Masoumeh Fathi

Essays on Power Laws and Financial Markets

ACTA WASAENSIA 560

University of Vaasa
VAASAN YLIOPISTO

ISBN 978-952-395-206-5 (print)
978-952-395-207-2 (online)

ISSN 0355-2667 (Acta Wasaensia 560, print)
2323-9123 (Acta Wasaensia 560, online)

URN https://urn.fi/URN:ISBN:978-952-395-207-2

PunaMusta Oy, Joensuu, 2025.

ACADEMIC DISSERTATION

To be presented, with the permission of the Board of the School of Accounting and Finance of the University of Vaasa, for public examination on the 25th of September, 2025, at noon.

Article based dissertation, School of Accounting and Finance, Finance.

Author	Masoumeh Fathi https://orcid.org/0000-0003-3506-6379
Supervisor(s)	Associate Professor Klaus Grobys University of Vaasa. School of Accounting and Finance, Finance.
	Professor Janne Äijö University of Vaasa. School of Accounting and Finance, Finance.
Custos	Associate Professor Klaus Grobys University of Vaasa. School of Accounting and Finance, Finance.
Reviewers	Professor Mika Vaihekoski University of Turku, Department of Finance.
	Professor Simone Alfarano Universitat Jaume I, Department of Economics.
Opponent	Professor Mika Vaihekoski University of Turku, Department of Finance.

Tiivistelmä

Tämä väitöskirja koostuu neljästä esseestä, jotka tutkivat potenssilakien moninaisten sovellusten merkitystä rahoitusmarkkinoilla keskittyen niiden rooliin arvopapereiden hinnoittelussa, riskienhallinnassa ja markkinadynamiikan ennustamisessa. Se käsittelee kahta keskeistä potenssilakien tyyppiä rahoituskirjallisuudessa—potenssilakeja, joita käytetään tuottojakaumien häntien mallintamiseen, sekä potenssilakeja, joilla mallinnetaan arvopapereiden hintaprosessien supereksponentiaalista kasvua—ja pyrkii näin vastaamaan perinteisten perinteisten normaalijakaumaoletukseen perustuvien mallien rajoituksiin sekä niiden soveltamiseen rahoitusmarkkinoiden käyttäytymisen mallintamisessa.

Kolme ensimmäistä esseetä analysoivat tuottojakaumien häntien potenssilakiluonnetta keskittyen toteutuneisiin ja vaihteluväliin perustuviin variansseihin eri omaisuusluokissa ja markkinoilla. Ensimmäinen essee tutkii Faman ja Frenchin osakefaktorien toteutuneita variansseja. Testaamalla eksplisiittisesti vakiintunutta lognormaalista jakaumaa potenssilakeja vastaan tutkimus osoittaa, että toteutuneet faktorivarianssit noudattavat raskashäntäistä jakaumaa, joka on yhdenmukainen potenssilakien kanssa. Lisätestit osoittavat, että Mandelbrot'n äärettömän varianssin hypoteesia ei voida hylätä. Toinen essee laajentaa analyysiä valuuttamarkkinoille (FX) keskittyen kymmenen suurimman valuuttaparin (G10) toteutuneisiin variansseihin. Se vahvistaa potenssilakikäyttäytymisen esiintymisen, jossa varianssien varianssi on määrittelemätön, korostaen perinteisten volatiliteetin arviointimenetelmien riittämättömyyttä valuuttamarkkinoiden riskienhallinnassa. Kolmas essee tutkii vaihteluväliin perustuvia toteutuneita variansseja eri omaisuusluokkien, kuten osakkeiden, hyödykkeiden ja kryptovaluuttojen välillä. Se havaitsee johdonmukaisen potenssilakieksponentin säätelevän markkinoiden välistä vaihteluväliin perustuvaa toteutunutta varianssia, mikä viittaa universaaliin riskidynamiikkaan, jota voivat ohjata systeemiset tai käyttäytymiseen liittyvät tekijät, kuten laumakäyttäytyminen. Neljäs essee laajentaa potenssilakien soveltamista hintaprosesseihin hyödyntämällä log-periodista potenssilakisingulariteettimallia (LPPLS) tunnistaakseen kuplien muodostumista ja ennustaakseen regiinivaihdosten ajoitusta, jotka johtavat romahdukseen vakiintuneessa rahoitusmarkkinoiden anomaliassa (ts. osakehintojen momentum-ilmiö). Tutkimus vahvistaa LPPLS-signatuurien esiintymisen ennen momentum-vetoisia kuplia, mikä korostaa mallin tehokkuutta varhaisessa havaitsemisessa ja riskienhallinnassa.

Asiasanat: potenssilait, varianssi, raskashäntäisyys, häntäriski, ääretön varianssi, laumakäyttäytyminen.

Abstract

This doctoral dissertation consists of four essays that examine the implications of diverse applications of power laws in financial markets, focusing on their role in asset pricing, risk management, and the prediction of market dynamics. It explores two primary types of power laws in financial literature—power laws employed to model the tails of return distributions and power laws used to model super-exponential growth in price processes of financial assets—to address key limitations of traditional Gaussian-based models and their applications in modeling financial market behavior.

The first three essays analyze the power-law characteristics of the tails of return distributions, focusing on realized and range-based variances across diverse financial assets and markets. The first essay investigates the realized variances of Fama and French equity factors. By explicitly testing the well-established log-normal distribution against power laws, the study documents that realized factor variances exhibit heavy-tailed behavior consistent with power laws. Further tests provide evidence that Mandelbrot's infinite variance hypothesis cannot be rejected. The second essay extends the analysis to foreign exchange (FX) markets, focusing on realized variances of Group of Ten (G10) currency pairs. It confirms the presence of power-law behavior, with undefined variances of variances, emphasizing the inadequacy of traditional volatility estimation techniques in FX risk management. The third essay examines the range-based variances across multiple asset classes, including equities, commodities, and cryptocurrencies. It finds a consistent power-law exponent governing range-based cross-market variance, suggesting a universal risk dynamic potentially influenced by systemic or behavioral factors like herding. The fourth essay extends the application of power laws to price processes, employing the Log-Periodic Power-Law Singularity (LPPLS) model to identify bubble formations and predict the timing of regime switches that lead to a crash in a well-established financial market anomaly (viz, stock price momentum). The study confirms the presence of LPPLS signatures preceding momentum-driven bubbles, emphasizing the model's effectiveness in early detection and risk mitigation.

Keywords: power laws, variance, heavy tails, tail risk, infinite variance, herding.

ACKNOWLEDGEMENT

Pursuing a PhD has been one of the most meaningful and challenging experiences of my life, filled by many ups and downs. In my second year, I became a mother, and my life turned into a balancing act between pursuing my PhD, motherhood, and managing all the other responsibilities life brings. Being away from the university brought its own challenges as well. However, I tried my best to stay balanced, even with all the challenges I faced. I'm not sure how successful I was, but I always tried to be the best version of myself. Now, standing at the end of this journey with a lovely five-year-old daughter by my side, I can see how those hard days shaped me into someone stronger. I'm truly grateful to everyone who stood by me and supported me in moving forward.

My deep gratitude goes first to my first supervisor, Associate Professor Klaus Grobys, for his exceptional guidance, unlimited support, expertise, and insightful feedback throughout every stage of this dissertation. I could not have completed this journey without his encouragement and commitment. Working under his supervision has been both an honor and an invaluable learning experience. I've learned a lot from him and hope to continue learning as I move forward.

Second, I would like to express my sincere gratitude to Professor Janne Äijö, my second supervisor, for his invaluable support throughout my doctoral journey. His guidance, encouragement, and practical insights were a vital part of this work. I am also deeply thankful for his commitment in securing funding for my studies, which made it possible for me to pursue and complete this research.

I am deeply grateful to the two pre-examiners of this dissertation, Professor Mika Vaihekoski from the University of Turku and Professor Simone Alfarano from Universitat Jaume I, for their valuable time, insightful comments and constructive feedback, all of which greatly contributed to improving this work.

I am also grateful to Professor James W. Kolari from Texas A&M University for co-authoring one of the papers included in this thesis.

I am especially grateful to Professor Sami Vähämaa, Head of the Department of Accounting and Finance, for fostering an inspiring academic environment by organizing regular research seminars, conferences, and various workshops that have greatly enriched the academic atmosphere of the department. I am particularly thankful for his warm and encouraging response to my initial email before I applied

for this PhD program. His encouragement at that early stage meant a lot to me and helped me decide to take this path.

I would also like to thank Dean Professor Marko Järvenpää and former Dean Professor Helinä Saarela for their constant support and for providing the facilities and opportunities that contributed significantly to my academic development.

I am also thankful to Associate Professor Tatiana King and Professor Timothy King for their excellent organization of seminars and conferences, as well as for their efforts in facilitating doctoral students' participation in academic events.

I am also grateful to Professor Mikko Leppämäki, Director of the Graduate School of Finance (GSF) and the Finnish Doctoral Program in Economics (FDPE), for providing valuable courses, workshops, and academic opportunities that greatly enriched the experience of doctoral students.

My heartfelt thanks are extended to Dr. Shaker Ahmad for his friendship over the years and the helpful advice he offered as a senior colleague throughout my studies. He generously shared his time and expertise, and I truly valued his readiness to assist whenever needed.

I am thankful as well to my colleagues and friends Maria Zhukova, Ulviyya Pekkarinen, and Marko Hanhimäki for their friendship and the many pleasant moments we shared throughout these years. Their presence made this journey more enjoyable and memorable.

I am also thankful to the Evald and Hilda Nissi Foundation, the Suomen Arvopaperimarkkinoiden Edistämissäätiö and the University of Vaasa for their financial support, which was essential in enabling me to focus on my research without financial constraints.

Above all, I extend my heartfelt thanks to my family for their love, encouragement, and support throughout this journey. I am thankful to my husband, Dr. Hossein Hafezi, who has always been by my side with patience and support. Finally, special thanks to my parents, who, though far away, are always close to my heart.

August 7, 2025

Masoumeh Fathi

This work is dedicated to my lovely daughter, who brought light to my life,

Nila.

Contents

Figures

Tables

Essays

[1] Fathi, M., Grobys, K., & Äijö, J. (2025). A common component of Fama and French factor variances. *The North American Journal of Economics and Finance, 75*(Part A), 102292. https://doi.org/10.1016/j.najef.2024.102292. CC BY.

[2] Fathi, M., Grobys, K., & Kolari, J. W. (2024). On the realized risk of foreign exchange rates: A fractal perspective. *Journal of Risk and Financial Management, 17*(2), 79. https://doi.org/10.3390/jrfm17020079. CC BY.

[3] Fathi, M., & Grobys, K. (2025). Modeling variance risk in financial markets using power-laws: New evidence from the Garman-Klass variance estimator. *Quantitative Finance*, 1–26. https://doi.org/10.1080/14697688.2025.2529485. CC BY-NC-ND.

[4] Fathi, M., Grobys, K., & Äijö, J. (2024). Momentum bubbles: A new perspective derived from log-periodicity. *Proceedings of the 29th International Conference on Forecasting Financial Markets (ffm29): Advances for Asset Management, Saïd Business School, University of Oxford. Proceedings of the 9th International Conference on Research in Management and Economics, Cambridge.*

1 INTRODUCTION

The aim of this doctoral dissertation is to explore the implications of diverse applications of power laws across various aspects of financial analysis, including asset pricing, risk management, and market behavior prediction. The dissertation is structured into four key essays. The first essay examines the risk associated with the Fama and French equity factors by analyzing their realized variances. The second essay evaluates the uncertainty in the FX market, utilizing power laws to model realized FX rate variances. The third essay investigates the Garman-Klass range-based variance of five key financial assets using power laws. Finally, the fourth essay employs the LPPLS model to recognize bubble formations and predict the timing of regime changes that lead to a crash within the popular momentum anomaly in financial markets.

Power laws have a long history in both natural and social sciences (Bak, 1996; Brown et al., 2002; Newman, 2005; West, 2017) and are recognized as universal phenomena across various disciplines, including physics, biology, earth and planetary sciences, economics, finance, computer science, demography, and the social sciences (Newman, 2005; West, 2017). There are two major classes of power laws commonly discussed within the finance literature. The first is the power laws governing the tails of return distributions, where the likelihood of extreme returns decreases as their size increases, reflecting the heavy-tailed nature of financial markets. It indicates as pareto rank/frequency distribution with a negatively-sloped line indicating the inverse power-law signature (Andriani & McKelvey, 2009). The second is power laws of price processes, which is a univariate scaling law that describes how asset prices grow over time. Even though both power laws—describing the tails in return distributions and modeling prices—are conceptually different, they are both derived from a common underlying theory of herding behavior and mutual imitative contagion in speculative markets (Lux, 1995). This dissertation explores the implications of different applications of both types of power laws within finance literature through its four essays to better understand and model financial market dynamics. While the first three essays focus on the application of the power laws to model the tails of return distributions, capturing extreme events and heavy-tailed behavior, the fourth essay examines the power laws in price processes, highlighting its role in modeling growth dynamics and predicting critical market events.

The early 20th century set the foundation for statistical modeling in financial research, beginning with Louis Bachelier (1900), who first proposed that asset prices follow a normal (Gaussian) distribution. Decades later, this idea influenced the development of Modern Portfolio Theory by Markowitz (1952) and the Capital Asset

Pricing Model (CAPM) by Sharpe (1964), both of which relied on the assumption of normally distributed asset returns. These models established a systematic approach to portfolio optimization and risk management. However, subsequent studies have documented that financial markets exhibit heavier tails than the Gaussian distribution, indicating that the Gaussian distribution may not always capture the complex dynamics of financial markets and tends to underestimate the probability of extreme events (Mandelbrot, 1963; Fama, 1965; Mantegna & Stanley, 1995; Cont, 2001; Lux & Marchesi, 1999; Gabaix et al., 2003; Taleb, 2007, 2020). The failure of prevailing financial models to predict and prepare financial sectors for economic downturns bears witness to this fact.

The inability of Gaussian-based models to accurately represent extreme market fluctuations highlights the need for alternative models that can account for the heavy-tailed features observed in financial data. Benoit Mandelbrot, widely known as the father of fractal geometry, was the first to propose a shift from Gaussian distributions to a new family of probability laws, known as stable Paretian distributions, in his seminal 1963 study (Mandelbrot, 1963). By analyzing cotton price changes, Mandelbrot found that the extreme tails of empirical distributions are typically more peaked than what would be expected from a Gaussian distribution, necessitating a different approach to the theory of random walks in speculative prices. Recognizing these deviations, he proposed the stable Lévy-Pareto distribution, a subset of power-law distributions, as a robust framework for understanding the overall distribution of averaged price changes. Notably, his analysis showed that the tail index alpha (α), which represents the heaviness of the tails, is approximately less than 2 ($\alpha \approx 1.7$), indicating that the population second moment (variance) is infinite. Building on Mandelbrot's work, Fama (1963, 1964) provided empirical validation for the stable Paretian hypothesis, noting that when the population variance of first differences is infinite, the sample variance becomes an unreliable measure, and statistical methods like least-squares regression may produce false conclusions. In the closing remarks of his 1963 review on Mandelbrot's stable Paretian hypothesis, Fama emphasized the necessity of developing "*more adequate statistical tools for dealing with stable Paretian distributions*".

Although strong evidence supported Mandelbrot's findings, economists avoided adopting them due to the limited availability of statistical tools for testing non-Gaussian distributions compared to the well-established techniques for Gaussian ones (Salazar, 2016). According to Salazar (2016), the lack of reliable methods to validate Lévy stable distributions made economists reluctant to invest effort in further developing Mandelbrot's ideas. This resistance persisted for decades, as the widespread use of Gaussian-based methods offered a more practical, yet less precise, framework for financial modeling. Nassim Nicholas Taleb, a pioneer in challenging

the conventional wisdom of using Gaussian distributions, directly critiques this tendency, arguing: "*There are a lot of theories on why things should be power-laws, as sort of exceptions to the way things work probabilistically. But it seems that the opposite idea is never presented: power-laws should be the norm, and the Gaussian a special case*." (Taleb, 2020, p. 91).

According to Taleb (2020), statistical estimation relies on two fundamental principles: the Central Limit Theorem (CLT) and the Law of Large Numbers (LLN). The CLT states that, for a sufficiently large sample size, the distribution of the sample mean approaches a Gaussian distribution, regardless of the original distribution of the data. The LLN, on the other hand, ensures that as the sample size increases, the sample mean converges to the true population mean, improving the accuracy of estimation. However, as Taleb (2020) emphasizes, these principles can fail in the presence of fat-tailed distributions due to the disproportionate impact of extreme values. Specifically, the CLT may not hold because heavy tails can result in infinite or undefined variance, which prevents convergence to a Gaussian distribution. Similarly, the LLN becomes unreliable as extreme values may dominate the sample, leading to slow or even non-existent convergence of the sample mean to the population mean. Consequently, statistical estimation becomes fundamentally flawed in such contexts.

This perspective challenges the foundational assumptions of modern portfolio theory (Markowitz, 1952) and its extensions, including CAPM and factor-based models, which rely on finite-variance distributions. While Taleb's critique exposes critical vulnerabilities in traditional models, it is important to note that several strands of literature have proposed frameworks that explicitly address tail risks. For instance, models based on stable Paretian distributions and other heavy-tailed frameworks, such as Conditional Value-at-Risk (CVaR) optimization, attempt to incorporate tail risks explicitly into portfolio construction (Rockafellar & Uryasev, 2000, 2002; Lim, Shanthikumar, & Vahn, 2011). Additionally, models grounded in Extreme Value Theory (Embrechts, Klüppelberg, & Mikosch, 1997) focus on the statistical behavior of rare events in the tails of distributions. Regime-switching frameworks (Hamilton, 1989) and stochastic volatility models such as Heston (1993) also allow for more flexible modeling of asset dynamics than traditional mean-variance approaches. While none of these models fully resolve the unpredictability emphasized in Taleb's arguments, they represent meaningful steps toward building financial models that are more robust to the realities of market turbulence and tail risk.

Building upon these developments, this dissertation contributes to the growing body of literature by examining both types of power laws—those governing the tails of return distributions and those describing price processes—to evaluate their

effectiveness in modeling complex market dynamics. It extensively explores the implications of diverse applications of power laws across various aspects of financial analysis, including asset pricing, risk management, and market dynamics prediction, particularly concerning extreme events and the formation of financial bubbles. Furthermore, the study seeks to address the inconsistency between theoretical models and the inherent characteristics of financial markets.

In addition, this dissertation focuses on the critical role of the second moment (variance) in financial analysis, particularly in understanding and modeling realized risk through the first three essays. As highlighted by Andersen et al. (2003), the second moment of asset returns is the most prominent and empirically dynamic aspect of the return distribution. This emphasis also stems from variance's crucial role as one of the two key elements in evaluating financial assets within Markowitz's mean-variance framework. Markowitz's (1952) seminal work laid the foundation for the CAPM and, later, multifactor models, including the influential Fama-French models (Fama & French, 1992, 1993, 2015, 2017, 2018, 2020), which remain central to empirical finance. Given this foundational role, the analysis of variance is essential for validating the principles of modern portfolio theory and advancing related research. Thus, this dissertation's focus on variance, particularly through the first three essays, not only highlights its importance as a key element of these foundational theories but also aims to contribute to the ongoing debate in volatility modeling by exploring variance within the context of alternative statistical approaches.

While standard models in the finance literature often assume constant volatility and correlation, it is a stylized fact that these characteristics vary over time. This empirical observation has led to a significant body of research focused on understanding the behavior and distribution of market volatility. Traditionally, studies have used parametric models such as ARCH (Engle, 1982) and stochastic volatility (Taylor, 1982; Ghysels et al., 1996; Shephard, 2005), as well as implied volatility from options prices (Black & Scholes, 1973; Latane & Rendleman, 1976; Christensen & Prabhala, 1998), all of which require strong distributional assumptions, while implied volatility further depends on assumptions about the market price of volatility risk. Moreover, one alternative method involves using squared returns over the chosen return interval to generate model-free and unbiased estimates of ex post realized volatility. However, as highlighted by Andersen et al. (2001), squared returns are noisy measures and provide limited reliability in capturing true underlying volatility. To address these challenges, Andersen et al. (2001) played a pioneering role by formalizing the concept of realized volatility through a new method that calculates daily volatility as the sum of squared intraday returns, providing an aggregated and assumption-free measurement. This method eliminates the need for strong distributional assumptions and enables the measurement of actual volatility over fixed intervals such as daily or weekly periods. Subsequent research extended this approach to explore the distributional features of realized volatility (Barndorff-Nielsen & Shephard, 2002), address issues like

microstructure noise and asynchronous trading (Zhang et al., 2005), and separate volatility into continuous and jump components (Barndorff-Nielsen & Shephard, 2004). Today, realized volatility is widely used in empirical finance for accurately measuring and forecasting short-term market fluctuations.

Alternatively, range-based volatility estimators, derived from intraday price ranges, offer a more efficient alternative to traditional squared-return measures. Notable early contributions include the models proposed by Parkinson (1980) and Garman and Klass (1980). These estimators are particularly effective in reducing the impact of microstructure noise and low trading volume (Chou et al. 2010). This dissertation analyzes both realized and range-based volatility measures by applying the realized variance approach introduced by Andersen et al. (2001, 2003) in the first two essays and employing the Garman-Klass range-based estimator (Garman & Klass, 1980) in the third essay. Through these complementary methods, the dissertation provides a comparative perspective on volatility modeling across different asset classes and market conditions.

The first essay explores the risk of the Fama and French equity factors in terms of their realized variances. Specifically, it investigates whether realized factor variances exhibit power-laws behavior and, if so, to what degree. The focus on Fama and French factors is driven by their significance as a cornerstone in asset pricing models and their pervasive influence in empirical finance. The use of realized variances is motivated by their ability to effectively reflect the underlying risk dynamics of asset returns, capturing not only the magnitude but also the temporal dependencies of return fluctuations (Andersen et al., 2001). In addition, realized variances offer a more precise measure of volatility compared to traditional estimators, as they are less prone to certain biases and measurement errors (Barndorff-Nielsen & Shephard, 2002). Recognizing that empirical distributions often underestimates extreme events, the study focuses on modeling the tail behavior of variance—a region where rare but high-impact risks reside. Moreover, while several well-established studies suggest that realized asset volatility is approximately lognormally distributed (e.g., Andersen et al., 2001, 2001a, 2001b), this study empirically tests this claim in the context of realized equity factor variances.

The results show that realized factor variances exhibit strong power-laws behavior, effectively ruling out log-normal and exponential distributions as the underlying distributions generating these variances. Mandelbrot's infinite variance hypothesis on realized factor variances confirms that the null hypothesis of a power-law tail exponent $\hat{\alpha} = 1.9$ remains valid, suggesting a common component governing factor variance risk and its statistically undefined character. The paper concludes that the statistically significant factor premiums widely reported in the finance literature are

sample-specific phenomena, because, in finite samples, it is impossible to verify or observe them.

The second essay focuses on the realized risk of FX rates from a fractal perspective. It examines whether weekly realized variances of G10 currency pairs would follow power-law distributions as opposed to the time-honored log-normal distribution. The findings indicate that the realized variances of FX rates are governed by power-law distributions, implying risk homogeneity within specific asset classes. The results of the goodness-of-fit test suggest that the power-law model is preferred over competing distributions, including the log-normal and exponential distributions, for most FX rates. Further analysis shows that estimated power-law exponents for most realized FX variances are $2 < \hat{\alpha} < 3$, indicating that the variance of realized FX variance is statistically undefined. The results, based on weekly realized FX variances, align closely with Grobys' (2023c) findings from daily realized FX variances data, indicating consistency across time frequencies. This supports Mandelbrot's fractal perspective, which suggests that power-law exponents remain unchanged regardless of time scale, reflecting the fractal-like behavior of realized FX variances.

While the methodological framework in the second essay is similar to that of the first, the choice of weekly aggregation for realized variances—rather than monthly, as in the first essay—was deliberate and reflects key differences in market structure and volatility characteristics. The FX market operates continuously with high liquidity and more frequent trading activity than equity factor portfolios. Weekly data thus provide a better balance between capturing meaningful short-term risk dynamics and mitigating the noise inherent in higher-frequency observations. In contrast, the first essay focuses on Fama-French equity factors, where monthly aggregation is more consistent with how factor premiums are typically evaluated in the empirical asset pricing literature. Moreover, due to the shorter historical availability of FX data compared to the Fama-French factors, the use of weekly frequency in the second essay ensures a sufficient number of observations for reliably estimating the power-law tail exponents, which require large sample sizes to yield stable and meaningful results. Using weekly data in the second study also facilitates closer comparison with prior work on realized FX volatility, such as Grobys (2023c), and aligns with the goal of examining persistence and tail behavior across different temporal scales.

Expanding the scope of the earlier essays, the third essay examines the presence of power-law behavior across the range-based variance risks of distinct asset classes, highlighting its potential as a universal property in financial market risk. It investigates the range-based variance risk of five key financial asset markets—S&P 500, gold, crude oil, the USD/GBP exchange rate, and Bitcoin—using the noise-efficient Garman-Klass variance estimator (Garman & Klass 1980). The results reveal

that Garman-Klass variances exhibit heavy tails across all five asset markets, strongly supporting the power-law hypothesis with $\alpha \approx 2.8$, indicating a universal law governing the cross-sectional variances of otherwise unrelated asset markets. The goodness-of-fit tests provide strong evidence that power-law models are the most suitable representation of the data across all assets, while alternative distributions, including log-normal and exponential, were consistently rejected. The analysis of tail exponents confirms that values consistently fall within $2 < \hat{\alpha} < 3$, implying infinite variance of range-based variance, consistent with the findings from the two previous essays using realized variance of Fama-French factors and G10 currency pairs.

