, Gautam. 2016. Gold price forecasting using arima model. *Journal of Advanced Management Science* 4: 117–21. [CrossRef]

Banjade, Dhruba. 2020. Forecasting bitcoin price using artificial neural network. *SSRN Electronic Journal*. [CrossRef]

Bartolozzi, Marco, Stanislaw Drożdż, D. B. Leinweber, J. Speth, and Anthony William Thomas. 2005. Self-similar log-periodic structures in western stock markets from 2000. *International Journal of Modern Physics C* 16: 1347–61. [CrossRef]

Bergmeir, Christoph, and José M. Benítez. 2012. On the use of cross-validation for time series predictor evaluation. *Information Sciences* 191: 192–213. [CrossRef]

Blanchard, Olivier, and Mark Watson. 1982. *Bubbles, Rational Expectations and Financial Markets*. Technical Report. Cambridge: National Bureau of Economic Research. [CrossRef]
Blanchard, Olivier Jean. 1979. Speculative bubbles, crashes and rational expectations. *Economics Letters* 3: 387–89. [CrossRef]
Brée, David S., and Nathan Lael Joseph. 2013. Testing for financial crashes using the log periodic power law model. *International Review of Financial Analysis* 30: 287–97. [CrossRef]
Chai, Jian, Jiangze Du, Kin Keung Lai, and Yan Pui Lee. 2015. A hybrid least square support vector machine model with parameters optimization for stock forecasting. *Mathematical Problems in Engineering* 2015: 231394. [CrossRef]
CME Group. 2020. Market Corrections: Definitions, Causes and Examples. Available online: https://www.cmegroup.com/ (accessed on 25 June 2025).
Cont, Rama. 2001. Empirical properties of asset returns: Stylized facts and statistical issues. *Quantitative Finance* 1: 223–36. [CrossRef]
Cutler, David M., James M. Poterba, and Lawrence H. Summers.1989. What moves stock prices? *The Journal of Portfolio Management* 15: 4–12. [CrossRef]
Daniel, Kent, and Tobias J. Moskowitz. 2016. Momentum crashes. *Journal of Financial Economics* 122: 221–47. [CrossRef]
Fama, Eugene F., and Kenneth R. French. 1988. Permanent and temporary components of stock prices. *Journal of Political Economy* 96: 246–73. [CrossRef]
Filimonov, V., and D. Sornette. 2013. A stable and robust calibration scheme of the log-periodic power law model. *Physica A: Statistical Mechanics and Its Applications* 392: 3698–707. [CrossRef]
Greenwood, Robin, Samuel Hanson, Andrei Shleifer, and Jakob Ahm Sørensen. 2020. *Predictable Financial Crises*. Technical Report. Cambridge: National Bureau of Economic Research. [CrossRef]
Herwartz, Helmut. 2017. Stock return prediction under garch—An empirical assessment. *International Journal of Forecasting* 33: 569–80. [CrossRef]
Hossain, Altaf, and Mohammed Nasser. 2011. Recurrent support and relevance vector machines based model with application to forecasting volatility of financial returns. *Journal of Intelligent Learning Systems and Applications* 3: 230–41. [CrossRef]
Irina, Fátima, Villalba Padilla, and Miguel Flores-Ortega. 2013. Forecasting the variance and return of Mexican financial series with symmetric GARCH models. *Theoretical and Applied Economics* 20: 61–82.
Johansen, Anders, and Didier Sornette. 1999. Financial "anti-bubbles": Log-periodicity in gold and nikkei collapses. *International Journal of Modern Physics C* 10: 563–75. [CrossRef]
Johansen, Anders, and Didier Sornette. 2000. The nasdaq crash of april 2000: Yet another example of log-periodicity in a speculative bubble ending in a crash. *SSRN Electronic Journal*. [CrossRef]
Johansen, Anders, and Didier Sornette. 2001a. Bubbles and anti-bubbles in latin-american, asian and western stock markets: An empirical study. *International Journal of Theoretical and Applied Finance* 4: 853–920. [CrossRef]
Johansen, Anders, and Didier Sornette. 2001b. Significance of log-periodic precursors to financial crashes. *SSRN Electronic Journal*. [CrossRef]
Johansen, Anders, Didier Sornette, and Olivier Ledoit. 1999. Predicting financial crashes using discrete scale invariance. *The Journal of Risk* 1: 5–32. [CrossRef]
Johansen, Anders, Olivier Ledoit, and Didier Sornette. 2000. Crashes as critical points. *International Journal of Theoretical and Applied Finance* 3: 219–55. [CrossRef]
Liu, H., and Y. Hou. 2019. Application of bayesian neural network in prediction of stock time series. *Computer Engineering and Applications Journal* 55: 225–29.
Malkiel, Burton G. 2019. *A Random Walk Down Wall Street*, 12th ed. Discussion of adjusted prices and returns. New York: W. W. Norton & Company. [CrossRef]
Matplotlib. Matplotlib—Visualization with Python. Available online: https://matplotlib.org/ (accessed on 25 June 2025).
Moghaddam, Amin Hedayati, Moein Hedayati Moghaddam, and Morteza Esfandyari. 2016. Stock market index prediction using artificial neural network. *Journal of Economics, Finance and Administrative Science* 21: 89–93. [CrossRef]
Mohammadi, Hassan, and Lixian Su. 2010. International evidence on crude oil price dynamics: Applications of arima-garch models. *Energy Economics* 32: 1001–8. [CrossRef]
Mu, Guangyu, Nan Gao, Yuhan Wang, and Li Dai. 2023. A stock price prediction model based on investor sentiment and optimized deep learning. *IEEE Access* 11: 51353–67. [CrossRef]
Murkute, Amod, and Tanuja Sarode. 2015. Forecasting market price of stock using artificial neural network. *International Journal of Computer Applications* 124: 11–15. [CrossRef]
NumPy. numpy.linalg.lstsq—NumPy v2.3 Manual. Available online: https://numpy.org/doc/2.3/reference/generated/numpy.linalg.lstsq.html (accessed on 25 June 2025).
Oliva, Immacolata, and Roberto Renò. 2018. Optimal portfolio allocation with volatility and co-jump risk that markowitz would like. *Journal of Economic Dynamics and Control* 94: 242–56. [CrossRef]