The existence of such commonality across various aspects of financial analysis—from equity factors to FX and diverse asset classes—examined through the first three essays provides compelling evidence supporting power laws as a universal characteristic of financial market risk dynamics. This suggests a unique underlying risk dynamic, potentially driven by similar systemic or behavioral forces, such as herding behavior, consistent with the findings of Lux (1995). Identifying a universal cross-domain exponent emphasizes the significance of unified risk modeling approaches and introduces a new perspective to the literature on power-law behavior.

The fourth essay investigates whether power-law frameworks can serve as a basis for predicting market dynamics, with a particular focus on bubble formations and their critical points, through the application of the power laws of price processes. For this purpose, this study utilizes Sornette et al.'s (1996) LPPLS model, incorporating the power laws of price processes and log-periodic oscillations to identify and predict critical events in financial markets, such as bubbles and crashes. Specifically, the fourth essay examines the performance of LPPLS models in predicting critical market events, particularly momentum-driven bubbles. The key findings indicate that LPPLS signatures act as early warnings before significant momentum bubbles formation, validating the LPPLS model's ability to identify early signs of bubble formation in a zero-cost trading strategy. This demonstrates that power-law frameworks are not limited to describing market dynamics; they also offer practical applications for forecasting critical market events and managing associated risks. The paper argues that significant LPPLS signatures in the log-compounded returns of the momentum factor indicate that momentum bubbles are not random events but artifacts of herding behavior, which underpins the theoretical foundation of the LPPLS model.

The rest of this introductory chapter is structured as follows: Section 2 outlines the overall contributions of the dissertation and each individual essay, Section 3 presents a brief background relevant to the essays included in the dissertation, Section 4

discusses the research methods employed in this dissertation, and finally, Section 5 summarizes the essays.

2 CONTRIBUTION OF THE DISSERTATION

This dissertation contributes to the finance literature by exploring both types of power laws—the power laws employed to model the tails of return distributions and the power laws used to model super-exponential growth in price processes of financial assets—addressing key limitations of traditional approaches in finance literature and examining their applications in modeling financial market behavior. Within the power laws of tail distributions, it incorporates Mandelbrot's infinite variance hypothesis (Mandelbrot, 1963), which challenges conventional Gaussian-based models by emphasizing the heavy-tailed nature of financial data and its higher likelihood of extreme market events, highlighting the importance of accurately capturing tail risk. By focusing on the tails of return distributions, the first three essays of this dissertation advance the understanding of realized variances in financial markets and their implications for asset pricing and risk management. These essays critically evaluate the established argument in finance literature that realized asset variance follows a log-normal distribution (e.g., Andersen et al., 2001), contrasting this with power-law frameworks.

In addition, the fourth essay of this dissertation contributes to the finance literature by examining the application of the power laws in price processes within the framework of the LPPLS model. It examines how the power laws of price processes can capture the scaling behavior of asset prices, particularly their super-exponential growth leading up to critical points, such as bubbles and crashes. This study highlights the relevance of power-law frameworks in understanding price dynamics and their applicability in predicting critical market phenomena.

The first essay of this dissertation contributes to the finance literature by exploring the risk of well-established Fama and French equity factors and finds that realized factor variances exhibit strong power-laws behavior. The findings confirm the infinite theoretical variance hypothesis for all the realized factor variances analyzed in the study. While these findings contrast with the consensus in the pertinent literature observed by Lux and Alfarano (2016), stating that financial asset returns exhibit power-laws behavior with finite variances, they align with Mandelbrot's (1963) infinite theoretical variance hypothesis, which emerged from his early research on cotton price changes. Building on the observations of Grobys (2021), which hypothesized that replication failures documented in the finance literature (Hou et al., 2020) may stem from undefined risk in financial markets, this study argues that both replication failures and inconsistent methodological choices (Hou et al., 2020; Dick-Nielsen et al., 2023) could be attributed to undefined risk associated with factor variances. Specifically, it highlights that statistical inferences derived from methodologies like ordinary least squares (OLS) in factor models, which assume the

existence of finite variances, are likely to produce misleading results, aligning with Fama's (1963) implications regarding the consequences of infinite variances.

Further, this study contributes to the methodological literature on risk assessment and volatility modeling by utilizing realized variances to provide a more precise understanding of the risk dynamics associated with asset returns, in line with Andersen et al. (2001) and Barndorff-Nielsen and Shephard (2002). According to Andersen et al. (2001), realized variances offer a more accurate and reliable measure of volatility by capturing both the magnitude and temporal dependencies of return fluctuations. Furthermore, the reduced susceptibility of realized variances to biases and measurement errors, compared to traditional volatility estimation methods (Barndorff-Nielsen & Shephard, 2002), highlights their importance in enhancing financial analysis and advancing risk modeling practices. Another contribution of this study is its novel perspective on analyzing realized variances, as realized variances capture information that standard models, such as generalized autoregressive conditional heteroskedasticity (GARCH) models—typically based on normality assumptions—cannot reveal (e.g., Bubák et al., 2011; Andersen et al., 2003, 2004).

Moreover, while established literature suggests that realized asset volatility is log-normally distributed (e.g., Andersen et al., 2001, 2001a, 2001b), this study is the first to explicitly test log-normality against power-law distributions for realized equity factor variances. In doing so, this paper contributes to the expanding empirical literature on power laws (Grobys, 2021, 2023c) by explicitly testing log-normal distributions against proposed power-law frameworks—an issue that has previously remained unexamined.

The findings of this essay carry important implications for understanding risk dynamics in financial markets. The consistent power-laws behavior observed in realized factor variances, with a shared exponent of $\alpha \approx 2$, suggests that these patterns are not market-specific phenomenon but are driven by systematic market behavior, possibly herding behavior, which can deviate prices from fundamental values, consistent with the findings of Lux (1995). The study concludes that the factor investing industry is exposed to a considerably higher level of extreme risk than earlier believed as indicated by the power-law exponent captures via extrapolation low-probability deviations not seen in the data (Taleb, 2020).

The second essay contributes to the finance literature by exploring the uncertainty in the FX market and finds that power laws provide a better data fit and more reliable probabilities for the arrivals of extreme events in the realized FX variances. Literature incorporating realized volatility to model FX risk has identified several issues with realized variances where standard models like GARCH often fail to capture, such as asymmetric volatility in response to exchange rate shocks (Wang & Yang, 2009), long-

memory effects in volatility dynamics (Corsi, 2004), the need to separate continuous volatility from jumps to avoid biased estimates (Andersen et al., 2005), the dynamic and crisis-sensitive nature of volatility spillovers (Bubák et al., 2011), and the time-varying, non-Gaussian characteristics of realized volatility itself (Corsi et al., 2008). Addressing these challenges, this essay contributes to the literature on power laws in financial economics by building on the findings of Calvet and Fisher (2004) and Lux et al. (2014), who demonstrated that power-law models often outperform GARCH models in forecasting future volatility. The essay discuses that power-law models provide a more robust framework for capturing the scaling behavior and persistence observed in market realized volatility.

This study extends recent power-laws modeling literature by Grobys (2021, 2023c), which demonstrated the effectiveness of power-law functions in capturing realized variances across various asset classes, including FX rates, the S&P 500, commodities, and cryptocurrencies. The results of the second essay align with Grobys' (2021) findings, consistently showing that the power-law null model is valid for most realized variances, with tail exponents ($2 < \hat{\alpha} < 3$) indicating infinite variance in the realized FX variances. This research makes contributions in several distinct ways. Methodologically, while Grobys (2023c) applied Bayes' rule to provide evidence on the distribution associated with the most extreme events, this study extends that approach by explicitly performing statistical tests using probability density functions. In terms of data frequency, this study incorporates lower-frequency data, offering a novel perspective compared to previous studies (Grobys, 2021, 2023c). Furthermore, in comparing power-law distributions with alternative candidates (log-normal and exponential distributions), this research employs the goodness-of-fit test of Clauset et al. (2009) to evaluate their suitability, providing a more robust alternative to the Bayes' rule-based approach used in Grobys (2023c), which only identifies the distribution associated with the most extreme events.

This study further contributes to the literature by documenting the fractal-like properties of realized FX variance, as the findings with weekly data align closely with the dynamics observed in the daily data reported by Grobys (2023c). This study aligns with Mandelbrot's (2008) argument that price movements in the currency market are wild, marked by such uncertainty that the second moment of realized variances does not converge. Moreover, given the high frequency of false discoveries in financial economics, scientific replications of existing research have been strongly encouraged in studies such as Hou et al. (2020) and Serra-Garcia and Gneezy (2021). By modeling realized FX rate variances using an alternative but related methodology to those employed by Grobys (2023c), this study aligns with the principles outlined

by Hamermesh (2007) [1] by providing replication findings that enhance the understanding of the fractal-like behavior of FX variances. Further, consistent with the findings of the first essay, the second essay also confirms the suitability of power laws over the log-normal distribution, providing additional evidence that realized asset variances align more closely with power laws in contrast to the conclusion of well-established studies, which suggest that realized volatility follows a log-normal distribution (e.g., Andersen et al., 2001a, 2001b, 2003). The paper concludes that the FX market is significantly riskier than previously believed, as the presence of infinite variance in realized FX variances challenges traditional assumptions of finite risk and the reliability of conventional statistical models, such as those based on normality.

The third essay investigates the range-based variance risk of five major financial asset markets—S&P 500, gold, crude oil, the USD/GBP exchange rate, and Bitcoin—using the Garman-Klass variance estimator (Garman & Klass 1980) and finds that Garman-Klass variances exhibit heavy tails across all five markets, strongly supporting the power-law hypothesis. This study makes several key contributions by addressing important methodological, theoretical, and empirical gaps, advancing the understanding of variance in financial markets and their relevance to asset pricing and risk management. The first contribution of this study lies in its methodological refinement: building on Grobys (2021), who used the Parkinson estimator (Parkinson, 1980) to analyze range-based variance across unrelated markets, this study adopts the noise-efficient Garman–Klass estimator (Garman & Klass, 1980), in line with Molnár's (2012) finding that the Garman–Klass estimator outperforms alternative estimators like Parkinson's when addressing noise in high-frequency data. This adjustment enhances the precision of range-based variance estimates, providing a more reliable foundation for investigating heavy-tailed distributions like power laws. The second contribution of this study is the explicit comparison of log-normal and exponential distributions with the power-law distribution across all range-based asset variances, which was not conducted by Grobys (2021). The rejection of alternative models in favor of the power-law, based on comprehensive goodness-of-fit tests applied to range-based variance across unrelated asset markets, extends the literature that critiques the log-normal assumption of realized volatility proposed by Andersen et al. (2001), and aligns with the findings of the two previous essays in this dissertation. Third, while previous literature has tested common power-law behavior of realized and range-based variances within related asset categories, such as equity factors (Fathi et al. 2025) and foreign exchange markets

[1] Hamermesh (2007) emphasizes the critical role of replication in ensuring the validity and robustness of empirical findings in economics, advocating for reproducibility, robustness checks, and transparency to build cumulative scientific knowledge. He highlights that replication not only confirms findings but also refines or challenges existing conclusions, contributing to the integrity of research.

(Grobys 2024), no prior research has explored this commonality across unrelated markets; thus, this study contributes to this literature by addressing this gap through an examination of commonality across unrelated markets and offers a novel cross-sectional analysis by exploring whether range-based variances in unrelated financial markets—such as equities, commodities, exchange rates, and cryptocurrencies—are governed by a common power-law.

Collectively, essays 1, 2, and 3 of this dissertation, with a focus on the power laws of the tail of return distributions, extend the finance literature by providing robust empirical evidence that realized variances across diverse financial markets—equity factors, FX, and major asset classes—follow power-law distributions rather than the traditionally assumed log-normal framework. In addition, Mandelbrot's infinite variance hypothesis holds for the equity market represented by Fama-French factors, but for the other markets (FX and major asset classes) the variance of variance and higher moments are found to be infinite. These findings highlight Mandelbrot's infinite variance hypothesis and provide various empirical examples of the inadequacy of traditional log-normal and finite-moment frameworks in capturing the heavy-tailed behavior and extreme risks in financial markets. The commonality observed within each market (equity factors in the first essay and FX in the second essay) as well as across multiple markets (S&P 500, gold, crude oil, USD/GBP, and Bitcoin in the third essay) provides evidence for a unified risk dynamic, likely driven by systemic or behavioral forces such as herding behavior, in line with Lux (1995). Recognizing this unified risk dynamic highlights the need for alternative frameworks that more effectively capture the interconnected nature of financial market risks. In addition, the heavy-tailed nature of realized variances has substantial implications for portfolio optimization, asset pricing, and risk management, necessitating frameworks that effectively account for the inherent extremity in financial data. By advancing a robust foundation for understanding variance risks within broad market categories, this dissertation contributes to a more precise and unified approach to financial modeling, particularly in the context of infinite variance dynamics.

Finally, the fourth essay takes a different perspective by utilizing the power laws of price processes to examine whether it provides a framework for identifying and predicting critical points in financial markets, such as bubbles and crashes. Using the LPPLS model, the essay investigates the predictability of momentum-driven bubbles and finds that LPPLS signatures act as early warnings before significant bubble formations. The findings contribute to the field of financial market analysis by providing empirical support for the predictive power of the LPPLS model concerning significant momentum bubbles and crashes. While prior research has utilized the LPPLS model to examine stock market bubbles (e.g., Johansen et al., 1999; Vandewalle et al., 1999), cryptocurrency bubbles (e.g., Wheatley et al., 2018), and sovereign debt

crises (Geraskin & Fantazzini, 2013), this study broadens this scope, offering novel insights into momentum-driven market behavior. Additionally, the study contributes to the literature by examining the effect of the calibration window on the precision of LPPLS model predictions for critical market events. While Sornette (2017) highlights the increasing precision of the LPPLS model as it approaches a crash, Brée, Challet, and Peirano (2013) raise concerns about the reliability of its critical time estimates due to significant fluctuations when input data change. Addressing these concerns, this research is the first to evaluate the reliability of LPPLS predictions based on the calibration window, using Grobys' (2023) recently proposed three-stage estimation approach.

The study also contributes to the literature by exploring the predictability of momentum bubbles and their potential management. Building on the argument by Barroso and Santa-Clara (2015) that momentum strategies carry the risk of sudden crashes, this study suggests that momentum risk can be mitigated if bubbles and subsequent crashes are predictable. By demonstrating how investors could adjust positions to manage these risks, the study offers novel insights into the practical implications of bubble predictability for momentum strategies. The study also provides a behavioral explanation for the occurrence of momentum bubbles. The LPPLS model is derived from the theoretical framework of imitation or herding behavior among traders. Consequently, the study argues that evidence of significant LPPLS signatures in the momentum strategy implies that momentum bubbles are not random events but rather artifacts of herding behavior, in line with Lux (1995) and Sornette and Cauwels (2014).

3 THEORETICAL BACKGROUND

This section provides a brief overview of the background underlying this dissertation and the four essays included.

3.1 Power laws of prices versus power laws of return distribution tails

Power laws play a crucial role in understanding the dynamics of financial markets, as they provide insight into both price behavior and market risks. In finance, two key types of power laws are particularly important: the power laws used to model super-exponential growth in price processes of financial assets and the power laws employed to model the tails of return distributions. While the power laws of a price process describe transient super-exponential growth in asset prices, the power laws of return distributions describe the probability and impact of extreme market events. Thus, it is crucial to distinguish between these two types, as they provide important information on market behavior and risk management.

The power law of a price process is typically not a probability distribution in the traditional sense, but rather a scaling relationship that describes how prices grow over time. In standard financial models, it is generally assumed that asset price growth follows a stochastic proportional process, driven by the mechanism of compound returns or interest rates. This implies that, apart from its variations, the price of a stock is typically expected to grow at a consistent exponential rate of return, on average, as shown in the left plot of Figure 1. The general form of this power laws is:

$$P(t) = p_0 t^{\alpha}, \quad \alpha > 0, \tag{1}$$

where $P(t)$ is the price at time t, p_0 is the initial price, and α is the scaling exponent which determines the rate of the growth process. This relationship describes how prices grow in financial markets or financial assets, often reflecting super-exponential growth due to compounding effects. It provides a framework for understanding the inherent self-reinforcing dynamics of market growth, which align with a positive tail exponent α, reflecting the overall long-term upward trend in prices.

Figure 1. Power laws of price (left) versus power laws of the tail of return (right).

In contrast, the power laws of returns describe the tail of the return distribution exceeding a certain threshold ($x > x_{min}$), which is often characterized by heavy tails. This tail behavior follows a power law with a decreasing pattern, indicating that large deviations are rare but possible. For values below the threshold ($x < x_{min}$), the return distribution is governed by a different, thin-tailed process. The general form of the power law for the tail of return distributions is often expressed as:

$$P(x > x_{min}) \sim x^{-\alpha}, \quad \alpha > 1, \tag{2}$$

where $P(x > x_{min})$ represents the cumulative distribution function (CDF). It gives the probability that the return x exceeds a threshold x_{min}, and α is the tail exponent that quantifies the heaviness of the tail. The right plot in Figure 1 depicts the power laws of the tails of return distributions, illustrating that as returns increase, their probability decreases following a power-law relationship, as described by Equation (2). This illustrates the fat-tailed nature of return distributions, where extreme returns, such as 6-sigma events—referred to as discontinuities by Mandelbrot (2008)—are virtually impossible under a Gaussian distribution due to the rapid decay of tail probabilities, yet they occur as an inherent feature of financial markets.

This dissertation explores the implications of the applications of both types of power laws across its four essays. The first three essays, focus on the power laws of tails of return distributions, as it offers a robust framework for understanding and modeling the heavy-tailed characteristics of financial data. This emphasis highlights the limitations of traditional Gaussian-based models in addressing tail risks and emphasizes the significance of power-law distributions in accurately capturing the probability and impact of extreme market events. However, in the last paper, the focus shifts to the power laws of price, as applied within the framework of the LPPLS model. This model utilizes the scaling behavior of prices to predict critical market

events, such as bubbles and crashes, emphasizing the importance of power-law dynamics in understanding and forecasting price movements.

3.1.1 Power-law distributions

Power laws (scaling law) describe relationships in which a relative change in one quantity leads to a proportional relative change in another, regardless of the initial magnitudes of those quantities. This scale-invariance property allows power laws to effectively capture systems characterized by extreme variability and rare events. This class of distributions are widely observed in various natural and social phenomena, ranging from the distribution of earthquake magnitudes and city populations to wealth and financial market behaviors.

In economics, the study of power-laws behavior dates back to the work of Vilfredo Pareto in the nineteenth century (Pareto, 1896). Pareto observed that wealth and income distributions consistently follow a power-laws pattern, with a small fraction of the population holding a disproportionately large share of total wealth. This phenomenon, known as the Pareto principle or the 80/20 rule, remains central to understanding inequality at wealth redistribution. Beyond wealth distribution, power laws are evident in other economic phenomena, such as the sizes of firms, price fluctuations, trading volume, market capitalization, and many other aspects of financial and economic systems.

These distributions are characterized by heavy tails, meaning that extreme events, though rare, are more likely to occur compared to other types of distributions like the normal or exponential. The economic significance of power-law distributions lies in its ability to capture the dynamics of systems where extreme events have a major effect. Consequently, understanding power-laws behavior is crucial for risk management and financial stability. Traditional Gaussian-based models, which assume finite variance and low probability of extreme events, often underestimate the likelihood of financial crises or market crashes. Power laws, by contrast, provide a framework for accurately modeling such tail risks, capturing the higher probabilities of extreme events and their disproportionate impact on financial systems.

3.1.2 Mathematical formulation and moments of the power-law distribution

The mathematical formulation of a power-law function, used in the first three essays following Grobys (2021), is defined as:

$$p(x) = Cx^{-\alpha}, \tag{3}$$

where $C = (\alpha - 1)x_{MIN}^{\alpha-1}$ with $\alpha \in \{\mathbb{R}_+ | \alpha > 1\}$, x represents the variable governed by a power-law process with $x \in \{\mathbb{R}_+ | x_{MIN} \leq x < \infty\}$, x_{MIN} is the minimum value of x that is governed by a power-law process, and α is the tail exponent that determines the heaviness of the distribution's tail.

For the observed values of x exceeding x_{MIN} $(X > x_{MIN})$, the conditional expected first moment is defined as:

$$E[X|X > x_{MIN}] = \int_{x_{MIN}}^{\infty} xp(x)dx = \frac{(\alpha-1)}{(\alpha-2)}x_{MIN}. \tag{4}$$

The conditional expected second moment, or the variance for $X > x_{MIN}$, is expressed as:

$$E[X^2|X > x_{MIN}] = \int_{x_{MIN}}^{\infty} x^2p(x)dx = \frac{(\alpha-1)}{(\alpha-3)}x_{MIN}^2. \tag{5}$$

Consequently, higher conditional expected moments of order k follow the general form:

$$E[X^k|X > x_{MIN}] = \frac{(\alpha-1)}{(\alpha-1-k)}x_{MIN}^k. \tag{6}$$

Equations (4) and (5) indicate that the theoretical mean of the power-law distribution is defined only when $\alpha > 2$, and the variance exists only for $\alpha > 3$.

3.1.3 Power-laws versus Gaussian distribution

The Gaussian distribution, a familiar bell curve as shown in Figure 2, has been widely used in financial modeling since it was introduced by the French mathematician Louis Bachelier in 1900 (Bachelier, 1900) to describe the price movements of French government bonds. The main characteristic of this distribution is that observations move around the mean, while the likelihood of deviations decrease exponentially by moving further away from the mean, making outliers increasingly rare and seemingly negligible. The 68-95-99.7 rule provides a clear understanding of this variation: approximately 68% of observations fall within one standard deviation (1-sigma events) of the mean, about 95% lie within two standard deviations, and nearly 99.7%

are within three standard deviations (see Figure 2). This stands in contrast to real-world financial events, such as the stock market crash of October 1987, known as Black Monday, when the S&P 500 index experienced a 22-sigma decline—an event that, under a Gaussian distribution, would have odds of less than one in 10^{50} (Mandelbrot, 2008), effectively making it impossible to observe in a "Gaussian World." When the tails of a distribution become fatter—similar to what occurs in financial markets—the probability of an event falling within one standard deviation of the mean increases significantly, often ranging between 75% and 95%, compared to 68% for a standard normal distribution (Taleb, 2020).

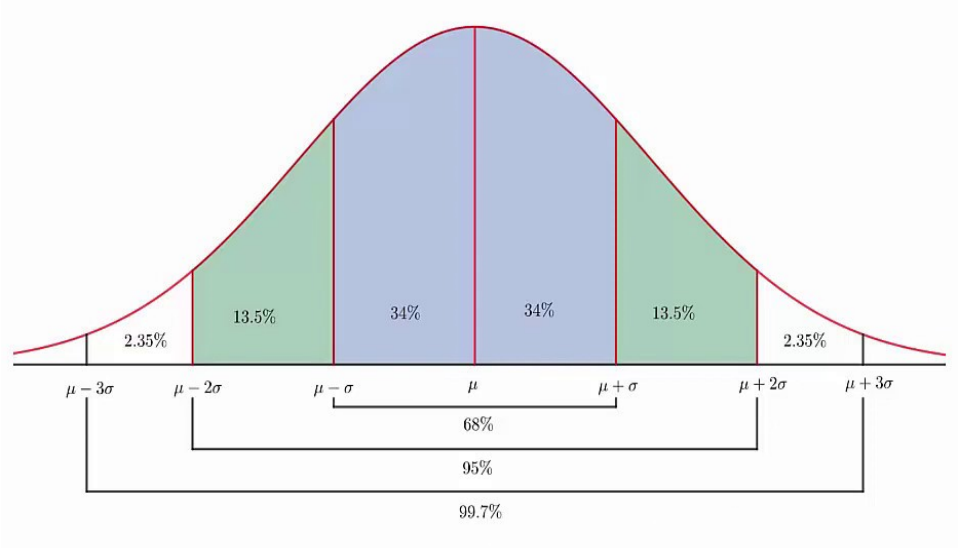

Figure 2. Gaussian distribution.

According to Taleb (2020), thicker tails in a distribution lead to higher peaks, smaller shoulders, and a higher frequency of extreme deviations from the mean, as illustrated by Figure 3, which compares the kernel density estimator of 30-minute price increments for S&P 500 index futures (red curve) with a Gaussian distribution (blue curve). This is because probabilities must sum to 1; thus, increasing the probability mass in one area necessarily reduces it in another. Consequently, the top 1% of observations in fat-tailed distributions holds more probability mass than in a normal distribution because the tails decay more slowly, making extreme events more frequent and impactful. For example, using a sample of 10,000,000 random values drawn from a standard normal distribution, the top 1% and 2% of observations contributed approximately 3.35% and 6.07% to the cumulative total, respectively. In contrast, using the same sample size, the top 1% and 2% of observations in a power-law distribution (with $\alpha = 2.5$ and $x_{min} = 1$) contributed approximately 6.30% and 9.55% to the cumulative total. This highlights the fundamental difference between

Gaussian and fat-tailed distributions, particularly in their ability to model extreme events.

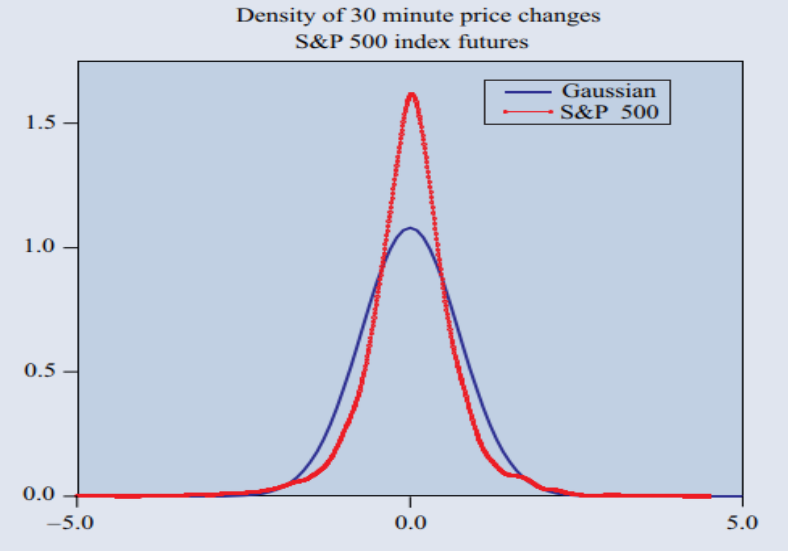

Figure 3. Kernel estimator of the density of 30-minute price increments. Reprinted from Cont (2001).

3.2 Mandelbrot's infinite variance hypothesis

Mandelbrot's infinite variance hypothesis is a concept to explain the behavior of certain financial and economic data that exhibit heavy-tailed distributions. In his seminal work (Mandelbrot, 1963), Mandelbrot demonstrated that empirical financial data often exhibit heavy tails, where extreme values (such as large price changes or returns) occur more frequently than predicted by Gaussian models. These heavy tails reflect the occurrence of rare but impactful events, such as market crashes or sudden price increases, which traditional models fail to adequately capture. Unlike Gaussian distributions, which assume finite variance and exponential tail decay, heavy-tailed distributions are better modeled by power-law distributions. For these distributions, the probability of an event is proportional to a power of the magnitude of the event. The infinite variance implies that the risk of extreme events cannot be bounded, making traditional risk management techniques, which rely on finite variance assumptions, insufficient for accurately assessing or mitigating risks.

Mandelbrot not only identifies the inadequacy of Gaussian-based models but also introduces the concepts of fractal and self-similar structures in financial time series. In his later work (Mandelbrot, 1982), Mandelbrot showed that financial time series often exhibit self-similarity, meaning that patterns observed at one time scale (e.g.,

daily returns) recur at different time scales (e.g., weekly or monthly returns). This self-similarity reflects the fractal nature of financial markets, where price movements display similar statistical properties across different time scales. The fractal-like structure emphasizes the limitations of Gaussian-based models, as Gaussian distributions assume independent events with thin tails and fail to account for the self-similarity and fat-tailed behavior observed in financial markets. This is a core principle of the CLT, which states that the sum (or average) of a large number of independent, identically distributed (*i.i.d.*) random variables with finite variance will tend to follow a Gaussian distribution, regardless of the original distribution of the variables. In other words, Gaussian distributions are stable under aggregation—their shape remains consistent when combining data across different time scales. Although Gaussian distributions retain their form under aggregation, they fail to capture the self-similar, scale-invariant behavior characteristic of fractal models, where statistical patterns persist and repeat across multiple time scales. In contrast, power-law distributions exhibit true scale-invariance, meaning they preserve their fundamental shape under rescaling. This property aligns with the fractal characteristics observed in financial markets, where extreme events and price fluctuations follow similar statistical patterns across different time frames. Importantly, power-law distributions account for the fat tails and frequent extreme events that Gaussian models underestimate. Thereby, while crashes are extreme negative events, their display of a power-law signature implies that markets operate without constraints, allowing for unrestricted movement both upward and downward (Andriani & McKelvey, 2009).

Formally, a time series $X(t)$ is self-similar if it satisfies the relation:

$$X(at) \stackrel{d}{=} a^H X(t), \tag{7}$$

where H is the Hurst exponent, a parameter that measures the persistence or memory in the system. Values of H different from 0.5 suggest long-term dependence, further challenging the classical assumption of *i.i.d*. returns.

Another crucial aspect of the infinite variance hypothesis is its connection to stable Lévy processes, which Mandelbrot proposed as a more appropriate framework for modeling financial returns. Unlike the Gaussian distribution, stable Lévy distributions can capture the skewness and heavy tails observed in empirical data, making them a valuable alternative to Gaussian-based models.

The characteristic function of a stable Lévy process is given by:

$$\phi(k) = E\left[e^{ikX(t)}\right] = \exp\left(-\sigma^{\alpha} \mid k \mid^{\alpha} \left(1 - i\beta sgn(k) tan\frac{\pi\alpha}{2}\right) + i\mu k\right), \tag{8}$$

where, α governs the tail thickness, β controls the skewness, k represents the frequency variable or Fourier transform variable and σ and μ represent the scale and location parameters, respectively. For $1 < \alpha < 2$, the distribution exhibits infinite variance, making it suitable for capturing the extreme and unpredictable fluctuations observed in financial markets, as empirically demonstrated by Mandelbrot (1963), who found a tail exponent of 1.7 for cotton prices.

Mandelbrot's hypothesis highlights how traditional risk models, which assume finite variance, systematically underestimate both the likelihood and severity of extreme events, leading to insufficient risk management strategies. By contrast, models that incorporate heavy-tailed distributions and self-similar structures provide a more realistic representation of market dynamics, enabling better preparation for rare but catastrophic events. In practice, employing models consistent with the infinite variance hypothesis can effectively enhance portfolio management, derivative pricing, and systemic risk assessment.

3.3 LPPLS model

This section provides the theoretical background for the fourth essay, which utilizes the power laws of the price process within the framework of the LPPLS model. The LPPLS model introduced by Sornette et al. (1996) is a mathematical framework designed to analyze and predict financial bubbles by capturing their super-exponential growth and oscillatory behavior. This model aims to estimate the critical time when a bubble is likely to burst, providing insights for risk management and market analysis. The LPPLS model is built on the dynamics of two distinct market agents: rational investors, who understand the fundamental asset values, and noise traders, whose herding behavior can drive prices away from these values. Within this framework, the rational investors continue to invest despite inflated prices, expecting higher returns unless a market bubble bursts and leads to a crash. Eventually, a bubble bursts when imitation among traders increases to a critical point, and all the (noise) traders shift to selling, as outlined by Johansen et al. (2000).

The LPPLS model combines two elements: a power-law growth to model the rapid price acceleration and a log-periodic term to account for oscillations that become more frequent as a potential crash approaches. The power-law growth implies that financial bubbles exhibit super-exponential growth in prices as they approach a critical time t_c, at which the bubble is likely to burst, with the power-law component capturing this accelerated growth. The log-periodic oscillations imply that in addition to the power-law growth, financial bubbles often show periodic fluctuations in price. These oscillations become more frequent as the critical time t_c is approached.

According to Geraskin and Fantazzini (2013), the LPPLS model can be defined as Equation (9):

$$log\left[E\left(p(t)\right)\right] \approx\ A + B(t_c - t)^{\beta} -\ (1 + C\cos[\omega\log(t_c - t) + \varphi]), \qquad (9)$$

where $log\left[E\left(p(t)\right)\right]$ denotes the logarithm of expected asset price at time t, A represents the logarithm of the expected asset price at the critical time, B measures the rate of increase in the logarithm of the asset price as it approaches the critical time t_c, and β is a parameter that captures the rate of super-exponential growth. The coefficient C measures the relative magnitude of the oscillations along the growth path, ω denotes the frequency of these oscillations during the bubble phase, and φ is the phase adjustment parameter.

The fourth essay investigates the capability of LPPLS in predicting momentum bubbles. This focus stems from the unique nature of momentum bubbles, which differ from those driven by technological innovation or macroeconomic shifts. Momentum bubbles are fueled by the synchronized behavior of investors adopting similar strategies, often disregarding the intrinsic value of assets. Such behavior is particularly prevalent during periods of high uncertainty, as investors tend to follow the crowd to avoid the perceived risks of acting independently (Avery & Zemsky, 1998). Consequently, herding behavior has been recognized as a significant factor in the escalation of these speculative bubbles, as it increases price movements and market volatility (Lux, 1995). As markets approach a critical point, price movements exhibit power-laws behavior, characterized by accelerating growth and increasing volatility (Sornette, 2003, 2017). In this regime, small changes in market conditions can cause disproportionately large price fluctuations, making the system highly sensitive to disturbances (Sornette & Cauwels, 2014). Thus, once the endogenous system approaches the critical point, not only does the price process follow a power laws, but price changes themselves are also governed by a power-law distribution—a finding first documented in the seminal study by Mandelbrot (1963), who identified that price changes of speculative financial assets exhibit power-laws behavior.

4 RESEARCH METHODS

This section outlines the methodologies employed in this dissertation across the four essays included. As discussed in the previous section, this research explores two types of power laws within the finance framework: the power laws used to model the tails of return distributions and the power laws used to model super-exponential growth in price processes of financial assets. While the first three essays applied the concept of the power laws to the tail of return distributions, the fourth essay utilized the power laws in price dynamics, within the LPPLS model, to analyze and predict critical market events such as bubbles and crashes.

4.1 Methods used for power laws of the tail of return distributions

Fitting power-law distributions to empirical data is a complex process that requires careful consideration of statistical methods and assumptions. There are different methods to fit power laws to empirical distributions, each with its strengths and limitations, depending on the characteristics of the data and the desired accuracy of the fit. These approaches typically provide an estimate for the scaling parameter α and, in some cases, also determine the threshold x_{min}. Traditional approaches, such as the visual methods using the linear fit of the log-log plot and the linear fitting using the least-squares regression, are simple and visually intuitive but are prone to biases, particularly in the tail region (Goldstein et al., 2004). A more robust approach, focusing on the mathematical properties of the distribution, maximizes the likelihood function to estimate parameters directly from the data. However, since power-law distributions hold only above a certain threshold (x_{min}), careful selection of x_{min} is crucial for accurate estimation. To address this, Clauset et al. (2009) introduced a method based on the Kolmogorov-Smirnov (KS) statistic, which minimizes the difference between the empirical and theoretical CDFs, providing a more reliable approach for evaluating the fit of power-law models. This approach provides precise parameter estimates, particularly in the limit of large sample sizes.

The first three essays in this dissertation apply the methodology proposed by Clauset et al. (2009), which integrates maximum- likelihood estimation (MLE) with goodness-of-fit tests based on the KS statistic. The MLE method for estimating the power-law tail exponent, following Clauset et al. (2009), is given by:

$$\hat{\alpha} = 1 + N\left(\sum_{i=1}^{N} \ln\left(\frac{x_i}{x_{MIN}}\right)\right)^{-1}, \tag{10}$$

where $\hat{\alpha}$ is the estimated tail exponent, N is the number of observations exceeding x_{MIN}, and other terms maintain their previously defined meanings. The selection of x_{MIN} is determined by minimizing the distance between the empirical data and the power-law model (Clauset et al., 2009). This is achieved using the KS distance, D, which measures the maximum difference between the CDFs of the observed data and the fitted model. The distance D is given by:

$$D = MAX_{x \geq x_{MIN}} |S(x) - P(x)|, \tag{11}$$

where $S(x)$ represents the CDF of the empirical data for observations $x \geq x_{MIN}$, and $P(x)$ is the CDF of the fitted power-law model. The optimalx_{MIN} is the value that minimizes D. Additionally, as described by Clauset et al. (2009), the standard deviation of the estimated power-law exponents is expressed as:

$$\sigma = \frac{\hat{\alpha}-1}{\sqrt{N}} + O\left(\frac{1}{N}\right). \tag{12}$$

To verify the appropriateness of a power-law fit, Clauset et al. (2009) introduced a goodness-of-fit test based on the KS distance, D, which compares the empirical data to synthetic data generated from a power-law distribution using specific x_{MIN} and $\hat{\alpha}$ values. The null hypothesis of this test is that the empirical data and the synthetic data generated from a power-law distribution using the estimated x_{MIN} and $\hat{\alpha}$ values belong to the same distribution, suggesting the power-law model is a plausible fit. While the alternative hypothesis suggests that the empirical and synthetic data originate from different distributions, implying that other models should be considered.

Additionally, model comparison methods are employed to validate the fit and ensure the appropriate application of the power laws following Clauset et al. (2009). Incorporating validation methods, such as comparing power-law fits with alternative models (e.g., log-normal or exponential distributions), ensures that the power-law hypothesis is applied appropriately. These comparative approaches help assess whether power-laws behavior accurately describes the data or if another distribution offers a better fit, thereby enhancing the robustness of the analysis.

While the findings presented in the first three essays provide robust empirical evidence of power-law behavior in the realized variances of equity factors, FX rates, and major asset classes using the method proposed by Clauset et al. (2009), it is important to acknowledge a critical limitation of this approach. Although this method has become standard practice in empirical research due to its systematic use of MLE and goodness-of-fit test, it is sensitive to the selection of the lower bound x_{MIN} , which determines where the power-law regime begins, as Clauset et al. (2009) themselves acknowledge. According to Deluca and Corral (2013), this method can fail to

correctly identify true power-law behavior, especially when the data includes a shoulder region—a transitional zone between the bulk of the distribution and the heavy tail. This concern is particularly relevant in finite samples where the distinction between the shoulder and tail may be ambiguous. Including data from this region can lead to underestimation of the power-law exponent, thereby mischaracterizing the tail behavior. Recognizing this limitation, future research could examine the robustness of the results by incorporating alternative approaches for threshold selection—such as the optimal sample fraction methods proposed by Drees and Kaufmann (1998)—which offer refined control over the trade-off between bias and variance in tail exponent estimation.

4.2 Methods used for power laws of price process

The fourth essay in this dissertation employed the power-laws framework within the LPPLS model to analyze and predict critical market events, with a particular focus on momentum bubbles. The application of the power laws in the price process is evident in the first term of the LPPLS model in Equation (9). The term $A + B(t_c - t)^{\beta}$ is a simple power-law model for financial log-prices that indicates how prices grow over time. This term captures the super-exponential growth of asset prices approaching a critical time t_c, explaining their scaling behavior, where A is the logarithm of the expected asset price at the critical time t_c. The parameter β, the power-law exponent controlling faster-than-exponential price growth, typically ranges between 0 and 1. It governs the curvature of this growth, with smaller values indicating slower acceleration and values closer to 1 reflecting rapid, compounding increases. This simple power-law term captures the dynamics of speculative bubbles by showing how prices deviate from normal growth patterns due to self-reinforcing mechanisms, such as investor herding and overconfidence (Sornette & Cauwels, 2014).

The calibration of the LPPLS model in the fourth essay is conducted using a three-stage estimation approach, as detailed by Grobys (2023). This approach facilitates the step-by-step improvement of model parameters, starting with a straightforward approximation and progressing toward a more complex and accurate representation of the dynamics captured by the LPPLS model. This methodology is designed to overcome the challenges of estimating highly nonlinear parameters by simplifying the LPPLS model to a form that can be calibrated using standard software.

The fourth essay utilizes daily data of 10 momentum portfolios from November 1926 to November 2023, sourced from Kenneth R. French's data library. The momentum factor, winners-minus-losers (WML), is calculated as the difference between the returns of past winners and losers, based on cumulative 12 months returns and

updated daily using decile breakpoints. The analysis focuses on compounded returns, as they effectively capture long-term growth dynamics and align with the LPPLS model's ability to detect super-exponential patterns associated with bubbles.

The first stage of the estimation process, as described by Grobys (2023), is referred to "the worst-case scenario," which simplifies the LPPLS framework by assigning predefined values to certain parameters. As mentioned earlier, this stage incorporates Sornette's (2017) basic power-law model for financial logarithmic prices, represented as:

$$ln[p(t)] = A + B(t_c - t)^{\beta}, \tag{13}$$

where $ln[p(t)]$ is the natural logarithm of the price $p(t)$ at time t, with the other variables defined as before. Additionally, for this model specification, as outlined by Grobys (2023), the following constraints are applied: $0.1 \leq \beta \leq 0.9$, $A > 0$ and $B < 0$.

Replacing the price $P(t)$ with the compounded return of momentum factor, while keeping the exponent β fixed at predefined values from the set $\bar{\beta} \in \{0.2, 0.4, 0.6, 0.8\}$ and setting $B = -1$, Equation (13) simplifies into the following expression:

$$ln[comp(t)] = ln[comp(T)] - (t_c - t)^{\bar{\beta}}, \tag{14}$$

where $comp(t)$ represents the compounded return of the momentum factor at time t and $[comp(T)]$ in association with $t_c = T + 1$ indicates that the last observed compounded return of the momentum strategy for a given sample serves as the compounded return at the critical time t_c. This implies that the critical event is expected to occur in the subsequent time period.

In the second stage, the model is optimized by minimizing the sum of squared residuals for Equation (14), subject to the following constraints:

$$\begin{cases} t_c \geq T + 1 \\ 0.1 \leq \beta \leq 0.9 \end{cases} \tag{15}$$

In the third stage, the optimized parameter values derived from the second stage, denoted as $\Phi_1^* = (A^*, B^*, t_c^*, \beta^*)$, are used as initial inputs to calibrate a more advanced version of the LPPLS model, incorporating log-periodic oscillations as:

$$ln[comp(t)] = Ln[comp(T)] - (t_c - t)^{\bar{\beta}} - (1 + C\cos[\omega\, Ln(t_c - t) + \varphi]), \tag{16}$$

where coefficient C measures the relative scale of the oscillations along the growth path, ω defines the angular frequency of these oscillations during the bubbles, and φ serves as the phase adjustment parameter. This step refines the initial estimates by

including additional parameters (C, ω, ϕ) , which capture the oscillatory deviations around the power-law growth, adding an additional constraint to the conditions in Equation (15):

$$\{5 \leq \omega \leq 15. \tag{17}$$

According to Sornette (2017), the parameter Φ is left unrestricted, as imposing constraints on the phase parameter Φ is not considered meaningful. Additionally, the constraint on ω follows Sornette's (2017) empirical findings, which suggest the customary use of the range $5 < \omega < 15$. Finally, the final optimized parameter vector for advanced LPPLS model is stored in $\Phi_2^* = (A^{**}, B^{**}, t_c^{**}, \beta^{**}, C^{**}, \omega^{**}, \Phi^{**})$.

5 SUMMARY OF THE ESSAYS

This dissertation includes four essays, with the author's contributions to each essay detailed in Table 1.

Table 1. Author's contributions across four essays include in the dissertation.

	Essays and Authors			
Area of Contribution	**1**	**2**	**3**	**4**
Defining the research design	A1, A2	A1, A2	A1, A2	A1, A2
Collecting and managing research data	A1	A1, A2	A1	A1
Data analysis	A1, A2	A1, A2	A1	A1
Methodology and research methods	A1, A2	A1, A2	A1, A2	A1, A2
Acquisition of research resources	A1, A2	A1, A2	A1, A2	A1, A2
Programming	A1, A2	A1, A2	A1	A1, A2
Guidance in the research process	A2, A3	A2	A2	A2, A3
Verification and analysis of findings	A1, A2	A1, A2	A1, A2	A1, A2
Visualization of results	A1, A2, A3	A1, A2, A4	A1, A2	A1, A2, A3
Writing the body text of the original article	A1, A2	A1, A2	A1, A2	A1, A2
Editing the article at different stages	A1, A2, A3	A1, A2, A4	A1, A2	A1, A2, A3

Note:

Author 1 (A1): Masoumeh Fathi, Doctoral Student

Author 2 (A2): Klaus Grobys, Associate Professor of Finance {First Supervisor}

Author 3 (A3): Janne Äijö, Professor of Finance {Second Supervisor}

Author 4 (A4): James W. Kolari, Professor of Finance

Essay 1: A common component of Fama and French factor variances.

Essay 2: On the realized risk of foreign exchange rates: a fractal perspective.

Essay 3: Modeling variance risk in financial markets using power-laws: new evidence from the Garman-Klass variance estimator.

Essay 4: Momentum bubbles; a new perspective derived from log-periodicity.

5.1 A common component of Fama and French factor variances

The first essay of this dissertation examines the risk of the Fama and French equity factors through their realized variances. Despite the widespread use of Fama and French factor models in financial analysis, an increasing body of research has highlighted their diminishing significance and reliability in explaining asset returns (Dichev, 1998; Chan et al., 2000; Hirshleifer, 2001; Schwert, 2003; Fama & French, 2015; Grobys & Kolari, 2022; Cakici et al., 2023). While earlier studies primarily focused on the reliability of factor premiums, this paper is the first to investigate the risk associated with well-established Fama and French equity factors by analyzing their realized variances. Thus, the aim of this study is to evaluate how these variances influence the explanatory power of the factors over time. Thereby, the study hypothesizes that the realized variances of these factors are too high to permit convergence in finite samples, resulting in unreliable *t*-statistics and consequently limiting the statistical significance of factor premiums.

The study utilizes daily observations of six Fama and French factors include excess market factor (MKT), size factor (SMB), value factor (HML), investment factor (CMA), profitability factor (RMW), and momentum factor (MOM) to compute their monthly realized variances calculated as the sum of squared deviations of daily returns from their monthly means, following Merton (1980). The power-law models were fitted using the MLE, consistent with Clauset et al. (2009). The results of this study show that realized factor variances indeed exhibit strong power-laws behavior with tail exponents falling within the range $2 < \hat{\alpha} < 3$. In addition, the results of the goodness-of-fit test for the power-laws null hypothesis indicate that power laws are indeed plausible distributions for modeling realized factor variances. Next, the study investigates whether the theoretical means of realized factor variances are statistically defined, given that the estimated power-law exponents are consistently estimated being close to $\alpha \approx 2$ for all realized factor variances. However, before conducting this test, the study addresses potential issues with standard errors, which are typically underestimated due to assumptions of independence, and employs blocks bootstraps (Grobys, 2023a; Grobys & Juntilla, 2021) to account for volatility clustering and unknown dependence structures. The results indicate that once unknown dependency structures (i.e., volatility clustering) are considered, the standard deviations are substantially larger than those calculated under *i.i.d.* assumptions. Despite these adjustments, the bootstrapped means of power-law exponents remain close to the original estimates, confirming that the estimates obtained via Clauset et al.'s (2009) MLE approach are unbiased.

Next, using the robust standard errors derived from block bootstrapping, the study tests the infinite theoretical variance hypothesis, with the null hypothesis (H_0): $\alpha \leq 2$ indicating infinite theoretical variance. The results show that the hypothesis $\alpha \approx 2$ cannot be rejected for any realized factor variances, indicating a notable commonality across all variances. These findings contrast with the existing literature discussed in Lux and Alfarano (2016) and align with Mandelbrot's (1963) infinite theoretical variance hypothesis. Additionally, goodness-of-fit tests provide strong evidence against log-normal and exponential distributions as the underlying distributions for generating realized factor variances. These findings strongly contrast with earlier research on realized volatility, which suggests that realized asset volatility is generally log-normally distributed (e.g., Andersen et al., 2001; Andersen et al., 2001a, 2001b).

To further validate the results, a novel joint test was performed to evaluate whether the Fama and French factor variances share a common component, represented by a common cross-sectional power-law exponent. The results indicate that the null hypothesis of a universal power-law exponent cannot be rejected for $1.7 < \alpha < 2.5$, providing strong evidence of a common source of risk across the realized factor variances. These findings emphasize the presence of a common power-laws behavior in the Fama and French factor variances, highlighting a shared component in their risk dynamics. The study also employed one-sigma tests from Grobys (2023a), which adopt a different approach compared to the method introduced by Clauset et al. (2009). While Clauset et al. (2009) assume the power-law model as the null hypothesis, Grobys' method considers the power-law model as the alternative, allowing better control over type-1 errors. The results from one-sigma tests strongly align with and confirm the earlier findings derived from Clauset et al.'s (2009) approach.

Additional evidence indicates that the dynamics of realized factor variances can be explained using Grobys' (2023b) recently developed multifractal model of asset invariances. This model, which successfully generates factor variances for all Fama and French equity factors, employs a multiplicative cascade framework based on binomial bending with a probability of $p = 0.70$. Consequently, the most significant finding of the paper is that realized factor variances exhibit a consistent power-law exponent, approximately $\alpha \approx 2$, highlighting a fundamental commonality across these variances. Further, the study examines the co-dependencies among realized factor variances through Grobys'(2023c) recently proposed concept of co-fractality. The results show that, for 11 out of 15 realized factor variance pairs, the co-fractality parameter is estimated at $\hat{\lambda}^{CF} > 0.80$, indicating a significant degree of (weak) co-fractality among the realized factor variances. This implies that risk diversification opportunities are limited, because extreme events occurring in the cross-section of

realized variances tend to coincide. Finally, additional analyses, including a sample split test and tests to address autocorrelation in realized variances and discretization errors, were conducted to further validate the findings of the study, with all the tests consistently confirming the robustness of the previous results.