Rajpal, Shlok Sagar, Rajesh Mahadeva, Amit Kumar Goyal, and Varun Sarda. 2025. Improving forecasting accuracy of stock market indices utilizing attention-based lstm networks with a novel asymmetric loss function. *AI* 6: 268. [CrossRef]

Rounaghi, Mohammad Mahdi, and Farzaneh Nassir Zadeh. 2016. Investigation of market efficiency and financial stability between s&p 500 and london stock exchange: Monthly and yearly forecasting of time series stock returns using arma model. *Physica A: Statistical Mechanics and Its Applications* 456: 10–21. [CrossRef]

Sarda, Varun, Yamini Karmarkar, and Neha Lakhotia Sarda. 2019. An empirical study applying log periodic structures for prediction of realty market crashes in india. *Vision: The Journal of Business Perspective* 23: 357–63. [CrossRef]

Shu, Ming. 2024. Detection of financial bubbles using a log-periodic power law singularity (LPPLS) model. *WIREs Computational Statistics* 16: e1649. [CrossRef]

Singh, Avi, Rajesh Mahadeva, Varun Sarda, and Amit Kumar Goyal. 2025. Quantitative modeling of speculative bubbles, crash dynamics, and critical transitions in the stock market using the log-periodic power-law model. *International Journal of Financial Studies* 13: 195. [CrossRef]

Sornette, Didier. 2003. Critical market crashes. *Physics Reports* 378: 1–98. [CrossRef]

Sornette, Didier. 2017. *Bubbles, Crises, and Crashes in Emergent Markets*. Princeton: Princeton University Press. [CrossRef]

Sornette, Didier, and Peter Cauwels. 2015. Financial bubbles: Mechanisms and diagnostics. *Review of Behavioral Economics* 2: 279–305. [CrossRef]

Tirole, Jean. 1985. Asset bubbles and overlapping generations. *Econometrica* 53: 1497. [CrossRef]

Virtanen, Pauli, Ralf Gommers, Travis E. Oliphant, Matt Haberland, Tyler Reddy, David Cournapeau, Evgeni Burovski, Pearu Peterson, Warren Weckesser, Jonathan Bright, and et al. 2020. SciPy 1.0: Fundamental algorithms for scientific computing in python. *Nature Methods* 17: 261–72. [CrossRef] [PubMed]

Wes McKinney. 2010. Data Structures for Statistical Computing in Python. Paper presented at 9th Python in Science Conference, Austin, TX, USA, June 28–July 3; pp. 56–61. [CrossRef]

West, Kenneth D. 1987. A specification test for speculative bubbles. *The Quarterly Journal of Economics* 102: 553. [CrossRef]

World Bank. Mexico Overview: Development News, Research, Data. Available online: https://www.worldbank.org/en/country/mexico (accessed on 5 June 2025).

Xiu, Yuxuan, Guanying Wang, and Wai Kin Victor Chan. 2021. Crash diagnosis and price rebound prediction in nyse composite index based on visibility graph and time-evolving stock correlation network. *Entropy* 23: 1612. [CrossRef] [PubMed]

Yahoo Finance. IPC Mexico (m̂xx) Historical Data. Available online: https://sg.finance.yahoo.com/quote/%5EN225/ (accessed on 5 June 2025).

Zhou, Wei-Xing, and Didier Sornette. 2004. Antibubble and prediction of china's stock market and real estate. *Physica A: Statistical Mechanics and Its Applications* 337: 243–68. [CrossRef]

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.