From a theoretical perspective, this study provides evidence that both documented replication failures and inconsistent methodological choices in finance literature (Hou et al., 2020; Dick-Nielsen et al., 2023) may stem from undefined risk associated with factor variances, because, as pointed out by Fama (1963), statistical conclusions drawn from OLS or other methods that depend on the existence of variance are likely to result in misleading inferences. The paper discusses that data with extremely fat tails, as observed in the Fama and French factor variances, include extreme values that can heavily impact the outcomes of statistical analysis. In this context, the findings highlight the need for a thorough re-evaluation of the econometric methods employed in financial decision-making. From a practical perspective, the findings suggest that the factor investing industry experiences a significantly higher level of extreme risk than previously assumed, as the power-law exponent extrapolates low-probability deviations not observed in the data.

5.2 On the realized risk of foreign exchange rates: a fractal perspective

The second essay of this dissertation examines the risk associated with FX rates of G10 currencies. Uncertainty in the FX market has received extensive attention in academic research, primarily because its market capitalization is significantly larger than that of other financial markets. While traditional models like GARCH have been widely applied to study FX market volatility (e.g., Jorion 1995; Baillie and Bollerslev 1991; Bollerslev and Melvin 1994), they often fail to adequately capture the heavy tails and extreme events observed in financial data. In addition, although Andersen et al. (2003, 2004) showed that reduced-form time series models for realized volatility perform better than GARCH-type models in forecasting future volatility, other researchers have highlighted challenges in using realized variances to model FX risk (e.g., Wang & Yang, 2009; Corsi, 2004; Andersen et al., 2005; Bubák et al., 2011; Corsi et al., 2008).

To address these limitations and enhance our understanding of extreme events arriving in FX markets, the second essay adopts a fractal perspective and employs power laws to model realized FX variances pairs, including AUD/USD, EUR/USD, GBP/USD, NOK/USD, NZD/USD, USD/CAD, USD/CHF, USD/JPY, and USD/SEK. This study explores the fractal-like properties of realized FX variances by extending the

work of Grobys (2023c), which focused on daily realized FX variances. If realized FX variances exhibit fractal-like properties, it is expected that the estimates derived from weekly data will closely align with those documented for daily data in Grobys (2023c).

The findings indicate that power-law models provide a plausible fit for the realized variance of G10 currency pairs. The results indicate that six out of nine realized FX variances have tail exponents in the range of $2 < \hat{\alpha} < 3$, implying that the variance of these variances is statistically undefined. Further analyses using the goodness-of-fit tests from Clauset et al. (2009), show that power-law models outperform log-normal and exponential models in most cases, emphasizing the presence of heavy tails indicating higher risk than previously assumed for the FX market. Additionally, the sample-split analysis confirms the previous results, showing consistent power-laws behavior with $2 < \hat{\alpha} < 3$ for all realized variance series across different economic periods. Moreover, the close alignment between the results from weekly data in this study and the findings derived from daily data in Grobys (2023c) provides strong evidence for fractal-like behavior in realized FX variances, consistent with Mandelbrot's (2008) proposition that power-laws behavior is manifested in invariance of the power-law exponent across time frequencies. The paper concludes that in the presence of undefined variance for most FX rate realized variances, statistical tests relying on finite variance may yield results that depend on the chosen time period or frequency, leading to inconsistent outcomes when analyzing extreme events in FX markets.

5.3 Modeling variance risk in financial markets using power-laws: new evidence from the Garman-Klass variance estimator

The third essay of this dissertation examines the range-based variance risk across five diverse asset markets: the S&P 500, gold, crude oil, the USD/GBP exchange rate, and Bitcoin. Utilizing the Garman-Klass variance estimator (Garman & Klass 1980), the study models these variances within a power-laws framework to capture extreme market risks and aims to address a critical question: does the population mean exist for range-based variances in financial markets? Range-based variance is preferred because, unlike model-based residual variance, it is directly observable and independent of model assumptions. If the distribution of range-based variances displays extremely heavy tails, it suggests that the variance estimates may be unstable and potentially lack finite moments. As a result, if input variances are not well-defined, the residual variance in models like CAPM or Fama-French may also be undefined, undermining the validity of classical inference tools such as standard

errors, *t*-statistics, and *R*-squared. Building on Grobys' (2021) work, which used the Parkinson variance estimator, this study adopts the Garman-Klass variance estimator to account for noise in high-frequency data (Molnár, 2012) by incorporating high, low, opening, and closing prices.

The results show that Garman-Klass range-based variances exhibit heavy tails across all five asset markets, closely aligning with power-law distributions. Furthermore, power-law fitting using MLE confirmed that the tail exponents for all assets fall within the range $2 < \hat{\alpha} < 3$, suggesting that the variances of variances are infinite. As Mandelbrot (1967) pointed out, such variance can be considered "*so large that it may in practice be assumed infinite*," challenging the core assumption of finite variances in conventional econometric models and exposing the limitations of statistical methods like *t*-tests, which are derived from the OLS regression framework. The paper discusses that the evidence of sample dependency, as noted in Fama's (1963) analysis of Mandelbrot's infinite variance hypothesis, underscores the limitations of moment-dependent statistical techniques in addressing the heavy-tailed distributions commonly observed in financial markets. Further, goodness-of-fit tests strongly support power-law distributions as plausible models for the Garman-Klass range-based variances of these assets. Additionally, these tests consistently favor power-law models over log-normal and exponential alternatives, emphasizing their suitability for modeling extreme financial risks. These findings contrast with the literature suggesting that realized asset volatility follows a log-normal distribution (e.g., Andersen et al., 2001; Andersen et al., 2001a, 2001b), but they align with studies proposing that a Pareto distribution better describes volatility in financial markets (Renò & Rizza, 2003; Grobys, 2021, 2023c, 2024).

In addition, a joint test identifies a common power-law exponent ($\alpha \approx 2.8$) across all five unrelated asset markets, indicating a shared risk dynamic likely driven by systemic or behavioral factors such as herding. Specifically, the null hypothesis of a shared exponent cannot be rejected for the range $2.5 < q < 3.1$, indicating a unified component affecting variance risks across the S&P 500, gold, crude oil, USD/GBP exchange rate, and Bitcoin. Moreover, to account for the potential impact of autocorrelation on tail index estimates, autoregressive models of order p were applied to the Garman-Klass range-based variances for each asset. The results showed significant autocorrelation patterns, and even after adjusting for higher-order autocorrelation, the innovation processes of annualized daily range-based variances exhibited extremely heavy tails. The power-law exponents estimated from the absolute values of the innovation processes closely align with those observed in the original data, providing strong evidence of genuine power-laws behavior. In addition, Clauset et al.'s (2009) goodness-of-fit test is applied to the absolute values of the innovation processes, confirming the persistence of power-laws behavior for

all assets except gold, thereby strengthening the validity of the power-law estimates for realized variances. The goodness-of-fit test results, with zero *p*-values in all cases, reject log-normal and exponential distributions, supporting the power-law distribution. In addition, other robustness tests—including a sample-split test, replication of the analysis using realized variance, alternative estimation of power-law exponents via a Bayesian inference approach instead of the KS-based method, using an alternative model comparison metric such as the Akaike Information Criterion (AIC), rolling window estimations, and extending the analysis to government bond markets—confirmed the previous results and reinforced the reliability and generalizability of the findings.

As a concluding remark, the paper emphasizes that effectively infinite variances and the rejection of log-normality highlight the limitations of traditional tools, which can significantly underestimate tail risks in extreme market environments. In addition, consistent with Grobys (2021), the study suggests that replication failures in financial research may result from the infinite variance of variances, which reduces the reliability of traditional statistical methods.

5.4 Momentum bubbles: A new perspective derived from log-periodicity

The fourth essay of this dissertation examines the predictability of momentum bubbles utilizing the LPPLS model. It explores LPPLS signatures prior to major momentum downturns and evaluates how the model's predictive accuracy changes as bubbles approach their peak. Although momentum strategies are replicable in expanded samples (Hou et al., 2020), they are prone to bubbles and crashes. For example, Daniel and Moskowitz (2016) documented extreme market dynamics between July and August of 1932 and March and May of 2009. Due to the momentum strategy's popularity and ability to capture market trends and deliver risk-adjusted returns, this study explores whether momentum bubbles can be predicted to facilitate more informed decision-making in momentum strategies. To achieve this, the study applies the LPPLS model introduced by Sornette et al. (1996), which is particularly well-suited for analyzing financial bubbles as it is derived from the theoretical framework of imitation or herding behavior among traders.

This study utilizes daily observations from ten portfolios formed using the momentum strategy, covering the period from 1926 to 2023.The analysis focuses on compounded returns from the equity momentum strategy involving a long position in the winner decile and a short position in the loser decile. Next, using the Epsilon Draw-down Method outlined by Johansen & Sornette (1998, 2001, 2010) and

focusing on bubbles with a minimum interval of four years between them, the study identifies five bubbles with peaks on December 21, 1929; March 10, 2000; July 14, 2008; January 20, 2016; and April 16, 2020, as potential candidates for LPPLS model predictions. Subsequently, the LPPLS model's prediction accuracy is analyzed using a three-stage estimation approach, following the method proposed by Grobys (2023), with calibration windows ending 3, 6, 9, and 12 months prior to each bubble peak to evaluate whether accuracy improves as the peak approaches. The three-stage approach allows for a gradual improvement of the model parameters, progressing from a simple approximation to a more complex and accurate representation of the underlying dynamics captured by the LPPLS model. Finally, to test the robustness of the results, the residuals are analyzed using the Augmented Dickey–Fuller (ADF) test, following Lin et al. (2014), who proposed that residuals from the LPPLS model should exhibit stationary, mean-reverting behavior. This is because non-stationary residuals may indicate a spurious regression, thereby affecting the clarity of statistical analyses.

The results of this study validate the presence of LPPLS signatures preceding each of the five bubbles analyzed, consistent with the findings for broad market indices, as documented in Sornette (2017). By applying the LPPLS model and evaluating prediction accuracy at 3, 6, 9, and 12 months intervals before each bubble peak, these signatures are shown to effectively serve as early indicators of future bubbles. The findings indicate that while the prediction accuracy of the LPPLS model does not consistently improve as the bubble approaches its peak, there is still a partial increase in accuracy. These findings align with previous literature suggesting that the LPPLS model's prediction accuracy improves as the critical time approaches (Sornette, 2017), although the results do not fully confirm this trend in all cases. Additionally, a cross-sectional comparison of predictions for all five bubbles shows that the LPPLS model tends to provide more accurate predictions when using data from 9 months prior to the critical point compared to data from 3, 6, or 12 months beforehand. Moreover, the robustness test indicates that the LPPLS model's residuals follow a stationary mean-reverting process, confirming the statistical significance of the LPPLS signature and strengthening confidence in the model's predictive capabilities.

The results of this study confirm the probabilistic rather than the deterministic nature of the LPPLS model, as emphasized by Sornette (2017), and acknowledge the model's sensitivity to parameter selection, noted by Zhou and Sornette (2006) and Brée et al. (2013). It also aligns with previous research on the predictability of momentum bubbles and crashes (e.g., Daniel & Moskowitz, 2016; Barroso & Santa-Clara, 2015; Chabot et al., 2014), which found these phenomena to be predictable, particularly when momentum performs well or interest rates are low. Contributing to this body of research, this study offers new insights by showing that the LPPLS

model can detect its signatures before the onset of momentum-driven bubble bursts, thereby confirming its capability to capture pre-peak dynamics. In addition, by analyzing the prediction accuracy using calibration windows ending 3, 6, 9, and 12 months before each bubble peak, the study provides a comprehensive view of the model's performance, aligning with previous literature that acknowledges the inherent limitations and challenges of forecasting complex dynamic systems (Sornette & Johansen, 2001; Cont, 2001; Sornette, 2003).

The study also offers a behavioral explanation for the occurrence of momentum bubbles, linking them to herding behavior among traders. The LPPLS model is based on the theoretical framework that emphasizes imitation dynamics in financial markets. Therefore, the presence of significant LPPLS signatures in the log-compounded return of the momentum factor suggests that momentum bubbles are not purely random events but rather the result of herding behavior. Essentially, individual traders tend to follow the herd instead of making independent decisions. Such behavior can result in the formation of speculative bubbles, which may lead to crashes when they burst. The study argues that the optionality effect identified in Daniel and Moskowitz's (2016) study has a behavioral cause rooted in imitation among traders. Furthermore, this study provides practical implications for portfolio management and risk assessment. It argues that by integrating the LPPLS model with momentum strategies, decision-making can be enhanced, allowing for timely adjustments to portfolios in anticipation of possible reversals.

References

Andersen, T. G., Bollerslev, T., & Diebold, F. X. (2005). *Roughing it up: Including jump components in the measurement, modelling, and forecasting of return volatility* (NBER Working Paper No. 11775). National Bureau of Economic Research.

Andersen, T. G., Bollerslev, T., & Meddahi, N. (2004). Analytical evaluation of volatility forecasts. *International Economic Review, 45*(4), 1079–1110.

Andersen, T. G., Bollerslev, T., Diebold, F. X., & Ebens, H. (2001). The distribution of realized stock return volatility. *Journal of Financial Economics, 61*(1), 43–76.

Andersen, T. G., Bollerslev, T., Diebold, F. X., & Labys, P. (2001a). Modeling and forecasting realized volatility. *Econometrica, 71*(2), 579–625.

Andersen, T. G., Bollerslev, T., Diebold, F. X., & Labys, P. (2001b). The distribution of realized exchange rate volatility. *Journal of the American Statistical Association, 96*(453), 42–55.

Andersen, T. G., Bollerslev, T., Diebold, F. X., & Labys, P. (2003). Modeling and forecasting realized volatility. *Econometrica, 71*(2), 579–625.

Andriani, P., & McKelvey, B. (2009). *From Gaussian to Paretian thinking: Causes and implications of power laws in organizations*. *Organization Science, 20*(6), 1053–1071.

Avery, C., & Zemsky, P. (1998). Multidimensional uncertainty and herd behavior in financial markets. *American Economic Review, 88*(4), 724–748.

Bachelier, L. (1900). Théorie de la spéculation. *Annales Scientifiques de l'École Normale Supérieure, 17*(3), 21–86.

Baillie, R. T., & Bollerslev, T. (1991). Intra-day and inter-market volatility in foreign exchange rates. *Review of Economic Studies, 58*(3), 565–585.

Bak, P. (1996). *How nature works: The science of self-organized criticality*. Springer-Verlag.

Barndorff-Nielsen, O. E., & Shephard, N. (2004). Power and bipower variation with stochastic volatility and jumps. *Journal of Financial Econometrics*, *2*(1), 1–37.

Barndorff-Nielsen, O. E., & Shephard, N. (2002). Econometric analysis of realized volatility and its use in estimating stochastic volatility models. *Journal of the Royal Statistical Society: Series B (Statistical Methodology), 64*(2), 253–280.

Barroso, P., & Santa-Clara, P. (2015). Momentum has its moments. *Journal of Financial Economics, 116*, 111–120.

Black, F., & Scholes, M. (1973). The pricing of options and corporate liabilities. *Journal of Political Economy*, *81*(3), 637–654.

Bollerslev, T., & Melvin, M. (1994). Bid–ask spreads and volatility in the foreign exchange market: An empirical analysis. *Journal of International Economics, 36*, 355–372.

Brée, D. S., Challet, D., & Peirano, P. P. (2013). Prediction accuracy and sloppiness of log-periodic functions. *Quantitative Finance, 13*(2), 275–280.

Brown, J. H., Gupta, V. K., Li, B. L., Milne, B. T., Restrepo, C., & West, G. B. (2002). The fractal nature of nature: Power laws, ecological complexity, and biodiversity. *Philosophical Transactions of the Royal Society of London. Series B: Biological Sciences, 357*(1421), 619–626.

Bubák, V., Kočenda, E., & Žikeš, F. (2011). Volatility transmission in emerging European foreign exchange markets. *Journal of Banking and Finance, 35*, 2829–2841.

Cakici, N., Fieberg, C., Metko, D., & Zaremba, A. (2023). Do anomalies really predict market returns? New data and new evidence. *Review of Finance.* Advance online publication.

Calvet, L. E., & Fisher, A. J. (2004). Regime-switching and the estimation of multifractal processes. *Journal of Financial Econometrics, 2*(1), 44–83.

Chabot, B., Ghysels, E., & Jagannathan, R. (2014). *Momentum trading, return chasing, and predictable crashes* (Working Paper No. 20660). National Bureau of Economic Research.

Chan, L. K. C., Karceski, J., & Lakonishok, J. (2000). New paradigm or same old hype in equity investing? *Financial Analyst Journal, 56*(4), 23–36.

Christensen, B. J., & Prabhala, N. R. (1998). The relation between implied and realized volatility. *Journal of Financial Economics, 50*(2), 125–150.

Clauset, A., Shalizi, C. R., & Newman, M. E. J. (2009). Power-law distributions in empirical data. *SIAM Review, 51*(4), 661–703.

Cont, R. (2001). Empirical properties of asset returns: Stylized facts and statistical issues. *Quantitative Finance, 1*(2), 223–236.

Corsi, F. (2004). A simple long memory model of realized volatility. *Capital Markets: Asset Pricing & Valuation eJournal.*

Corsi, F., Mittnik, S., Pigorsch, C., & Pigorsch, U. (2008). The volatility of realized volatility. *Econometric Reviews, 27*(1–3), 46–78.

Daniel, K., & Moskowitz, T. J. (2016). Momentum crashes. *Journal of Financial Economics, 122*(2), 221–247.

Deluca, A., & Corral, Á. (2013). Fitting and goodness-of-fit test of non-truncated and truncated power-law distributions. *Acta Geophysica, 61*(6), 1351–1394

Dichev, I. D. (1998). Is the risk of bankruptcy a systematic risk? *Journal of Finance, 53*(3), 1131–1147.

Dick-Nielsen, J., Feldhütter, P., Pedersen, L. H., & Stolborg, C. (2023). *Corporate bond factors: Replication failures and a new framework*. SSRN.

Drees, H., & Kaufmann, E. (1998). Selecting the optimal sample fraction in univariate extreme value estimation. *Stochastic Processes and their Applications, 75*(2), 149–172.

Embrechts, P., Klüppelberg, C., & Mikosch, T. (1997). *Modelling extremal events: For insurance and finance* (1st ed.). Springer.

Engle, R. F. (1982). Autoregressive conditional heteroskedasticity with estimates of the variance of United Kingdom inflation. *Econometrica*, *50*(4), 987–1007.

Fama, E. F. (1963). Mandelbrot and the stable Paretian hypothesis. *Journal of Business, 36*(4), 420–429.

Fama, E. F. (1964). *The distribution of the daily differences of the logarithms of stock prices* (Doctoral dissertation, University of Chicago). University of Chicago.

Fama, E. F. (1965). Random walks in stock market prices. *Financial Analysts Journal, 21*(5), 55–59.

Fama, E. F., & French, K. R. (1992). The cross-section of expected stock returns. *Journal of Finance, 47*(2), 427–465.

Fama, E. F., & French, K. R. (1993). Common factors in the returns on stocks and bonds. *Journal of Financial Economics, 33*(1), 3–56.

Fama, E. F., & French, K. R. (2015). A five-factor asset pricing model. *Journal of Financial Economics, 116*(1), 1–22.

Fama, E. F., & French, K. R. (2017). International tests of a five-factor asset pricing model. *Journal of Financial Economics, 123*(3), 441–463.

Fama, E. F., & French, K. R. (2018). Choosing factors. *Journal of Financial Economics, 128*(2), 234–252.

Fama, E. F., & French, K. R. (2020). *The value premium* (Fama-Miller Working Paper No. 20-01). SSRN.

Fathi, M., Grobys, K., & Kolari, J. W. (2024). On the realized risk of foreign exchange rates: A fractal perspective. *Journal of Risk and Financial Management*, *17*(2), 79.

Gabaix, X., Gopikrishnan, P., Plerou, V., & Stanley, H. E. (2003). A theory of power-law distributions in financial market fluctuations. *Nature, 423*(6937), 267–270.

Garman, M. B., & Klass, M. J. (1980). On the estimation of security price volatilities from historical data. *The Journal of Business, 53*(1), 67–78.

Geraskin, P., & Fantazzini, D. (2013). Everything you always wanted to know about log-periodic power laws for bubble modeling but were afraid to ask. *The European Journal of Finance, 19*(5), 366–391.

Ghysels, E., Harvey, A. C., & Renault, E. (1996). Stochastic volatility. In G. S. Maddala & C. R. Rao (Eds.), *Handbook of statistics* (Vol. 14, pp. 119–191). Elsevier.

Goldstein, M. L., Morris, S. A., & Yen, G. G. (2004). Problems with fitting to the power-law distribution. *The European Physical Journal B, 41*(2), 255–258.

Grobys, K. (2021). What do we know about the second moment of financial markets? *International Review of Financial Analysis, 78,* Article 101891.

Grobys, K. (2023). A finite-time singularity in the dynamics of the US equity market: Will the US equity market eventually collapse? *International Review of Financial Analysis, 89,* Article 102787.

Grobys, K. (2023a). A fractal and comparative view of the memory of bitcoin and S&P 500 returns. *Research in International Business and Finance, 66,* Article 102021.

Grobys, K. (2023b). A multifractal model of asset (in)variances. *Journal of International Financial Markets, Institutions and Money, 85,* Article 101767.

Grobys, K. (2023c). Correlation versus co-fractality: Evidence from foreign exchange rate variances. *International Review of Financial Analysis, 86,* Article 102531.

Grobys, K. (2024). A universal exponent governing foreign exchange rate risks. *International Review of Financial Analysis, 95*(Part B), Article 103422.

Grobys, K., & Juntilla, J.-P. (2021). Speculation and lottery-like demand in cryptocurrency markets. *Journal of International Financial Markets, Institutions and Money, 71,* Article 101289.

Grobys, K., & Kolari, J. W. (2022). Choosing factors: The international evidence. *Applied Economics, 54,* 633–647.

Hamermesh, D. S. (2007). Viewpoint: Replication in economics. *Canadian Journal of Economics, 40*(3), 715–733.

Hamilton, J. D. (1989). A new approach to the economic analysis of nonstationary time series and the business cycle. *Econometrica, 57*(2), 357–384.

Heston, S. L. (1993). A closed-form solution for options with stochastic volatility with applications to bond and currency options. *The Review of Financial Studies*, *6*(2), 327–343.

Hirshleifer, D. (2001). Investor psychology and asset pricing. *Journal of Finance, 56*(4), 1533–1597.

Hou, K., Xue, C., & Zhang, L. (2020). Replicating anomalies. *Review of Financial Studies, 33*(5), 2019–2133.

Johansen, A., & Sornette, D. (1998). Stock market crashes are outliers. *The European Physical Journal B - Condensed Matter and Complex Systems, 1*(2), 141–143.

Johansen, A., & Sornette, D. (2001). Large stock market price drawdowns are outliers. *Journal of Risk, 4*(2), 69–110.

Johansen, A., & Sornette, D. (2010). Shocks, crashes, and bubbles in financial markets. *Brussels Economic Review (Cahiers Economiques de Bruxelles, 53,* 201–253.

Johansen, A., Ledoit, O., & Sornette, D. (2000). Crashes as critical points. *International Journal of Theoretical and Applied Finance, 3*(2), 219–255.

Johansen, A., Sornette, D., & Ledoit, O. (1999). Predicting financial crashes using discrete scale invariance. *Journal of Risk, 1*(4), 5–32.

Jorion, P. (1995). Predicting volatility in the foreign exchange market. *Journal of Finance, 50*(2), 507–528.

Latane, H. A., & Rendleman, R. J. (1976). Standard deviations of stock price ratios implied in option prices. *Journal of Finance*, *31*(2), 369–381.

Lim, A. E. B., Shanthikumar, J. G., & Vahn, G.-Y. (2011). *Conditional value-at-risk in portfolio optimization: Coherent but fragile*. *Operations Research Letters, 39*(3), 163–171.

Lin, L., Ren, R. E., & Sornette, D. (2014). The volatility-confined LPPL model: A consistent model of 'explosive' financial bubbles with mean-reverting residuals. *International Review of Financial Analysis, 33,* 210–225.

Lux, T. (1995). Herd behaviour, bubbles and crashes. *The Economic Journal, 105*(431), 881–896.

Lux, T., & Alfarano, S. (2016). Financial power-laws: Empirical evidence, models, and mechanisms. *Chaos, Solitons and Fractals, 88,* 3–18.

Lux, T., & Marchesi, M. (1999). Scaling and criticality in a stochastic multi-agent model of a financial market. *Nature, 397*(6719), 498–500.

Lux, T., Morales-Arias, L., & Sattarhoff, C. (2014). A Markov-switching multifractal approach to forecasting realized volatility. *Journal of Forecasting, 33*(7), 532–541.

Mandelbrot, B. (1963). The variation of certain speculative prices. *Journal of Business, 36*(4), 394–419.

Mandelbrot, B. (1967). *The variation of some other speculative prices*. *The Journal of Business, 40*(4), 393–413.

Mandelbrot, B. (1982). *The fractal geometry of nature*. W.H. Freeman and Company.

Mandelbrot, B. B. (2008). *The (mis)behavior of markets: A fractal view of risk, ruin, and reward*. Profile Books.

Mantegna, R. N., & Stanley, H. E. (1995). Scaling in financial markets. *Nature, 376*(6535), 46–49.

Markowitz, H. (1952). Portfolio selection. *The Journal of Finance, 7*(1), 77–91.

Merton, R. C. (1980). On estimating the expected return on the market: An exploratory investigation. *Journal of Financial Economics, 8*(4), 323–361.

Molnár, P. (2012). Properties of range-based volatility estimators. *International Review of Financial Analysis, 23*, 20–29.

Newman, M. E. J. (2005). *Power laws, Pareto distributions, and Zipf's law*. *Contemporary Physics, 46*(5), 323–351.

Pareto, V. (1896). *Cours d'économie politique: Professé à l'Université de Lausanne* (Vol. 1). F. Rouge.

Parkinson, M. (1980). The extreme value method for estimating the variance of the rate of return. *The Journal of Business, 53*(1), 61–65.

Renò, R., & Rizza, R. (2003). Is volatility lognormal? Evidence from Italian futures. *Physica A: Statistical Mechanics and Its Applications, 322*, 620–628.

Rockafellar, R. T., & Uryasev, S. (2000). Optimization of conditional value-at-risk. *Journal of Risk*, *2*, 21–41.

Rockafellar, R. T., & Uryasev, S. (2002). Conditional value-at-risk for general loss distributions. *Journal of Banking & Finance*, *26*(7), 1443–1471.

Salazar, B. (2016). Mandelbrot, Fama and the emergence of econophysics. *Cuadernos de Economía, 35*(69), 637–662.

Schwert, G. W. (2003). Anomalies and market efficiency. In G. M. Constantinides, M. Harris, & R. M. Stulz (Eds.), *Handbook of the economics of finance* (Vol. 1, pp. 939–974). North-Holland.

Serra-Garcia, M., & Gneezy, U. (2021). Nonreplicable publications are cited more than replicable ones. *Science Advances, 7*(5), eabd1705.

Sharpe, W. F. (1964). Capital asset prices: A theory of market equilibrium under conditions of risk. *The Journal of Finance, 19*(3), 425–442.

Shephard, N. (Ed.). (2005). *Stochastic volatility: Selected readings*. Oxford University Press.

Sornette, D. (2003). Critical market crashes. *Physics Reports, 378*(1), 1–98.

Sornette, D. (2017). *Why stock markets crash: Critical events in complex financial systems*. Princeton University Press.

Sornette, D., & Cauwels, P. (2014). Financial bubbles: Mechanisms and diagnostics. *Swiss Finance Institute Research Paper No. 14–28.*

Sornette, D., & Johansen, A. (2001). Large financial crashes. *Physica A: Statistical Mechanics and its Applications, 299*(1–2), 40–59.

Sornette, D., Johansen, A., & Bouchaud, J. P. (1996). Stock market crashes, precursors, and replicas. *Journal de Physique I, 6*(1), 167–175.

Taleb, N. N. (2007). *The black swan: The impact of the highly improbable*. Random House Publishing Group.

Taleb, N. N. (2020). *Statistical consequences of fat tails: Real world preasymptotics, epistemology, and applications. Papers and commentary.* STEM Academic Press.

Taylor, S. J. (1982). Financial returns modelled by the product of two stochastic processes—A study of daily sugar prices 1961–79. In O. D. Anderson (Ed.), *Time series analysis: Theory and practice 1* (pp. 203–226). North-Holland.

Vandewalle, N., Ausloos, M., Boveroux, P., & Minguet, A. (1999). Visualizing the log-periodic pattern before crashes. *European Physics Journal B, 9,* 355–359.

Wang, J., & Yang, M. (2009). Asymmetric volatility in the foreign exchange markets. *Journal of International Financial Markets, Institutions and Money, 19*(4), 597–615.

West, G. (2017). *Scale: The universal laws of growth, innovation, sustainability, and the pace of life in organisms, cities, economies, and companies*. Penguin Press.

Wheatley, S., Sornette, D., Huber, T., Reppen, M., & Gantner, R. N. (2018). Are bitcoin bubbles predictable? Combining a generalized Metcalfe's law and the LPPLS model. *Swiss Finance Institute Research Paper No. 18–22.*

Zhou, W.-X., & Sornette, D. (2006). Is there a real estate bubble in the US? *Physica A: Statistical Mechanics and Its Applications, 361*(1), 297–308.

North American Journal of Economics and Finance 75 (2025) 102292

Contents lists available at ScienceDirect

North American Journal of Economics and Finance

journal homepage: www.elsevier.com/locate/najef

A common component of Fama and French factor variances☆

Masoumeh Fathi [a,*], Klaus Grobys [a,b], Janne Äijö [a]

[a] Finance Research Group, School of Accounting and Finance, University of Vaasa, Wolffintie 34, 65200 Vaasa, Finland
[b] Innovation and Entrepreneurship (InnoLab), University of Vaasa, Wolffintie 34, 65200 Vaasa, Finland

ARTICLE INFO

JEL codes:
C02
C15
C18
G10
G17
G18

Keywords:
Co-fractality
Fama and French factors
Power laws
Realized variance
Risk

ABSTRACT

This is the first study that explicitly explores the risk of the Fama and French equity factors in terms of their realized variances. Our results show that realized factor variances exhibit strong power-law behavior. A striking commonality is that the power-law exponents are close to $\alpha \approx 2$ regardless of which factor variance is analyzed. Notably, our novel joint test designed to test Mandelbrot's infinite variance hypothesis in the cross-section of realized factor variances shows that the null hypothesis of $\alpha = 1.9$ cannot be rejected, which further corroborates the evidence that (a) there exist a common component governing factor variance risk, and (b) factor variance risk is statistically undefined. Further evidence derived from co-fractality analysis shows that (c) risk diversification appears to be very limited as factor variances tend to exhibit power-law behavior coincidently. We argue that our study has several theoretical and practical implications—especially due to the fact that factor investing reached $5 trillion in assets under management.

1. Introduction

In their well-known and often-cited 1992 paper, published in the well-respected *Journal of Finance*, Fama and French (1992, p. 464) declared the death of the capital asset pricing model (CAPM) by making the following bold statement:

> We are forced to conclude that the SLB model does not describe the last 50 years of average stock returns.[1] (Fama and French, 1992, p. 464)

In their paper, Fama and French (1992) provided strong evidence for a negative relation between size and average return and a positive relation between book-to-market equity and average return. Based on their findings, Fama and French (1993) proposed a three-factor model consisting of the excess market factor, a size factor that is a zero-cost portfolio long on small stocks and short on big stocks, and a value factor that is a zero-cost portfolio long on high book-to-market ratio stocks and short on low high book-to-market ratio stocks. This study has been cited more than 30,000 times and was considered 'groundbreaking research'.

☆ Valuable comments were received from Sami Vähämaa and Anupam Dutta at the 2023 Finance Research Seminar at the University of Vaasa. Further valuable comments were received from Tommi Sottinen, Seppo Hassi and Seppo Pynnönen at the 2023 Mathematics & Statistics Research Seminar at the University of Vaasa. In addition, insightful comments were received from Petri Jylhä and Samuli Knüpfer at the 2023 Winter Workshop in Finance of the Graduate School of Finance (GSF) and the Finnish Doctoral Program in Economics (FDPE).

* Corresponding author.
E-mail addresses: masoumeh.fathi@uwasa.fi (M. Fathi), klaus.grobys@uwasa.fi (K. Grobys), janne.aijo@uwasa.fi (J. Äijö).

[1] Note that this model has been widely-used since the 1960s and was proposed in highly-influential works by Sharpe (1964), Treynor (1962), Lintner (1965a,b) and Mossin (1966).

https://doi.org/10.1016/j.najef.2024.102292
Received 5 February 2024; Received in revised form 20 May 2024; Accepted 20 September 2024
Available online 24 September 2024

In 2015, Fama and French (2015) proposed an expanded asset pricing model consisting of five factors; that is, a profitability factor and an investment factor on top of the previous three equity factors. Unsurprisingly, linear factor models such as the capital asset pricing model (CAPM) or Fama and French factor models have become the cornerstone of empirical asset pricing research and are also widely-used in corporate finance studying capital budgeting, for instance.

What do we actually know about the first moments (e.g., the "factor risk premiums") of the often-used Fama and French factors? Whereas Dichev (1998) and Chan et al. (2000) document that the relative performance of small and large firms has been much smaller and often even negative since the early 1980 s, Hirshleifer (2001) argues that the U.S. small firm effect appears to have disappeared. Similarly, Schwert (2003) highlights that the value effect seems to have disappeared after the papers that highlighted them were published. Even Fama and French (2015) concluded that the value factor is redundant after accounting for the investment factor in their five-factor model. In this regard, Grobys and Kolari (2022) perform international tests on Fama and French factor models and conclude:

> Surprisingly, employing FF (2018) test for nested models, the three-factor model did not outperform the CAPM in NA, Europe, and Japan. This finding is interesting in view of the fact that approximately 30 years ago FF proposed the three-factor model in lieu of the CAPM's inability to describe the cross section of average returns. (Grobys and Kolari, 2022, p. 644)

Interestingly, 30 years after Fama and French (1992) declared the death of the CAPM, amounting research provided evidence for that their own proposed factor models can be queried. Moreover, a recent study of Dong Li Rapach & Zhou (2022) explores the information content of cross-sectional anomaly strategies. Employing a large pool of factor returns and various model specifications, the authors find that past returns on long–short anomaly portfolios help to forecast the market risk premium within the US market. On the other hand, another recent study from Cakici, Fieberg, Metko, and Zaremba (2023) investigates the same issue using new data from both US and international markets, and explores hundreds of anomalies in forty-two countries around the world. The authors rely on machine-learning techniques to forecast market equity premia based on anomaly portfolio returns. Contrary to Dong et al. (2022), Cakici et al. (2023) conclude that equity anomalies cannot predict market returns:

> Any apparent predictability lacks external validity in two critical aspects: stock market selection and anomaly sample. While some evidence may be spotted in individual markets—such as the USA—it originates from a handful of specific anomalies and depends heavily on seemingly unimportant methodological choices. (Cakici et al., 2023)

Overall, recent evidence confirms that we do not seem to know much about the reliability of factor premiums which should be—according to theory—a compensation for risk. What about factor risk?

This is the first study that explicitly explores the risk of the popular Fama and French equity factors in terms of their realized variances. Using realized variances offers a more accurate representation of the risk dynamics associated with asset returns, as they capture both the magnitude and the temporal dependencies of return fluctuations (Andersen et al., 2001). Further, realized variances have been shown to provide a more precise measure of volatility, as they are less susceptible to biases and measurement errors that can affect other volatility estimators (Barndorff-Nielsen & Shephard, 2002). We compute the realized variances for the excess market factor (MKT), size factor (SMB), value factor (HML), investment factor (CMA), profitability factor (RMW), and momentum factor (MOM). The sample period for MKT, SMB, and HML is from June 1926 to September 2022, whereas the sample period for RMW and CMA is from June 1963 to September 2022. The sample for MOM factor encompasses the period from November 1926 to September 2022. Following Mandelbrot (1963a) and Taleb (2020), who advocated the usage of power laws, a novel aspect of this study is that it models realized variances as power laws. We estimate factor risk-specific power-law exponents using the maximum-likelihood estimation (MLE) procedure as proposed by Clauset et al. (2009). Also, we test the power-law models using the goodness-of-fit test derived from Clauset et al. (2009). Since previous literature documented that realized volatility of financial assets is close to log-normal (e.g., Andersen, Bollerslev, Diebold, & Ebens, 2001; Andersen, Bollerslev, Diebold, and Labys, 2001a & 2001b), another important aspect of this study is that we perform empirical tests to clarify the plausibility of other standard distributions including the well-established lognormal.

This study contributes to the current literature in some important ways. For instance, this the first study that explores whether the risk of well-established equity factors is subject to power-law behavior. Lux and Alfarano (2016) provide an extensive literature review on the literature on power laws in financial economics. The authors argued that the pertinent literature gradually converged to the insight that financial asset returns are governed by power-law exponents significantly larger than 2 and mostly close to 3. As a result, the presumed existence of theoretical variances motivated the vast majority of finance scholars to use standard econometrics models such as ordinary least squares (OLS) which require the theoretical variance to exist. The question arises, if the theoretical variances of asset returns exists, why does the vast majority of asset pricing studies fail scientific replication, as pointed out in a recent study of Hou, Xue, and Zhang (2020)?[2] Further, one may wonder why does the literature on corporate bond factors, as documented by Dick-Nielsen,

[2] Note that Hou et al. (2020) is not the only study putting emphasis on replication failure in financial research. A recent study of Smith and Timmermann (2022) that uses a novel methodology for identifying pervasive and discrete changes in cross-sectional risk premia documents that size, value, and investment risk premia are insignificantly different from zero in the most recent sample. Furthermore, Chen and Velikov (2023) analyze 204 stock market anomalies by accounting for effective bid–ask spreads, post-publication effects, and the modern era of trading technology that began in the early 2000s. Remarkably, Chen and Velikov (2023) are forced to conclude that "expected returns are negligible." Another recent study of Huang, Li, Wang, and Zhou (2020) explore return predictability in terms of time series momentum. The authors argue that the performance of time series momentum is virtually the same as that of a strategy that is based on historical sample mean and does not require any predictability.

Feldhütter, Pedersen, and Stolborg (2023), suffer from replication failures or inconsistent methodological choices? In addition, why do firm size, industrial sector, geographical region, or investment incentives not correlate with the patterns of errors in capital budgeting, as reported by Soares, Coutinho, and Martins (2007)?[3]

In an attempt to understand the abovementioned issues, one novel aspect of this study is that it takes a new perspective by analyzing realized variances because realized variances contain information that standard models such as generalized auto-regressive conditional heteroskedasticity (GARCH) models – which are typically derived under the assumption of normality – cannot reveal (e.g., Bubák, Kočenda, and Žikeš, 2011; Andersen, Bollerslev, Diebold, and Labys, 2003; Andersen, Bollerslev, and Meddahi, 2004). Hence, our approach may reveal novel insights concerning the risk dynamics of the popular Fama and French equity factors. We hypothesize that a potential reason for why finance studies suffer from the aforementioned issues could be that the variance of equity factors is infinite—a result which would be in stark contrast to the pertinent literature (Lux and Alfarano, 2016).

Next, some recent literature has emerged that explores power-law behavior of asset market variances. For instance, Grobys (2021) investigates the dynamics of annualized daily realized variances for the S&P 500, gold, crude oil, the GBP/USD exchange rate, and Bitcoin. His findings indicate that the power-law null model cannot be rejected for any of the realized variances. Moreover, the estimated power-law exponents are within the range $2 < \widehat{\alpha} < 3$, indicating that the variance of those realized asset market variances is not defined. Grobys (2021) argues that a manifestation of this issue is that estimated *t*-statistics are subject to sample-specificity in finite samples. Moreover, Grobys (2021) concludes that the potential reason for replication failures in financial economics could be that research methodologies often employed in financial economics are not designed for that specific research environment. Moreover, in another recent study, Grobys (2023a) explores the dynamics of annualized daily realized variances for the G10 currencies. His findings indicate that the power-law null model cannot be rejected for the vast majority of realized foreign exchange rate variances. Strikingly, estimated power-law exponents are within the range $2 < \widehat{\alpha} < 3$ for all realized foreign exchange rate variances and estimated power-law exponents exhibit a high level of stability across various subsamples.

The current study extends this emerging field of literature in some important aspects. First, it is the first that explores realized variance risk for equity market factors which are long-short portfolios and therefore could exhibit very different risk dynamics as opposed to single assets. Given that factor investing reached about $5 trillion assets under management, this is an important issue that needs to be investigated for enabling us to accurately managing factor risks. Second, some well-established literature argues that realized asset volatility is close to lognormally distributed (e.g., Andersen, Bollerslev, Diebold, & Ebens, 2001; Andersen, Bollerslev, Diebold, and Labys, 2001a & 2001b). Whereas the studies of Grobys (2021; 2023a) do not explicitly test the lognormal distribution against proposed power-law processes, this is the first study to test this issue for realized equity factor variances.

Using a sample from June 1926 to September 2022 for MKT, SMB, and HML, June 1963 to September 2022 for RMW and CMA, and November 1926 to September 2022 for MOM, we compute the monthly realized factor variances in line with Merton (1980). Q-Q plots show that the tails of the quantiles associated with the factor's realized variance distributions do not align with the theoretical quantile derived from the log-normal. Regardless of which realized factor variance we consider, the mean excess plot shows an upward linear trend that is mainly linear in the first part, which typically indicates a heavy-tailed distribution such as a power law. Using maximum likelihood estimation and selecting the cutoffs in line with Clauset et al.'s (2009) approach—which is based on the optimal Kolmogoroff-Smirnov distance—estimated power-law exponents vary between $\widehat{\alpha} = 1.96$ for the realized variance of the value factor, and $\widehat{\alpha} = 2.41$ for the realized variance of the excess market factor. Goodness-of-fit tests provide strong evidence for that lognormal or exponential distributions are rejected as underlying distributions generating realized factor variances, whereas the power-law null model cannot be rejected for any of the realized factor variances. This result is in stark contrast to the literature on realized volatility (e.g., Andersen, Bollerslev, Diebold, & Ebens, 2001; Andersen, Bollerslev, Diebold, and Labys, 2001a & 2001b).

Since $\alpha \approx 2$ for all realized factor variances, we test Mandelbrot's (1963b) infinite theoretical variance hypothesis. To do so, we use blocks bootstraps to compute robust standard deviations of the estimated power-law exponents. Our findings strongly indicate that the infinite theoretical variance hypothesis cannot be rejected for any of the realized factor variances because estimated *p*-values are within the range $0.06 \leq p \leq 0.48$. While these findings are contrary to the pertinent literature—as documented in Lux and Alfarano (2016)—they are consistent with the Mandelbrot's (1963b) infinite theoretical variance hypothesis.

Further evidence suggests that the dynamics of realized factor variances can be priced by Grobys' (2023b) recently proposed multifractal model of asset invariances. Specifically, the corresponding multifractal model which is capable of generating the factor variances for all Fama and French equity factor incorporates a multiplicative cascade derived from binominal bending with probability $p = 0.70$. Grobys (2023b) showed that the same multifractal model is capable of pricing the realized asset variances for crude oil and the S&P 500. Hence, an interesting commonality is here that realized factor variances appear to share a common power-law exponent manifested in $\alpha \approx 2$. Finally, we explore co-dependencies across realized factor variances measured in terms of co-fractality. Our findings indicate that for 11 out of 15 the co-fractality parameter is estimated at $\widehat{\lambda}^{CF} > 0.80$, suggesting a high level of (weak) co-fractality among realized factor variances. This implies, in turn, that there is not much room for risk diversification because power-law behavior often coincides; that is, extreme events are likely to occur within the same parts of time series paths.

From a more practical point of view, our results imply that the factor investing industry is exposed to a considerably higher level of extreme risk than earlier believed because the power law exponent captures via extrapolation low-probability deviations not seen in

[3] In a recent study, Dessaint et al. (2021, p. 35) commented on valuation errors in capital budgeting that "acquirers experience significantly lower cumulative abnormal returns when announcing bids for low rather than high beta targets and [the authors] estimate that using the CAPM in this context leads to valuation errors (relative to the market's view) that correspond, on average, to 12% to 33% of the deal values.".

the data (Taleb, 2020). From a more theoretical perspective, the results of our study have some serious implications. As pointed out in Fama (1963):

> From a purely statistical standpoint, if the population variance of the distribution of first differences is infinite, the sample variance is probably a meaningless measure of dispersion. Moreover, if the variance is infinite, other statistical tools (e.g., least-squares regression) which are based on the assumption of finite variance will, at best, be considerably weakened and may in fact give very misleading answers. (Fama, 1963, p. 421)

In line with Fama's (1963) implications derived from infinite variances, we hypothesize that both documented replication failures and inconsistent methodological choices (Hou et al., 2020; Dick-Nielsen et al., 2023) could be manifestations of undefined risk in terms of factor variances because statistical inference derived from OLS or any other methodology that requires the variance to exist will inevitably result in misleading inference.

This study is organized as follows. The next section describes the data and provides some preliminary analysis. The third section provides the methodology, whereas the fourth section discusses the findings. The last section concludes.

2. Data and preliminary analysis

The data used in this study include daily observations of the MKT, SMB, HML, RMW, CMA, and MOM factors, obtained from the data library of Kenneth R. French. Following Merton (1980), monthly realized variances are computed as the sum of the squared deviation of daily returns from their monthly means according to Equation (1):

$$RV_{j,t} = \sigma^2_{j,t} = \sum_{i=1}^{n} \left(R_{i,j,t} - \overline{R}_{j,t}\right)^2, \tag{1}$$

where, $RV_{j,t}$ ($\sigma^2_{j,t}$) denotes the monthly realized variance of factor j in month t, $R_{i,j,t}$ denotes the return of factor j on day i of a specific month t, $\overline{R}_{j,t}$ denotes the corresponding average daily return of factor j in month t, and n denotes the number of trading days per month. Note that this or similar measures for monthly realized variances or volatilities derived from daily data are frequently used in finance research (e.g., Moreira and Muir, 2017; Cederburg, O'Doherty, Wang, and Yan, 2020; Barroso and Detzel, 2021). Due to data availability, the sample period for the factor variances of MKT, SMB, and HML is from June 1926 to September 2022 (1,155 monthly observations), whereas the sample period for RMW realized variance and CMA realized variance is from June 1963 to September 2022 (711 monthly observations). The sample for MOM factor variance encompasses the period from November 1926 to September 2022 (1,151 monthly observations).

Table 1 presents the descriptive statistics for the monthly realized variances of the Fama and French factors. The mean values show that MKT has the highest average realized variance, followed by MOM, HML, SMB, RMW, and CMA. The median values reveal a similar pattern, with MKT realized variance exhibiting the highest median value and CMA realized variance the lowest. In addition, the maximum and minimum values indicate a wide range of realized variances across all factors, with the largest variation occurring for MKT realized variance. The standard deviations highlight the high level of uncertainty in realized factor variances, with MKT realized variance having the highest standard deviation, and CMA realized variance the lowest. All realized factor variances show extremely high kurtosis values ranging from 27.34 for HML realized variance to 66.32 for MKT realized variance, which indicates the presence of extremely fat tails in the distribution. Furthermore, among all those factor variances, the skewness is found to be most profound for MKT realized variance.

Table 1
Summary statistics for the realized Fama and French factor variances.

	$\widehat{\sigma}^2_{MKT}$	$\widehat{\sigma}^2_{SMB}$	$\widehat{\sigma}^2_{HML}$	$\widehat{\sigma}^2_{RMW}$	$\widehat{\sigma}^2_{CMA}$	$\widehat{\sigma}^2_{MOM}$
Mean	24.41	7.27	8.01	3.05	2.74	12.44
Median	10.11	3.20	2.89	1.25	1.45	4.39
Maximum	717.39	208.01	153.89	78.34	54.56	389.96
Minimum	0.54	0.25	0.30	0.16	0.15	0.18
Std. Dev	51.04	16.59	16.36	6.00	4.36	26.99
Skewness	6.84	6.41	4.55	6.33	5.58	6.40
Kurtosis	66.32	53.81	27.34	58.28	47.13	61.84
Observations	1,155	1,155	1,155	711	711	1,151

Data used to compute factor variances for the Fama and French factors are collected from the Data Library of Kenneth R. French and cover the excess market return (MKT), size factor (SMB), value factor (HML), profitability factor (RMW), investment factor (CMA), and momentum factor (MOM). The first three factors are the original Fama and French three factors model (Fama and French, 1993), while RMW and CMA are from the original Fama and French five-factor model (Fama and French, 2015). The MOM factor represents the momentum factor as proposed by Carhart (1997). The sample period for the MKT, SMB, and HML realized variance data (e.g., $\widehat{\sigma}^2_{MKT}$, $\widehat{\sigma}^2_{SMB}$, $\widehat{\sigma}^2_{HML}$) spans the period July 1926 to September 2022 (1,155 monthly observations). The RMW and CMA realized variance data (e.g., $\widehat{\sigma}^2_{RMW}$, $\widehat{\sigma}^2_{CMA}$) cover the period from July 1963 to September 2022 (711 monthly observations), whereas the MOM factor variance data (e.g., $\widehat{\sigma}^2_{MOM}$) encompass the period from November 1926 to September 2022 (1,151 monthly observations).

Next, we continue to visually investigate the time series graphs of each factor's realized variance to gain a comprehensive understanding of their behavior over time. Fig. 1 shows the evolution of monthly realized variance for each factor over the sample period. The time series graphs are particularly informative, as they demonstrate that—although the data typically fluctuates around the sample mean—there are several sudden shifts in each of the figures. These sudden shifts serve as indicators for extreme events, such as the Great Depression from 1929 to 1939, the recession in the early 1990 s, the 2007–2008 financial crisis, the European sovereign debt crisis spanning from 2009 to 2010, and the COVID-19 pandemic in 2020.

In addition, Fig. 2 depicts the discrepancy between the theoretical log-normal distribution (black curve), and the distribution obeyed by each factor's monthly realized variance (red curve). Self-evidently, Fig. 2 shows that the log-normal distribution does not accurately interpolate the density function of the realized factor variance series'. In this regard, Taleb (2020, p. 23), points out that "as we fatten the tails we get higher peaks, smaller shoulders, and a higher incidence of a very large deviation." We clearly observe this issue in Fig. 2, where the peaks of the realized factor variance distributions exceed the theoretical peaks predicted by the log-normal distribution (red curve). We observe the same pattern for each realized factor variance.

Further visual evidence is provided by investigating the log-normal Q-Q plots, shown in Fig. 3. Fig. 3 reveals that the tails of the quantiles associated with the factor's realized variance distributions do not align with the theoretical quantile derived from the log-normal (red line). From Fig. 3 we observe that both ends of the Q-Q plots for all factor's variances deviate from the theoretical log-normal quantiles, whereas the center follows the red line. This pattern is indicative of a fat-tailed distribution implying a higher frequency of extreme events than what would be anticipated under a log-normal distribution.

As a final tool for visual inspection, we employ the mean excess plot (ME plot) as additional potentially indicative tool for identifying power-law processes. For a given threshold u, the mean excess of a random variable X is defined as,

$$E(X-u|X>u). \tag{2}$$

An upward linear trend in the mean excess plot signifies a power law, where the slope of the plot has a positive derivative with respect to the tail index parameter alpha (α). A larger value of α corresponds to a steeper slope in the mean excess plot. As depicted in Fig. 4, an upward linear trend that is mainly linear in the first part may indicate a heavy-tailed distribution such as a power law. However, further analysis is necessary to distinguish between these distributions and identify the specific distribution.

3. Methodology

3.1. Estimating the power-law exponents using maximum-likelihood estimation

Following Grobys (2021; 2023a), we model realized Fama and French factor variances using the following power-law function:

$$p(x)=Cx^{-\alpha}, \tag{3}$$

where $C=(\alpha-1)x_{MIN}^{\alpha-1}$ with $\alpha\in\{\mathbb{R}_+|\alpha>1\}$, $x=\sigma_{j,t}^2$ denotes the respective realized Fama and French factor variance provided $x\in\{\mathbb{R}_+|x_{MIN}\leq x<\infty\}$, x_{MIN} is the minimum value of realized Fama and French factor variance that is governed by a power-law process, and α is the magnitude of the specific tail exponent.[4] Choosing power laws for modeling the extremely fat tails is in line with Mandelbrot (1963a) who argued that there are strong pragmatic reasons to begin the study of economic distributions and time series by those that satisfy the law of Pareto. Moreover, Taleb (2020) highlighted that the tail exponent α of a power law function captures via extrapolation the low-probability deviation not seen in the data which plays a disproportionately large role in determining the mean of the theoretical distribution.

The expectation for the realized factor variance, provided $X>x_{MIN}$, defined as $E[X|X>x_{MIN}]$, is given by:

$$E[X|X>x_{MIN}]=\int_{x_{MIN}}^{\infty}xp(x)dx=\frac{(\alpha-1)}{(\alpha-2)}x_{MIN}. \tag{4}$$

The second moment, provided $X>x_{MIN}$, defined as $E\left[X^2|X>x_{MIN}\right]$ (viz., the variance of the variance), is defined as:

$$E\left[X^2|X>x_{MIN}\right]=\int_{x_{MIN}}^{\infty}x^2p(x)dx=\frac{(\alpha-1)}{(\alpha-3)}x_{MIN}^2. \tag{5}$$

Higher moments of order k are analogously defined as:

[4] This study follows the notation in Clauset et al. (2009). To simplify notation, the index i denoting the respective realized factor variance j is dropped. Note that some other studies parametrize the *tail probability* as $P(|X|>x)\propto x^{-\alpha}$ as opposed to parametrizing the *density* function as in Equation (3) (see, for instance, Lux and Alfrano, 2016). In our study, we use the maximum-likelihood estimation as described in detail in the often-cited study of Clauset et al. (2009) which gives us point estimates for α and corresponding goodness-of-fits tests derived from the function $p(x)=Cx^{-\alpha}$. Hence, in our study we parametrize the *density* function instead of the *tail probability*.

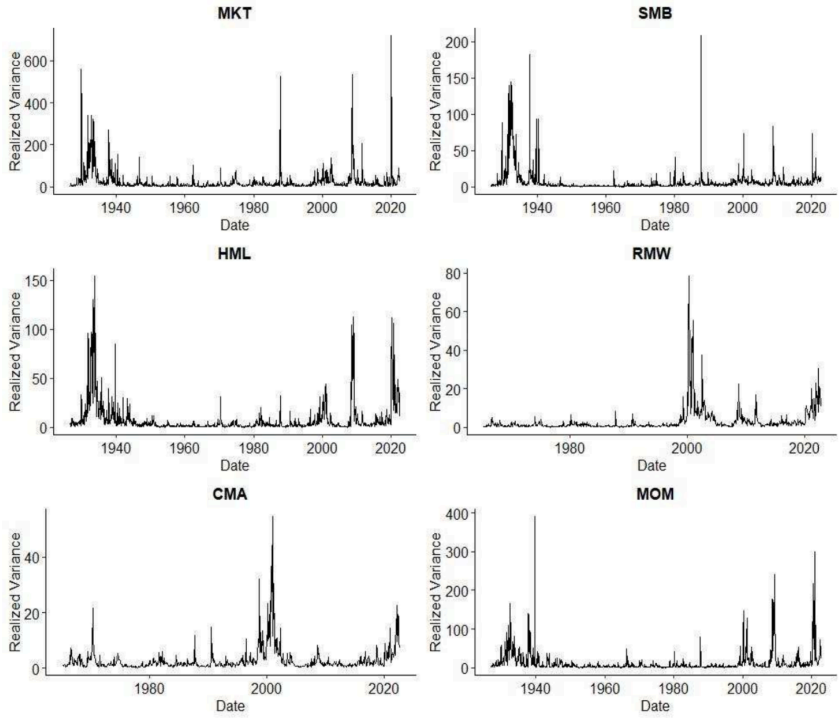

Fig. 1. Time series plot. The time series graph of monthly realized variance for excess market return (MKT), size (SMB) and value (HML) factors from July 1926 to September 2022 (1,155 monthly observations), profitability (RMW), investment (CMA) factors from July 1963 to September 2022 (711 monthly observations) and momentum factor (MOM) from November 1926 to September 2022 (1,151 monthly observations).

$$E\left[X^k | X > x_{MIN}\right] = \frac{(\alpha - 1)}{(\alpha - 1 - k)} x_{MIN}^k. \tag{6}$$

From Equations (4) and (5), we observe that the theoretical mean for realized factor variances only exists for $\alpha > 2$, whereas the variance of realized factor variance requires $\alpha > 3$.

Following White et al. (2008) and Clauset et al. (2009), who argued that maximum-likelihood estimation (MLE) performs best for estimating power law exponents, tail exponents are estimated as:

$$\widehat{\alpha} = 1 + N\left(\sum_{i=1}^{N} \ln\left(\frac{x_i}{x_{MIN}}\right)\right)^{-1}, \tag{7}$$

where $\widehat{\alpha}$ denotes the MLE estimator, $N \leq T$ is the number of observations exceeding x_{MIN}, and other notation is as before. As noted in Clauset et al. (2009), it is necessary to determine the relevant value for x_{MIN} to accurately estimate the most adequate tail exponent. More specifically, as shown in Equation (7), the MLE estimator depends on the chosen x_{MIN} and, hence, there are different MLE estimators from which to choose. In this regard, Clauset et al. (2009) observed that it is common practice to employ the $\widehat{\alpha}/x_{MIN}$-plot and choose the value for x_{MIN} beyond which $\widehat{\alpha}$ is stable.

Since this approach is somewhat subjective and can be sensitive to noise or fluctuations in the tail of the distribution, the authors proposed to select x_{MIN} based on minimizing the distance between the power-law model and the empirical data: Specifically, the Kolmogorov–Smirnov distance (D) is the maximum distance between the cumulative density functions (CDFs) of the data and the fitted model as defined by:

$$D = MAX_{x \geq x_{MIN}} |S(x) - P(x)|, \tag{8}$$

where $S(x)$ is the CDF of the data for the observation with value at least x_{MIN}, and $P(x)$ is the CDF for the power-law model that best fits the data in the region $x \geq x_{MIN}$. The estimate $\widehat{x}_{MIN}$ is then the value of x_{MIN} that minimizes D. Further, the standard deviation for estimated power-law exponents is according to Clauset et al. (2009) given by:

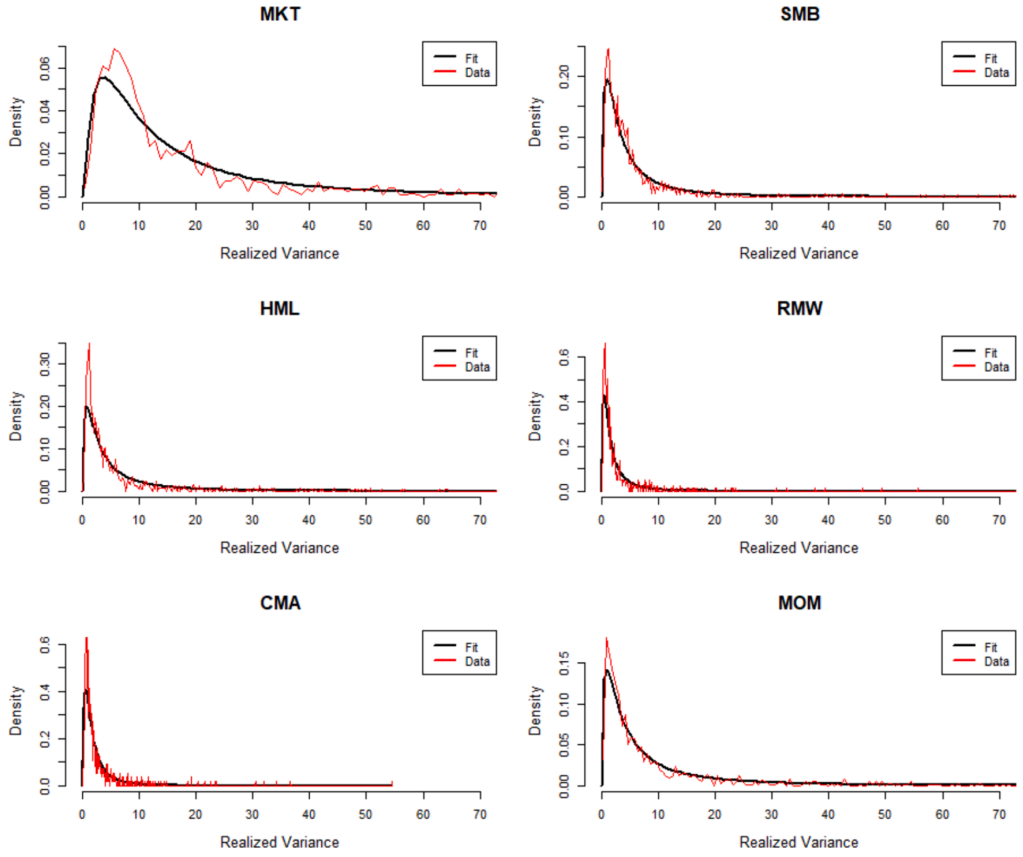

Fig. 2. Histogram plot. Theoretical Log-normal distribution (black curve) versus distribution followed by the factor's variance (red curve) for the monthly realized variance of the excess market return (MKT), size factor (SMB), value factor (HML), profitability factor (RMW), investment factor (CMA), and momentum factor (MOM). (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

$$\sigma = \frac{\widehat{\alpha} - 1}{\sqrt{N}} + O\left(\frac{1}{N}\right). \tag{9}$$

The descriptive statistics for estimated power law exponents are reported in Table 2. From Table 2 we observe that the tail exponent is estimated at $\widehat{\alpha} < 3$ for all Fama and French factor variances, indicating that the variance of realized factor variance does not exist for any of the factor's variances. This implies, in turn, that the distributions of Fama and French realized factor variances has no theoretically defined variance. In the context of realized asset variances, this result further implies that extreme market movements and tail risk might be more prevalent than assumed when modeling realized factor variance via the lognormal distribution.

Furthermore, the most striking result from Table 2 is that the estimated alpha for the HML realized variance is estimated at $\widehat{\alpha} < 2$, which indicates that not only the variance for this factor variance is undefined, but also the theoretical mean is undefined. This is an issue that has serious implications which will be discussed later.[5]

[5] In addition, from the results presented in last row of Table 2, one might question the limited number of observations governed by the power law distribution, suggesting that the power law explains only a small fraction of the time series observations. However, it is important to note that the shares of the top 20 percent of the cumulative total of the distributions are 68.01, 69.00, 72.98, 68.10, 60.84, and 72.59 percent for the realized factor variances of the MKT, SMB, HML, RMW, CMA, and MOM, respectively. This indicates that a small fraction of observations has a large impact on the entire dataset. These numbers are closely aligned with the well-known Pareto 80/20 rule, where the top 20 percent of the cumulative total accounts for 80 percent of the cumulative total of the distribution. This means that the vast majority of the distribution of realized factor variances can essentially be considered as noise—simply because it has minimal impact compared to the significant influence employed by the minority segment that is governed by a power law process.

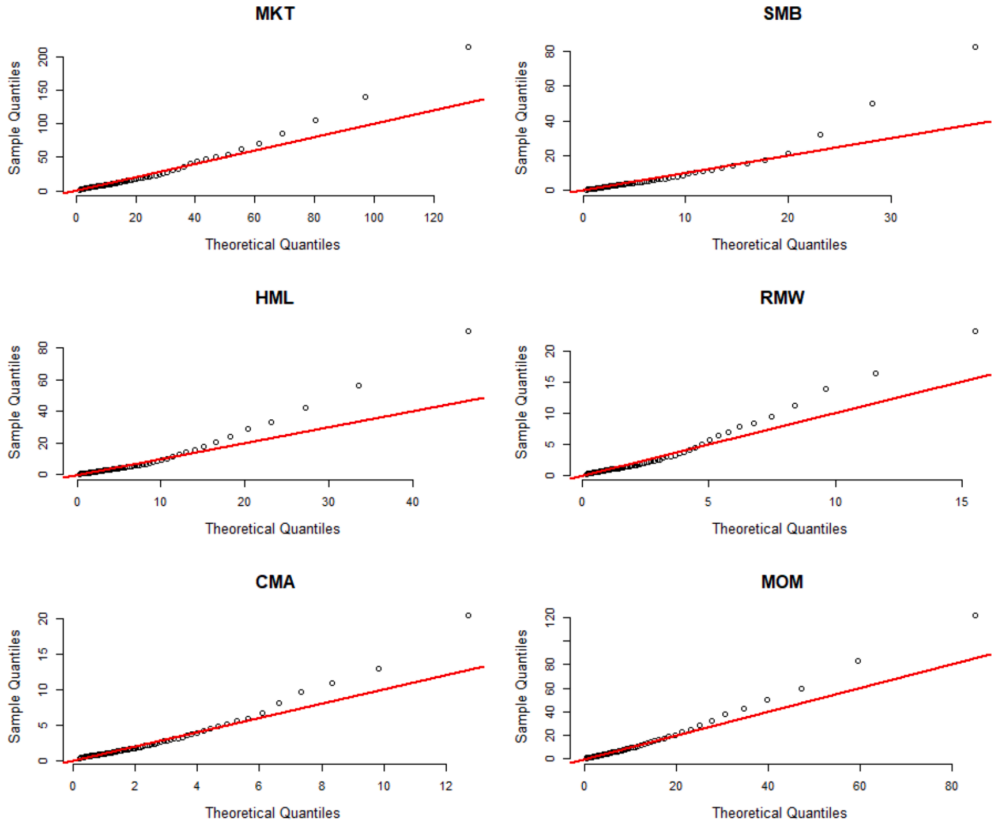

Fig. 3. Log-normal QQ Plot. The theoretical log-normal quantile (red line) versus factor's variance quantile for the monthly realized variance of excess market return (MKT), size factor (SMB), value factor (HML), profitability factor (RMW), investment factor (CMA), and momentum factor (MOM). (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

3.2. *Is the power-law null model plausible?*

The findings from the previous section should be approached with caution, as a tail exponent within the range $2 < \alpha < 3$ does not necessarily imply that the power law is an appropriate fit for the data. Therefore, it is important to verify whether the observed dataset genuinely exhibits a power-law distribution. Following Clauset et al. (2009), we test this hypothesis by goodness-of-fit test, derived from simulations using the Kolmogorov-Smirnoff distance D, as given in Equation (8), to determine whether the empirical data and the generated data from the power law distribution with specific x_{MIN}, and α values belong to the same distribution. The null hypothesis of this test asserts that the empirical data and the generated synthetic data from the power-law distribution, with specific values for x_{MIN} and α, belong to the same distribution, implying that the power law distribution provides a plausible fit to the empirical data.

The alternative hypothesis posits that the empirical data and the generated synthetic data do not belong to the same distribution, suggesting that the power-law distribution does not provide a plausible fit to the empirical data and other possible distributions should be considered. To perform this analysis, we generate large synthetic datasets of power-law distribution using the estimated parameters x_{MIN} and α from the previous section for each realized factor variance. Specifically, the p-value of this test is defined as the fraction $\#(D_S > D)/K$, where D_S is the estimated Kolmogorov-Smirnoff distance, as defined in Equation (8), for some synthetic dataset for given x_{MIN} and α values (as estimated for the original factor variances), D is the estimated Kolmogorov-Smirnoff distance for the original dataset, and K is the number of synthetic datasets generated.[6]

We choose a significance level of 5 %, that is, we do not reject the power law model if $\#\frac{D_S > D}{K} > 0.05$. To ensure the reliability and

[6] That is, each synthetic data set is individually fit to its corresponding power-law model, and the Kolmogorov-Smirnov distance D is computed for each model. Then, we compute the p-value by counting the fraction of the time that the resulting statistic exceeds the value for the empirical data.

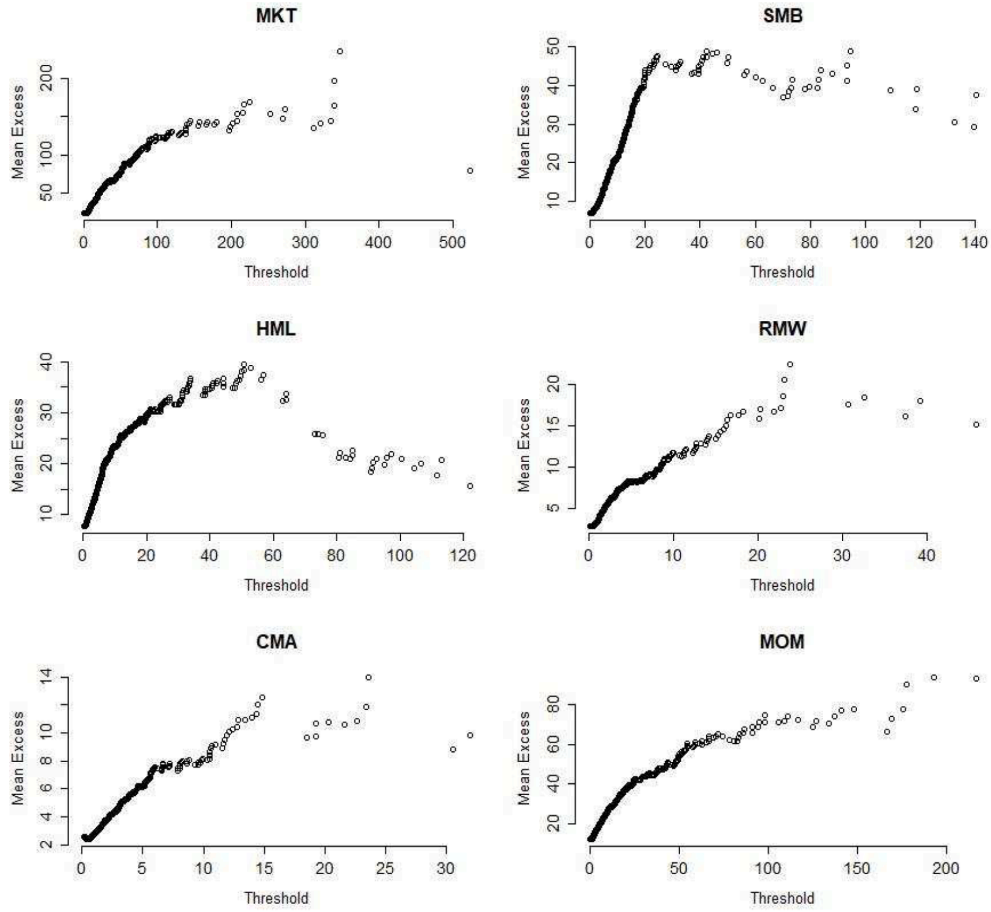

Fig. 4. Mean Excess Plot. This figure plots the sorted data on the x-axis against the corresponding values of the mean excess function for the monthly realized variance of the excess market return (MKT), size factor (SMB), value factor (HML), profitability factor (RMW), investment factor (CMA), and momentum factor (MOM). For a given threshold u of a random variable X, the mean excess is defined as $E(X-u|X>u)$. An upward linear trend in the mean excess plot is indicative of a Pareto distribution, where the slope of the plot has a positive derivative with respect to the tail index parameter alpha (α).

robustness of our findings, we simulated 2,500 synthetic datasets to examine whether the empirical data and the generated dataset belong to the same distribution.[7] Table 3 provides the results of the goodness-of-fit tests using our realized factor variance data.[8] The second column of Table 3 reports the p-value for testing the power-law null hypothesis. We observe that for all realized factor variances, the p-values exceed 5 %, which provides strong evidence for that power laws appear to be indeed plausible distributions for modeling realized factor variances.

3.3. *Estimating robust standard errors and testing Mandelbrot's infinite variance hypothesis*

Since the estimated power-law exponents are close to $\alpha \approx 2$ for all factor variances, a natural question that arises is whether the theoretical means of realized factor variances are statistically defined. This is an important issue to clarify because undefined theo-

[7] The choice of 2,500 simulations are based on Clauset et al. (2009) recommendation, which is intended to achieve a balance between computational feasibility and the precision of the p-value estimation.

[8] In this study, we utilized the code provided in the paper 'Fitting Power Law Distributions to Data' by Willy Lai. We would like to thank Willy Lai for their valuable contribution and for making the code available, which greatly facilitated our research and analysis. The paper can be accessed at: https://www.stat.berkeley.edu/~aldous/Research/Ugrad/Willy_Lai.pdf.

Table 2
Estimating power law exponents.

	$\widehat{\sigma}^2_{MKT}$	$\widehat{\sigma}^2_{SMB}$	$\widehat{\sigma}^2_{HML}$	$\widehat{\sigma}^2_{RMW}$	$\widehat{\sigma}^2_{CMA}$	$\widehat{\sigma}^2_{MOM}$
$\widehat{\alpha}$	2.41	2.16	1.96	2.00	2.39	2.10
$\widehat{\sigma}$	0.11	0.05	0.04	0.05	0.10	0.07
x_{MIN}	40.14	3.38	3.31	0.97	2.50	11.44
D	0.03	0.02	0.04	0.04	0.04	0.06
% observations for which $x \geq x_{MIN}$	13.85 %	48.31 %	44.85 %	64.14 %	28.83 %	24.07 %

This table reports the results of estimating power law exponents for the following realized factor variances: Excess market factor (MKT), size factor (SMB), value factor (HML), profitability factor (RMW), investment factor (CMA), and momentum factor (MOM). The $\widehat{\alpha}$ represents the estimated tail exponent and $\widehat{\sigma}$ is the estimated standard deviation. The lower threshold x_{MIN} is estimated using the optimized Kolmogorov-Smirnov distance (D) distance as proposed by Clauset et al. (2009). The estimated x_{MIN} is the value that corresponds to the optimal distance D. Additionally, for each factor variance, the fraction of observations governed by power law process is reported. The sample period for the MKT, SMB, and HML realized variances span the period July 1926 to September 2022 (1,155 monthly observations). The RMW and CMA realized variances cover the period from July 1963 to September 2022 (711 monthly observations), whereas the MOM factor variance encompasses the period from November 1926 to September 2022 (1,151monthly observations).

retical means would imply that standard econometric models based on the concept of correlation—such as ordinary least squares (OLS) or generalized methods of moments (GMM)—cannot be used to draw conclusions, as they would inevitably generate misleading statistical results (Fama, 1963).

It is important to note that the standard errors documented in Clauset et al. (2009) are derived under the assumption of *independence*. It is, however, a stylized fact that financial market data typically exhibit volatility clustering which suggest some type of dependence structure.[9] As a result, the standard deviations reported in Table 3 are underestimated, and hence, inadequate for hypothesis testing. To address this issue, we follow Grobys (2023c) and employ blocks bootstraps in line with Grobys and Juntilla (2021) for estimating robust standard deviations.

Denoting the selected block length as m, we implement a blocks bootstrap procedure such that $E[m] = T^{1/3}$ which is the suggested block length for variance estimation in Godfrey (2009). Then from a given data vector $\boldsymbol{x}$, blocks m are randomly drawn which are distributed as a geometric distribution $m \sim GEO(p)$ with $E[m] = \frac{(1-p)}{p}$. For example, for the overall data sample, for the realized variances of the MKT, SMB, HML, and MOM factor, we employ $p = 0.0909$, whereas we employ $p = 0.1111$ for the realized variances of the RMW and CMA factor. Employing this bootstrap approach, the blocks drawn from a given data vector $\boldsymbol{x}$ vary in lengths. Randomly drawn blocks m from data vector $\boldsymbol{x}$ are stacked in vector $\boldsymbol{x}^b$ as:

$$\boldsymbol{x}_i^b = \begin{bmatrix} m_1 \\ m_2 \\ m_3 \\ \vdots \end{bmatrix}.$$

Note again that the length of the blocks $m_1, m_2, \ldots$ varies. The procedure is stopped when the length of the artificial vector $\boldsymbol{x}^b$ exhibits a length that exceeds T. Observations exceeding T are cut off; that is, every artificial data vector $\boldsymbol{x}^b$ has the same length as the original data vector $\boldsymbol{x}$. Using this blocks bootstrap procedure, for each data vector, $B = 1,000$ artificial data vectors are constructed,

$$\left[\boldsymbol{x}^1 \quad \boldsymbol{x}^2 \quad \cdots \quad \boldsymbol{x}^B\right],$$

and point estimates for α are obtained for each bootstrapped data vector $\boldsymbol{x}^1, \boldsymbol{x}^2, \ldots, \boldsymbol{x}^B$ using Clauset et al.'s (2009). For each realized factor variance, the estimated tail exponents of the bootstrap data are stacked in a 1xB vector,

$$\left[\widehat{\alpha}^1 \quad \widehat{\alpha}^2 \quad \cdots \quad \widehat{\alpha}^B\right].$$

Finally, the corresponding bootstrapped standard error $\widehat{\sigma}_{BOOT}$ is computed for the bootstrapped $\widehat{\alpha}^1, \widehat{\alpha}^2, \ldots, \widehat{\alpha}^B$ for each given realized factor variance. The descriptive statistics for the blocks bootstrapped tail exponents are reported for each realized factor variance in Table 4. Comparing the bootstrapped standard deviations with the ones reported in Table 2, it becomes evident that once (unknown) dependence structures (i.e., volatility clustering) are accounted for, standard deviations are substantially larger. As an example, the ordinary standard deviation for the realized HML realized variance is estimated at $\widehat{\sigma} = 0.05$, whereas the bootstrapped counterpart is estimated at $\widehat{\sigma}^{boot}_{HML} = 0.16$; that is, once dependence structures are accounted for, the estimation uncertainty is about three times larger.

[9] While traditional finance research typically refers to the concept of autocorrelation in this context, we would rather like to refer to dependency as the concept of correlation requires the variance to be defined.

Table 3
Goodness-of-fit test.

Realized factor variance	Power law	Log-normal	Exponential
$\widehat{\sigma}^2_{MKT}$	0.84	0.01	0.00
$\widehat{\sigma}^2_{SMB}$	0.59	0.00	0.00
$\widehat{\sigma}^2_{HML}$	0.69	0.00	0.00
$\widehat{\sigma}^2_{RMW}$	0.42	0.03	0.00
$\widehat{\sigma}^2_{CMA}$	0.50	0.03	0.00
$\widehat{\sigma}^2_{MOM}$	0.06	0.00	0.00

This table reports the results from goodness-of-fit test, by Kolmogorov Smirnoff test to examine if our empirical data and the generated data from the power law distribution with particular x_{min} and α belong to the same distribution in the second column. The table also represents the goodness-of-fit test for log- normal and exponential distributions in the third and fourth columns respectively. The null hypothesis of this test asserts that the empirical data and the generated synthetic data from the power law distribution, with specific x_{MIN} and α values, belong to the same distribution, implying that the power law distribution provides a plausible fit to the empirical data.

Moreover, it becomes evident that the means of the bootstrapped power-law exponents are close to the estimates for the original data. However, this should not come as a great surprise because the point estimates for the exponents obtained via Clauset et al.'s (2009) MLE approach should be unbiased.

Because the point estimates for the tail exponents of the realized factor variances are close to $\alpha \approx 2$, a natural question that arises is: Can we reject the infinite mean hypothesis? Recall that Fama (1963) argued that correlation-based methodologies will inevitably result in very misleading statistical inference if the theoretical variance is undefined. Hence, it may be not surprising that Mandelbrot's (1963b) seminal study on cotton price changes—ascertaining that cotton price changes do not exhibit a defined theoretical variance—resulted in an enormous amount of follow-up research. Interestingly, in a literature review covering power laws in financial economics, Lux and Alfarano (2016) argued that other studies raised doubts over the validity of the infinite variance hypothesis by questioning the stability-under-aggregation property of these estimates, and the pertinent literature gradually converged to the insight of an exponent $\alpha \geq 3$. These studies, however, do not employ realized variances. As pointed out from Barndorff-Nielsen and Shephard (2002) realized variances provide a more precise measure of volatility, as they are less susceptible to biases and measurement errors that can affect other volatility estimators. As a result, we test the infinite theoretical variance hypothesis for all realized factor variances as follows:

$$H_0 : \alpha \leq 2 \text{ versus } H_1 : \alpha > 2\,.$$

In Panel B of Table 4 we report the *p*-values based on the empirical distribution of the bootstrapped power-law exponents. The *p*-value is given by the fraction of power-law exponents for which $\widehat{\alpha} \leq 2$ is satisfied. We observe from Panel B of Table 4 that the *p*-values are ≥ 0.05 for all realized factor variances. In Panel C we report the *p*-values derived from *t*-tests based on the bootstrapped standard deviations. The *p*-values are virtually the same, and hence, strongly confirm this evidence. Overall, it appears that realized factor variances are governed by a power-law process and regardless of which realized factor variance we consider, the hypothesis $\alpha \approx 2$ cannot be rejected which suggest a remarkable commonality across realized factor variances.

To further strengthen our results, we introduce a joint test which is able to assess whether Fama and French factor variances share a common component manifested in a common power-law exponent. First, we note that due to our bootstrapping approach $COV(\widehat{\alpha}_i, \widehat{\alpha}_j) = 0 \; \forall \; i \neq j$. As a consequence, we can define the following covariance matrix:

$$\widehat{\Sigma}_{\widehat{\boldsymbol{\alpha}}} = \begin{pmatrix} \widehat{\sigma}^2_{\widehat{\alpha}_{boot},MKT} & 0 & 0 & 0 & 0 & 0 \\ 0 & \widehat{\sigma}^2_{\widehat{\alpha}_{boot},SMB} & 0 & 0 & 0 & 0 \\ 0 & 0 & \widehat{\sigma}^2_{\widehat{\alpha}_{boot},HML} & 0 & 0 & 0 \\ 0 & 0 & 0 & \widehat{\sigma}^2_{\widehat{\alpha}_{boot},RMW} & 0 & 0 \\ 0 & 0 & 0 & 0 & \widehat{\sigma}^2_{\widehat{\alpha}_{boot},CMA} & 0 \\ 0 & 0 & 0 & 0 & 0 & \widehat{\sigma}^2_{\widehat{\alpha}_{boot},MOM} \end{pmatrix},$$

where $\widehat{\sigma}^2_{\widehat{\alpha}_{boot},MKT}, \ldots, \widehat{\sigma}^2_{\widehat{\alpha}_{boot},MOM}$ denote the bootstrapped variances obtained via block bootstrap as outlined earlier. Using the vector of estimated power-law exponents from Table 2, we define the following test statistic:

$$\widehat{\lambda} = (\widehat{\boldsymbol{\alpha}} - q1)' \widehat{\Sigma}_{\widehat{\boldsymbol{\alpha}}}^{-1} (\widehat{\boldsymbol{\alpha}} - q1), \tag{10}$$

where the covariance matrix $\widehat{\Sigma}_{\widehat{\boldsymbol{\alpha}}}$ has the dimension 6x6, $\widehat{\boldsymbol{\alpha}}$ is the 6x1 vector of estimated power-law exponents, 1 is a 6x1 vector of ones

Table 4
Summary statistics for blocks bootstrapped power law exponents.

Panel A. Descriptive statistics for the bootstrapped tail exponents.						
	$\widehat{\alpha}^{boot}_{MKT}$	$\widehat{\alpha}^{boot}_{SMB}$	$\widehat{\alpha}^{boot}_{HML}$	$\widehat{\alpha}^{boot}_{RMW}$	$\widehat{\alpha}^{boot}_{CMA}$	$\widehat{\alpha}^{boot}_{MOM}$
Mean	2.25	2.22	2.01	2.23	2.49	2.21
Median	2.21	2.20	1.99	2.13	2.45	2.13
Maximum	3.22	3.18	3.14	4.00	4.39	4.01
Minimum	1.78	1.84	1.68	1.67	1.79	1.73
Std. Dev.	0.24	0.18	0.16	0.33	0.32	0.33
Skewness	0.80	0.81	1.43	1.75	1.09	1.44
Kurtosis	3.50	4.56	8.59	7.34	5.99	5.24
Jarque-Bera	117.52	210.45	1,642.99	1,296.71	571.39	553.84
Probability	0.00	0.00	0.00	0.00	0.00	0.00
Observations	1,000	1,000	1,000	1,000	1,000	1,000
Panel B. Estimated *p*-values for testing the infinite mean hypothesis using the empirical distribution of bootstrapped power law exponents.						
p-value	0.14	0.11	0.54	0.24	0.05	0.36
Panel C. Estimated *p*-values for testing the infinite mean hypothesis using bootstrapped standard deviations.						
p-value	0.15	0.11	0.48	0.24	0.06	0.26

Denoting the selected block length as m, we implement a blocks bootstrap procedure such that $E[m] = T^{1/3}$. From a given data vector $\boldsymbol{x}$, blocks m_j are randomly drawn which are distributed as a geometric distribution $m_j \sim GEO(p)$ with $E[m_j] = \frac{(1-p)}{p}$. Employing this bootstrap approach, the blocks drawn from a given data vector $\boldsymbol{x}$ vary in lengths. Randomly drawn blocks m_j from data vector $\boldsymbol{x}$ are stacked in vector $\boldsymbol{x}^b$ as: $\boldsymbol{x}_i^b = \begin{bmatrix} m_1 \\ m_2 \\ m_3 \\ \vdots \end{bmatrix}$. The procedure is stopped when the length of the artificial vector $\boldsymbol{x}^b$ exhibits a length that exceeds T. Observations exceeding T are cut off; that is, every artificial data vector $\boldsymbol{x}^b$ has the same length as the original data vector $\boldsymbol{x}$. Using this blocks bootstrap procedure, for each data vector, $B = 1,000$ artificial data vectors are constructed, $[\boldsymbol{x}^1 \quad \boldsymbol{x}^2 \quad \cdots \quad \boldsymbol{x}^B]$, and point estimates for α are obtained for each bootstrapped data vector $\boldsymbol{x}^1, \boldsymbol{x}^2, \ldots, \boldsymbol{x}^B$ using Clauset et al.'s (2009) approach. For each factor variance, the estimated tail exponents of the bootstrap data are stacked in a 1xB vector, $[\widehat{\alpha}^1 \quad \widehat{\alpha}^2 \quad \ldots \quad \widehat{\alpha}^B]$. Finally, the corresponding bootstrapped standard error $\widehat{\sigma}_{BOOT}$ is computed for the bootstrapped $\widehat{\alpha}^1, \widehat{\alpha}^2, \ldots, \widehat{\alpha}^B$ for each given realized factor variance. The descriptive statistics for the blocks bootstrapped tail exponents are reported in Panel A, whereas Panel B and C report the *p*-values for testing the infinite mean hypothesis.

and q is the hypothesized common power-law exponent. The estimated test statistic, denoted as $\widehat{\lambda}$, is under the null hypothesis distributed as $\chi^2(6)$. The test statistic is iteratively estimated covering the economically important interval $q = (1.2, 1.3, \ldots, 3.1, 3.2)$. Note acceptance of the null hypothesis would suggest that a common power-law exponent exists that governs the realized factor variance risk of the Fama and French factors.

The results are reported in Table 5. Note that using a 5 percent significance level, the corresponding critical value for the chi-square distribution is $\chi^2_{0.95}(6) = 12.59$. Strikingly, the null hypothesis cannot be rejected for $1.7 < \alpha < 2.5$ which strongly confirms the presence of a common exponent governing the realized factor variances. Since the *p*-value for testing the null hypotheses $H_0 : \alpha = 1.9$ or $H_0 : \alpha = 2.0$ exceeds 5 percent, Mandelbrot's infinite variance hypothesis is confirmed because both hypotheses assume that the theoretical mean of the realized factor variances is undefined. Furthermore, the optimal α governing Fama and French realized factor variances is $\alpha = 2.1$ as this test statistic generates the largest *p*-value corresponding to 0.74. Finally, because the null hypothesis is clearly rejected for all joint tests for which $\alpha > 2.4$, we can clearly rule out that the variance of realized factor variances exists. Overall, a key finding is here that realized factor variances have a common component that is manifested in sharing the same power-law behavior—that is, the common component is manifested here in a common power-law exponent governing Fama and French factor variances.

3.4. Testing alternative distributions against power laws

Our results derived from Clauset et al.'s (2009) goodness-of-fit test do not guarantee that power-law distributions are ideal matches for the data, as there might be other distributions that, perhaps, could fit the data better. Therefore, for each realized factor variance, we compare the corresponding power-law distribution with the log-normal distribution to identify the most appropriate fit for the data. It is important to note that earlier research postulated that realized volatility is close to lognormally distributed (e.g., Andersen, Bollerslev, Diebold, & Ebens, 2001; Andersen, Bollerslev, Diebold, and Labys, 2001a & 2001b). Therefore, in this section we examine whether the lognormal distribution is a plausible fit to the data. According to Clauset et al. (2009), we test this hypothesis using a similar goodness-of-fit test with the same significance level of 5 percent like outlined earlier. Thereby, the null model which is assumed is the lognormal distribution. The results from the goodness-of-fit tests for lognormal distributions are presented in column 3 of

Table 5
Testing for a common power-law exponent governing Fama and French factor variance risk.

q	$\widehat{\lambda}$	p-value
1.2	**103.80**	**0.00**
1.3	**83.43**	**0.00**
1.4	**65.36**	**0.00**
1.5	**49.60**	**0.00**
1.6	**36.15**	**0.00**
1.7	**25.00**	**0.00**
1.8	**16.17**	**0.01**
1.9	9.64	0.14
2	5.43	0.49
2.1	3.52	0.74
2.2	3.91	0.69
2.3	6.62	0.36
2.4	11.64	0.07
2.5	**18.96**	**0.00**
2.6	**28.59**	**0.00**
2.7	**40.53**	**0.00**
2.8	**54.78**	**0.00**
2.9	**71.34**	**0.00**
3	**90.21**	**0.00**
3.1	**111.38**	**0.00**
3.2	**134.86**	**0.00**

To explore whether the exists a common component governing power-law behavior of the realized Fama and French factor variances, the following estimated test statistic is proposed: $\widehat{\lambda} = (\widehat{\boldsymbol{\alpha}} - q1)'\widehat{\Sigma}_{\hat{\boldsymbol{\alpha}}}^{-1}(\widehat{\boldsymbol{\alpha}} - q1)$, where the covariance matrix $\widehat{\Sigma}_{\hat{\boldsymbol{\alpha}}}$ has the dimension 6x6, $\widehat{\boldsymbol{\alpha}}$ is the 6x1 vector of estimated power-law exponents, 1 is a 6x1 vector of ones and q is the hypothesized common power-law exponent. The estimated test statistic denoted as $\widehat{\lambda}$ is under the null hypothesis distributed as $\chi^2(6)$. The test statistic is iteratively estimated covering the economically important interval $q = (1.2, 1.3, \ldots, 3.1, 3.2)$. The corresponding test statistic is under the null hypothesis distributed as $\chi^2(6)$. **Bold** figures indicate statistical significance on a 5 % level.

Table 3. Strikingly, we observe that the p-values for all of the six realized factor variances are less than 0.05. We interpret this as strong evidence for rejecting the assumption of time-honored lognormality of realized factor variances. Consequently, these findings suggest that finance researchers may need to move beyond the lognormal assumption to more accurately understand and predict the behavior of factor risk. In addition, to further strengthen our analysis, we also conduct goodness-of-fit tests using the exponential distribution as the null model. The exponential distribution serves as alternative model for fitting data exhibiting a considerably fatter tail than the normal distribution. The results of the goodness-of-fit tests for exponential distributions are reported in column 4 of Table 3. The documented findings suggest that p-values are zero in all cases, indicating that the exponential distribution does not provide an adequate fit for the realized factor variances series.

3.5. Additional analysis

3.5.1. Further evidence from one-sigma tests

While the goodness-of-fit tests used in previous sections have the assumed model as the null model, a recent study from Grobys (2023c) proposes a test where the power-law model is under the alternative allowing us to have control over the type-1 error. Grobys' (2023c) notes that Taleb (2020) argues that a common feature of fat-tailed distributions is that more observations are within one-standard deviation than predicted by the normal distribution. Specifically, a stylized fact for the normal distribution is that, as $T \longrightarrow \infty$, we would expect that a fraction of 0.6827 of the observations are within one standard deviation from the mean. The rationale of Grobys' (2023c) one-sigma test is as follows: Distributions that exhibit theoretically defined mean and variance can be standardized which allows us to compute the fraction of observation that can be expected to occur within one standard deviation from the mean, as $T \longrightarrow \infty$. It can be shown that a fraction of 0.8189 is within one standard deviation from the mean for a standardized lognormal distribution (LGN). Hence, the following test statistic can be used:

$$\lambda = \frac{\left(x_{\leq \pm 1\sigma} - m_{\leq \pm 1\sigma}\right)^2}{m_{\leq \pm 1\sigma}} + \frac{\left(x_{> \pm 1\sigma} - m_{> \pm 1\sigma}\right)^2}{m_{> \pm 1\sigma}}, \tag{11}$$

where $m_{\leq \pm 1\sigma}$ denotes the expected number of observations occurring within one standard deviation from the mean of the lognormal,

$m_{>\pm 1\sigma}$, denotes the expected number of observations exceeding one standard deviation from the mean, whereas $x_{\leq \pm 1\sigma}$ and $x_{>\pm 1\sigma}$ are corresponding values for the observed distribution. In an early study, Pearson (1900) showed that this type of test statistic is distributed as chi-square with one degree of freedom. In Table A1 in the appendix we report the corresponding results derived from one-sigma tests. We observe from Table A1 that $\widehat{\lambda}$ exceeds the critical value of $\chi^2_{0.95}(1) = 3.84$ which is the corresponding critical value using a significance level of 5 percent. The *p*-values are equal to zero for all factor variances indicating that the well-established lognormal distribution can be ruled out as underlying data-generating process. These results strongly confirm our earlier results derived from Clauset et al.'s (2009) goodness-of-fit tests.

3.5.2. Evidence from a multifractal model of asset invariances

Next, recent work from Grobys (2023b) extends Mandelbrot's (2008) multifractal model of asset returns and proposed a multifractal model of asset invariances. Grobys (2023b) shows that his model is capable of pricing various realized asset market variances (e.g., S&P 500, Bitcoin, gold, crude oil, GBP/USD foreign exchange rate). We hypothesize that if realized factor variances exhibit the same power-law behavior, they should be generated by the same multifractal model. To explore whether the Fama and French factor variances can be explained by Grobys' (2023b) multifractal model of asset invariances, we first estimate the Hurst exponents for the realized factor variances using detrended fluctuation analysis (DFA). Specifically, the DFA used in this study can be summarized as follows: First, a given data series x_t is converted to the mean-centered cumulative sum:

$$\widetilde{x}_t = \sum_{t=1}^{T} x_t \,. \tag{12}$$

In the present research context,x_t is a vector of realized factor variances. Different time scales k are defined, that is, $k \in \{4, 8, 16, 32, 64, 128, 256, 512\}$. Depending on the defined time scale, data is split into epochs and for each epoch s, a time series regression is used to detrend the data. For example, employing $k = 512$ means that monthly data for $\widetilde{x}_t$ is split into two non-overlapping epochs. For each epoch s, the following regression is employed:

$$\widetilde{x}_t = \gamma_0 + \gamma_1 t + e_t \,, \tag{13}$$

where $t = 1, \cdots, 512$ for the first epoch and $t = 513, \cdots, 1024$ for the second epoch. Then for each respective epoch s, the root mean squared error (RMSE) is computed as:

$$RMSE_S = \sqrt{\frac{1}{T_s} \sum_{t}^{T_s} \widehat{e}_t} \,, \tag{14}$$

where $T_s = 512$. Finally, the estimates for $RMSE_S$ are averaged for each time scale k, giving us $\overline{RMSE}_k$. According to theory, the following relation holds:

$$\overline{RMSE}_k = ck^H. \tag{15}$$

The corresponding Hurst exponent is then estimated by implementing a linear fit between log-scales and log-$\overline{RMSE}_k$. If the data were independent, the ratio between numerator and denominator should be, according to the theory, 1:2, corresponding to a Hurst exponent of $H = 0.50$. Moreover, $H > 0.50$ is in line with long-term dependence, that is, a long memory of the stochastic process in which the data are persistent; however, $H < 0.50$ implies antipersistence, which is characterized by the tendency to revert back. Grobys (2023b) shows that the realized asset variances for the S&P 500 and crude oil can be priced by a multifractal model derived from a multiplicative cascade using binominal bending with a probability of $p = 0.70$. Specifically, this model generates power-law exponents with expectation $\overline{\widehat{\alpha}} = 1.9187$ ($\widehat{\sigma}_\alpha = 0.2916$) and $\overline{\widehat{H}} = 0.9176$ ($\widehat{\sigma}_H = 0.0557$). In Table A2, the point estimates for estimated power-law exponents and estimated Hurst exponents are reported. We see that the estimated power-law exponents are less than 1.96 standard deviations from $\alpha = 1.9187$ suggesting that the realized factor variances are in line with predictions generated from the multifractal model. While Hurst exponents suggest an even higher level of persistence than predicted from the multifractal model, a robustness check in Grobys (2023b) provides evidence for that for smaller samples, Hurst exponents corresponding to $H \approx 1.20$ are still in the support of the multifractal model. Overall, the evidence suggests the presence of a common underlying mechanism that generates the power-law behavior which we empirically observe for the realized Fama and French factor variances.

3.5.3. Evidence from co-fractality analysis

To which extent do realized factor variances co-move? Note that a common power-law exponent governing realized factor variances does not necessarily imply that factor variances are in power-law regimes at the same time. Since correlations are undefined if variances are infinite, co-dependencies between realized factor variances can be investigated using Grobys's (2023a) recently proposed concept of co-fractality.[10] Since co-fractality in its strong form is perhaps less relevant for empirical data, we focus here solely on co-fractality in its weak form. Specifically, let us define n_1 and n_2 as the number of observations that are governed by power laws for

[10] The concept of co-fractality is related to the concept of tail dependence and discussed in detail in Grobys' (2023a) study.

processes $x_{i,t}$ and $x_{j,t}$. Then, the (weak) co-fractality statistic, defined as λ^{CF}, is $\lambda^{CF} = 0$ if for two processes $x_{i,t}$ and $x_{j,t}$, none of the power-law observations coincide, whereas $\lambda^{CF} = 1$ if all of the observations (e.g., $MIN\{n_1, n_2\}$) which are governed by power laws coincide.[11] Self-evidently, as $\lambda^{CF} \to 1$, there is not much room for risk diversification because power-law behavior coincides; in other words, extreme events are likely to occur within the same parts of time series paths.

To explore co-fractality for realized factor variances, we consider the sample from July 1963 to September 2022 because we do not have observation for the realized variances of RMW and CMA in the ex-ante July 1963 period. As a result, our sample consists of 711 monthly observations. The estimates $\widehat{\lambda}^{CF}$ are reported in Table A3 in the appendix. We see that the number of sample observations which are governed by a power-law process varies between 77 months for the realized variance of the MKT factor and 456 for the realized variance of the RMW factor. Next, we observe from Table A3 that for 7 out of 15 realized factor variance pairs, $\widehat{\lambda}^{CF} > 0.90$, whereas for 11 out of 15 $\widehat{\lambda}^{CF} > 0.80$, suggesting a high level of (weak) co-fractality among realized factor variances. Overall, our results suggest a high level of co-movement in the tails of the realized factor variances.

This commonality aligns well with the theory of the market model. According to this model, the return of an asset or portfolio can be expressed as a function of market return and a unique idiosyncratic component, formulated as:

$$R_t = \alpha + \beta M_t + \varepsilon_t \,. \tag{16}$$

When applying this model to SMB (Small Minus Big) factor returns, we get:

$$SMB_t = (\alpha_s - \alpha_B) + (\beta_s - \beta_B) M_t + \left(\varepsilon_{s,t} - \varepsilon_{B,t}\right), \tag{17}$$

where S is the portfolio of small stocks and B is the portfolio of big stocks. Consequently, the variance of SMB is expressed as:

$$Var(SMB_t) = \underbrace{(\beta_s - \beta_B)^2}_{>0} var(M_t) + var(\varepsilon_s - \varepsilon_B). \tag{18}$$

This suggests that the SMB factor's variance is essentially a scaled version of market variance—influenced by the difference in market sensitivities (betas) of small and big stocks, alongside the variance of differential idiosyncratic terms. The correspondence of our empirical findings with the market model implies a deeper, systematic behavior in factor variances, suggesting that the commonality observed in factor variances—including that of all factors analyzed—is not an anomaly but rather a reflection of fundamental market dynamics, as captured by the market model. This insight highlights the significance of the observed common component in factor variance risk. However, it is essential to emphasize that this model is applicable only for values of $x < x_{min}$ because variance is undefined for $x \geq x_{min}$. The interesting thing is here that it turns out that even the observations $x \geq x_{min}$ exhibit a high level of co-dependence measured in terms of co-fractality.[12]

3.5.4. *Are estimated power-law exponents reliable? Evidence from sample-split tests*

Next, one might wonder whether the results reported in this study are subject to sample-specificity themselves, which is, of course, a valid concern. To explore this issue, we split each subsample into two non-overlapping subsamples of equal length. Specifically, the first subsample spans from July 1926 to October 1972 for the realized MKT, SMB, and HML realized variances, July 1963 to February1993 for the realized RMW and CMA realized variances, and November 1926 to September 1974 for the realized MOM factor variance.

The second subsample spans from November 1972 to September 2022 for the realized MKT SMB, and HML realized variances, March 1993 to September 2022 for the realized RMW and CMA realized variances, and October 1974 to September 2022 for the realized MOM factor variance. For each subsample and realized factor variance, we implement the blocks bootstrap procedure as detailed in section 3.3. The descriptive statistics for the estimated power-law exponents are reported in Panels A and B of Table A4 in the appendix, whereas Panel C reports the results from z-tests for testing parameter equality. The two-sample z-tests are given by,

$$\widehat{z} = \frac{(\widehat{\alpha}_{1i} - \widehat{\alpha}_{2i})}{\sqrt{\left(0.5\widehat{\sigma}^2_{\widehat{\alpha}_{1i}} + 0.5\widehat{\sigma}^2_{\widehat{\alpha}_{2i}}\right)}} \,, \tag{19}$$

where $\widehat{\alpha}_{1i}$ ($\widehat{\alpha}_{2i}$) denotes the estimate of α for sample 1 (sample 2) and factor variance i, and $\widehat{\sigma}^2_{\widehat{\alpha}_{1i}}$ ($\widehat{\sigma}^2_{\widehat{\alpha}_{2i}}$) denotes the corresponding estimated sample variance obtained from blocks bootstraps. Using a standard significance level of 5 percent, we observe from Panel C of Table A4 that for five out of six realized factor variances, we cannot reject the null hypothesis that the point estimates $\widehat{\alpha}_{1i}$ and $\widehat{\alpha}_{2i}$ are statistically the same. The realized SMB variance appears to be the only exception here. Overall, our results suggest that power-law exponents exhibit a high level of reliability, which appears to be contrary to GARCH-type models because Mandelbrot (2008)

[11] Recall that observations $x_{i,t} \geq x_{MIN,i,t}$ and $x_{j,t} \geq x_{MIN,j,t}$ are governed by power-law processes.

[12] Note that we use the market model only to show that the market model predicts co-movement of factor variances too; however, the market model is not designed for non-linearities. Since power-law behavior strongly indicate non-linearities, the tails of the realized variance distributions need to be investigated separately using, for instance, co-fractality analysis.

criticized the problem of changing parameters retrieved from GARCH-type models, arguing that "… many recent models of price variation try to explain the obviously shifting pattern of volatility by inserting parameters that change by the day, hour, and second; such are the GARCH family mentioned earlier." Our evidence suggests that unlike GARCH-type models, power-law models work with just a set of a few consistent parameters that typically remain constant over time and place. This is also in line with Calvet and Fisher (2004) or Lux, Morales-Arias, and Sattarhoff (2014), who found that power-law models usually outperform GARCH models in this regard.

3.5.5. Does the tail index change after controlling for autocorrelation in realized variances?

A reader could be concerned about the existence of a genuine power law in the factor variance distributions. It is well-known that time-varying volatility can induce power law-like behavior in the tails of asset returns. That implies that common statistical methods for estimating power laws will erroneously produce estimates indicating a power law when there is none; that is, the tails are thick but it is due to the fact that they are generated by a mixture of exponential-tail distributions rather than a genuine power law. Thus, a natural question that arises is whether the same problem arises for estimating the power-law behavior of realized factor variances. To explore this issue, we estimate autoregressive models of order p for each realized factor variance such as:

$$\widehat{\sigma}^2_{i,t} = \beta_{0,i} + \beta_{1,i}\widehat{\sigma}^2_{i,t-1} + \beta_{2,i}\widehat{\sigma}^2_{i,t-2} + \cdots + \beta_{p,i}\widehat{\sigma}^2_{i,t-p} + \varepsilon_{i,t}\,,$$

where $\varepsilon_{i,t}$ is the innovation process, p is the lag order, and $i \in \{MKT, SMB, HML, RMW, CMA, MOM\}$. The optimal lag-order p is obtained by inspection of the partial autocorrelation function.[13] Note that this type of model can be considered a version of an ARCH-type model where the equation modeling the conditional variance is modeled directly by using the own lags of the realized factor variances. Then, we use the absolute amount of the innovation processes, $|\widehat{\varepsilon}_{MKT}|$, $|\widehat{\varepsilon}_{SMB}|$, …, $|\widehat{\varepsilon}_{MOM}|$, and estimate the power-law exponents and implement Clauset et al.'s (2009) goodness-of-fit tests using the procedure outlined in section 3.

In Table 6 we report the point estimates for estimated autoregressive models of realized Fama and French factor variances. The results from Table 6 show that there are strong patterns of autocorrelation in the data-generating processes of realized factor variances. The order of autocorrelation varies between $p = 2$ (e.g., $\widehat{\sigma}^2_{MOM,t}$) and $p = 4$ (e.g., $\widehat{\sigma}^2_{SMB,t}$). The descriptive statistics of the innovation processes are reported in Table 7. From Table 7 it becomes evident that even after controlling for higher order autocorrelation, the innovation processes of realized factor variances exhibit extremely heavy tails, as indicated by kurtosis values varying between 34.11 (e.g., $\widehat{\varepsilon}_{HML}$) and 111.40 (e.g., $\widehat{\varepsilon}_{MKT}$), respectively. Finally, in Table 8 we report the estimated power-law exponents for the innovation processes of realized factor variances using the absolute amount, $|\widehat{\varepsilon}_{MKT}|$, $|\widehat{\varepsilon}_{SMB}|$, …, $|\widehat{\varepsilon}_{MOM}|$. Strikingly, the estimated power-law exponents are virtually the same which strongly supports our earlier finding of genuine power-law behavior.[14] Also, Clauset et al.'s (2009) GoF tests do not reject the power-law null models for the vast majority of realized factor variance innovations which further strengthens our main findings.

3.5.6. Are our results subject to discretization errors?

To examine the impact of different aggregation periods on the reliability of our findings suggesting power-law behavior of realized factor variances, we conduct a robustness check by using quarterly realized variance data derived from squared daily return data. This approach allows us to address potential concerns related to discretization errors that may influence the tail properties of the realized factor variance distributions across varying data frequencies. Thus, instead of using the 22-day window of non-overlapping squared daily returns to compute monthly realized variance, we employ a 66-day window of non-overlapping squared daily returns to calculate quarterly realized variance data.

The summary statistics for the quarterly realized variance are reported in Table A5. in the appendix. The results show that quarterly data exhibits similar statistical characteristics like monthly data. From Table A5, the mean values indicate that—as with the monthly realized factor variances—the MKT realized variance exhibits the highest sample average, followed by momentum MOM, HML, SMB, RMW, and CMA realized variances. Similarly, the median values reinforce this pattern, with MKT realized variance showing the highest median value and CMA realized variance the lowest. Additionally, the maximum and minimum values reveal a broad range of realized variances across all factors, with the most significant variation observed for MKT realized variance. The standard deviation is highest for realized MKT realized variance and lowest for realized CMA realized variance. All realized factor variances display extremely high kurtosis values, ranging from 17.12 for HML realized variance to 39.62 for RMW realized variance, which indicates the presence of exceptionally fat tails in the distribution. Moreover, the skewness is most pronounced for the realized RMW realized variance, suggesting notable asymmetry in its distribution.

Next, we estimate the power-law exponents for the quarterly realized variance data. The results are reported in Panel A of Table A6. Similar to the monthly data reported in Table 2, the tail exponent for all quarterly realized factor variances is estimated to be $\widehat{\alpha} < 3$ for all realized factor variances, indicating that the variance of factor variance does not exist for any of the factor's quarterly variances.

[13] Results are unreported to save space but are available upon request.

[14] As an example, let us consider the estimated power-law exponent for the market factor variance which is, according, to Table 2, $\widehat{\alpha} = 2.41$ with a corresponding 95% confidence interval $\alpha \in [1.94, 2.88]$ because the robust estimate for the standard deviation is, as reported in Table 4, $\widehat{\sigma} = 0.24$. The point estimate for the power-law exponent for the innovations process of the market factor variance is, according to Table 8, $\widehat{\alpha} = 2.19$ which is clearly within the 95% confidence interval for the power-law exponent for the market factor variance. That is, statistically, the point estimates are not different from each other. The same argument holds for the power-law exponents of the remaining realized factor variances too.

Table 6
Point estimates for autoregressive models of realized Fama and French factor variances.

	$\widehat{\beta}_{0,i}$	$\widehat{\beta}_{1,i}$	$\widehat{\beta}_{2,i}$	$\widehat{\beta}_{3,i}$	$\widehat{\beta}_{4,i}$	R-square
$\widehat{\sigma}^2_{MKT,t}$	9.02***	0.48***	0.02	0.13***		0.31
	(6.17)	(16.39)	(0.61)	(4.55)		
$\widehat{\sigma}^2_{SMB,t}$	1.52***	0.34***	0.15***	0.15***	0.15***	0.43
	(3.60)	(11.71)	(5.12)	(4.91)	(5.01)	
$\widehat{\sigma}^2_{HML,t}$	1.18***	0.44***	0.15***	0.27***		0.59
	(3.36)	(15.45)	(4.79)	(9.40)		
$\widehat{\sigma}^2_{RMW,t}$	0.63***	0.60***	0.02	0.18***		0.54
	(3.56)	(16.23)	(0.36)	(4.87)		
$\widehat{\sigma}^2_{CMA,t}$	0.44***	0.35***	0.18***	0.32***		0.55
	(3.24)	(9.64)	(4.89)	(8.83)		
$\widehat{\sigma}^2_{MOM,t}$	4.43***	0.38***	0.27***			0.32
	(5.96)	(13.34)	(9.33)			

*** Statistically significant on a 1 % level

Data used to compute realized Fama and French factor variances are collected from the Data Library of Kenneth R. French and cover the excess market return (MKT), size factor (SMB), value factor (HML), profitability factor (RMW), investment factor (CMA), and momentum factor (MOM). The first three factors are the original Fama and French three factors model (Fama and French, 1993), while RMW and CMA are from the original Fama and French five-factor model (Fama and French, 2015). The MOM factor represents the momentum factor as proposed by Carhart (1997). The sample period for the MKT, SMB, and HML realized variance data spans the period July 1926 to September 2022 (1,155 monthly observations). The RMW and CMA realized variance data cover the period from July 1963 to September 2022 (711 monthly observations), whereas the MOM factor variance data encompasses the period from November 1926 to September 2022 (1,151 monthly observations). The innovation process of realized Fama and French factor variances is estimated by running autoregressive models of order p: $\widehat{\sigma}^2_{i,t} = \beta_{0,i} + \beta_{1,i}\widehat{\sigma}^2_{i,t-1} + \beta_{2,i}\widehat{\sigma}^2_{i,t-2} + \cdots + \beta_{p,i}\widehat{\sigma}^2_{i,t-p} + \varepsilon_{i,t}$, where $\varepsilon_{i,t}$ is the innovation process, p is the lag order, and $i \in \{MKT, SMB, HML, RMW, CMA, MOM\}$. The optimal lag-order p is obtained by inspection of the partial autocorrelation function. This table reports the point estimates of the regressions as well as the coefficients of determination (R-square). The *t*-statistics are given in parentheses.

Table 7
Summary statistics for the innovations of realized Fama and French factor variances.

	$\widehat{\varepsilon}_{MKT}$	$\widehat{\varepsilon}_{SMB}$	$\widehat{\varepsilon}_{HML}$	$\widehat{\varepsilon}_{RMW}$	$\widehat{\varepsilon}_{CMA}$	$\widehat{\varepsilon}_{MOM}$
Mean	0.00	0.00	0.00	0.00	0.00	0.00
Median	−6.58	−1.30	−1.04	−0.45	−0.35	−3.50
Maximum	686.67	201.29	88.07	44.54	26.90	376.10
Minimum	−217.76	−63.99	−60.95	−27.50	−20.39	−151.42
Std. Dev	42.58	12.58	10.44	4.10	2.94	22.28
Skewness	8.39	7.60	3.72	4.92	3.31	7.06
Kurtosis	111.40	97.73	34.11	54.17	34.86	101.42
Observations	1,152	1,151	1,152	708	708	1,149

Data used to compute realized Fama and French factor variances are collected from the Data Library of Kenneth R. French and cover the excess market return (MKT), size factor (SMB), value factor (HML), profitability factor (RMW), investment factor (CMA), and momentum factor (MOM). The first three factors are the original Fama and French three factors model (Fama and French, 1993), while RMW and CMA are from the original Fama and French five-factor model (Fama and French, 2015). The MOM factor represents the momentum factor as proposed by Carhart (1997). The sample period for the MKT, SMB, and HML realized variance data (e.g., $\widehat{\sigma}^2_{MKT}$, $\widehat{\sigma}^2_{SMB}$, $\widehat{\sigma}^2_{HML}$) spans the period July 1926 to September 2022 (1,155 monthly observations). The RMW and CMA realized variance data (e.g., $\widehat{\sigma}^2_{RMW}$, $\widehat{\sigma}^2_{CMA}$) cover the period from July 1963 to September 2022 (711 monthly observations), whereas the MOM factor variance data (e.g., $\widehat{\sigma}^2_{MOM}$) encompass the period from November 1926 to September 2022 (1,151 monthly observations). The innovation process of realized Fama and French factor variances is estimated by running autoregressive models of order p: $\widehat{\sigma}^2_{i,t} = \beta_{0,i} + \beta_{1,i}\widehat{\sigma}^2_{i,t-1} + \beta_{2,i}\widehat{\sigma}^2_{i,t-2} + \cdots + \beta_{p,i}\widehat{\sigma}^2_{i,t-p} + \varepsilon_{i,t}$, where $\varepsilon_{i,t}$ is the innovation process, p is the lag-order, and $i \in \{MKT, SMB, HML, RMW, CMA, MOM\}$. The optimal lag-order p is obtained by inspection of the partial autocorrelation function. This table reports the descriptive statistics for the innovation process $\varepsilon_{i,t}$.

Additionally, Panel B of Table A6 reports the 95 percent confidence intervals for the estimated power-law exponents for the monthly realized variance data. The analysis shows that the point estimates for the estimated power-law exponents derived from quarterly realized variance data consistently fall within the confidence intervals derived from monthly data. This consistency across different sampling frequencies strengthens the reliability of our findings and supports the hypothesis of invariance in the relationship between realized variances at different data frequencies.

Finally, we perform Clauset et al.'s (2009) goodness-of-fit tests to the quarterly realized variance data. The results are reported in Table A7. The second column of Table A7 reports the *p*-value for testing the power-law null hypothesis. Similar to the findings derived from monthly data, the *p*-values for all quarterly realized factor variances exceed 5 percent. This not only provides strong evidence

Table 8
Estimating power law exponents for the innovation processes of realized factor variances.

	$\lvert\widehat{\varepsilon}_{MKT}\rvert$	$\lvert\widehat{\varepsilon}_{SMB}\rvert$	$\lvert\widehat{\varepsilon}_{HML}\rvert$	$\lvert\widehat{\varepsilon}_{RMW}\rvert$	$\lvert\widehat{\varepsilon}_{CMA}\rvert$	$\lvert\widehat{\varepsilon}_{MOM}\rvert$
$\widehat{\alpha}$	2.19	2.10	1.92	2.15	2.41	2.13
x_{MIN}	11.78	1.36	1.12	0.69	1.72	6.38
p-value (GoF)	0.24	0.05	0.02	0.73	0.14	0.05

This table reports the results of estimating power law exponents for the absolute amount of the innovation processes of the following realized factor variances: Excess market factor (MKT), size factor (SMB), value factor (HML), profitability factor (RMW), investment factor (CMA), and momentum factor (MOM). The $\widehat{\alpha}$ represents the estimated tail exponent and $\widehat{\sigma}$ is the estimated standard deviation. The lower threshold x_{MIN} is estimated using the optimized Kolmogorov-Smirnov distance (D) distance as proposed by Clauset et al. (2009). The estimated x_{MIN} is the value that corresponds to the optimal distance D. Additionally, for each factor variance, the fraction of observations governed by power law process is reported. The sample period for the MKT, SMB, and HML realized variances span the period July 1926 to September 2022 (1,155 monthly observations). The RMW and CMA realized variances cover the period from July 1963 to September 2022 (711 monthly observations), whereas the MOM factor variance encompasses the period from November 1926 to September 2022 (1,151monthly observations). The innovation process of factor variances is obtained by using the following autoregressive model: $\widehat{\sigma}^2_{i,t} = \beta_{0,i} + \beta_{1,i}\widehat{\sigma}^2_{i,t-1} + \beta_{2,i}\widehat{\sigma}^2_{i,t-2} + \cdots + \beta_{p,i}\widehat{\sigma}^2_{i,t-p} + \varepsilon_{i,t}$, where $\varepsilon_{i,t}$ is the innovation process, p is the lag-order, and $i \in \{MKT, SMB, HML, RMW, CMA, MOM\}$. The optimal lag-order p is obtained by inspection of the partial autocorrelation function. This table reports the estimated power-law exponent $\widehat{\alpha}$, the corresponding cutoff x_{MIN}, and the p-value for the goodness-of-fit (GoF) test outlined in section 3. To implement the GoF tests we use 2500 simulated data sets.

supporting the plausibility of the power-law null models but also underscores the robustness of our results across different time frequencies; that is, the distributional tail characteristics of realized factor variances are the same regardless of the aggregation period.

In addition, according to the results reported in column 3 of Table A7, the null hypothesis of the lognormal distribution is rejected for all realized factor variances except for SMB realized variance and MOM realized variance. However, according to Clauset et al. (2009), comparing p-values for the power law along with other possible distributions (here: lognormal or exponential) can provide a persuasive argument concerning the suitability of the power-law model. Explicitly, Clauset et al. (2009) stress that a higher p-value for the power law compared to other competing models, effectively eliminates the alternatives distributions. Thus, as reported in Table A7, the notably higher p-values for power laws relative to those for the lognormal distribution for SMB realized variance and MOM realized variance suggest that the power-law model more accurately captures the behavior of these specific realized factor variances too. Furthermore, from column 4 of Table A7 we observe that—consistent with the results retrieved from the analysis of monthly data—the p-values are zero or near to zero in all cases. This indicates that the exponential distribution does not either offer an adequate fit for quarterly factor variance data.

4. Discussion

4.1. Commonalities and contradictions with other studies

Using daily data, earlier studies found that the realized variances of the S&P 500, crude oil, gold, Bitcoin and foreign exchange rates are governed by power-law processes (Grobys, 2021; 2023a). The estimated tail exponents varied in their economic magnitude but were usually $\widehat{\alpha} > 2$ implying that that the theoretical mean of realized variances exist. Using monthly data, we test whether the realized variances of the Fama and French factors exhibit finite theoretical means. Statistically, we could not reject the infinite theoretical mean hypothesis for any of the realized equity factor variance. Given that Lux and Alfarano (2016) argue that the pertinent literature converged to the insight that financial asset returns exhibit power-law behavior in line finite variances, this is, perhaps a surprising finding; However, our results are in line with Mandelbrot's (1963b) early study, which used cotton price changes as a case study to demonstrate that returns on some speculative financial assets do not exhibit a theoretically defined variance.

One might wonder if a potential reason could be that we explicitly analyze realized variances as opposed to the absolute amount of asset returns. Note that our approach to investigate realized variances is motivated by well-established earlier literature arguing that realized variances contain information that standard models, which employ simple asset returns, cannot reveal (e.g., Bubák, Kočenda, and Žikeš, 2011; Andersen, Bollerslev, Diebold, and Labys, 2003; Andersen, Bollerslev, and Meddahi, 2004). It is interesting to note that a recent study of Grobys (2023c) uses the absolute amount of S&P 500 and Bitcoin returns covering the period July 28, 2013, until May 14, 2023. Using weekly data, that study cannot either reject the infinite variance hypothesis for both Bitcoin and S&P 500 returns. Therefore, our study is in line with these recent findings.

4.2. Power-law or lognormal?

Next, there is a well-established literature arguing that realized asset volatility is close to lognormally distributed (e.g., Andersen, Bollerslev, Diebold, & Ebens, 2001; Andersen, Bollerslev, Diebold, and Labys, 2001a & 2001b). While the studies of Grobys (2021; 2023a) do not explicitly test the lognormal distribution against proposed power laws, we explicitly test this issue for realized equity factor variances. Using Clauset et al.'s (2009) goodness-of-fit test our findings strongly suggest that the power-law null model cannot be rejected for any realized factor variance. However, one could perhaps argue that this type of test is often considered weak because it

imposes the power-law model under the null hypothesis. Therefore, we employ Grobys' (2023c) recently proposed one-sigma test where the lognormal distribution is the null model and the power-law model forms the alternative. Grobys' (2023c) one-sigma test strongly rejects the lognormal distribution, as indicated by p-values that are $p = 0.00$ for all tests carried out. Overall, our findings provide strong evidence for significant power-law behavior of realized factor variances. We believe, however, that this should not come as a great surprise, given that Paretian tails are considered stylized facts of financial asset returns. We argue that lognormally distributed variances are an illusion which may result in severe consequences: While the lognormal distribution has defined first and second moment, our analysis reveals that the underlying power-law distributions do not exhibit theoretically defined variances. Hence, we argue that the often-used lognormal model will inevitably underestimate (realized) factor risks by a substantial margin.

4.3. Implications for factor premiums and the validity of earlier research

Furthermore, the study of Grobys and Kolari (2022) found that the popular Fama and French three-factor model did not outperform the CAPM in North America, Europe, and Japan. This finding indicates that size and value factor premiums do not exist in other markets than the U.S. Are only U.S. investor compensated for risk? While Fama and French (2015) concluded that their earlier proposed value factor is redundant after accounting for the investment factor, a recent study of Cakici et al. (2023) uses new data from U.S. and global markets to re-examine market risk premium predictability by using equity anomalies. Employing machine learning, their study shows that (a) anomalies cannot predict aggregate market returns and (b) any ostensible evidence from the U.S. lacks external validity because it cannot be extended internationally nor does it hold for alternative anomaly sets. Moreover, other studies do not even find any evidence for the existence of significant factor premiums in extended samples (Dichev, 1998; Chan, Karceski, and Lakonishok, 2000; Hirshleifer, 2001; Schwert, 2003; Smith and Timmermann, 2022; Chen and Velikov, 2023).

Given our finding of undefined second moments, the finding of sample-specific factor premiums should not come as a great surprise: Note that Fama (1963) in his early study already pointed out that results derived from standard statistical methods based on the concept of correlation will provide very misleading answers if the variance does not exhibit a theoretically defined mean. Therefore, "statistically significant factor premiums", documented in abundance in the finance literature appear to be sample-specific phenomena, and therefore, inappropriate means for predicting returns.

A reader might wonder what are the practical implications of sample-specificity or how likely is it that two independent studies will find different results even though the data generating processes are the same? It is important to note that the most essential requirement for a factor is (a) the existence of positive factor premium which (b) should be statistically significant on at least a 5 percent level. Once these requirements are satisfied, a potential factor could be used for statistical tests. In this regard, it is interesting to recall that, in their seminal 1992 and 1993 studies, Fama and French proposed the size and value factor because the (excess) market factor alone failed to describe historical average stock returns. As mentioned earlier, the results form subsequent research cast doubts on the existence of the size premium in expanded samples (i.e., Dichev, 1998; Chan et al., 2000; Hirshleifer, 2001; Schwert, 2003). Performing international tests of the Fama and French (1992, 1993) three-factor model, Grobys and Kolari (2022) show that neither the size factor nor value factor matter for describing the cross section of expected stock returns. In view of this evidence, we employ the size factor as factor that is supposedly *not* priced; that is, a factor that does not meet the previously defined factor requirements. In fact, over the July 1963 to September 2022 sample, which used in the present study, the sample average of the size factor corresponded to 0.19 % per month with a t-statistic of 1.51 indicating that the size premium is statistically not different from zero.

Therefore, to explore the consequences of sample-specificity, we retrieve the size factor over the July 1963 to September 2022 sample from Kenneth French's website and employ blocks bootstraps, as described in section 3.3., to simulate 5,000 size factor evolutions.[15] For each synthetic size factor, we compute the factor premium (e.g., sample average) and the corresponding t-statistic. We plot these two metrics in Fig. 5.[16] We observe from Fig. 5 that even though the original size premium is statistically not different from zero, in 36.50 percent of the synthetic size factor samples, the size premium appears to be 'statistically significant' as indicated by t-statistics $> |1.96|$. Specifically, 1,824-out-of-5,000 synthetic samples exhibit size premiums of at least 0.18 % per month with corresponding t-statistics $> |1.96|$. Using one-sided tests, our findings indicate that in 47.70 percent of the synthetic size factor samples, the size premium appears to be 'statistically significant' as indicated by t-statistics > 1.65.

Next, we use trial-and-error to generate a data series of normally distributed random variables that exhibit sample t-statistic corresponding to 1.51—which is the same value of the t-statistic that we empirically observe for the size factor over our sample. Again, we use the same blocks bootstraps procedure as described above to generate synthetic samples and explore how many synthetic samples exhibit t-statistics $> |1.96|$ or > 1.65, respectively. Our findings (unreported) indicate that 26.30 percent of the synthetic samples exhibit t-statistics $> |1.96|$, whereas 34.22 percent of the synthetic samples exhibit t-statistics > 1.65. Overall, our findings show that in the presence of an infinite variance process, standard statistical inference is severely distorted.

Future research could employ simulations to study the behavior of the standard Fama and French methodology and predictive regressions under the assumptions that the variances of Fama and French factors do not exist and hence standard OLS inference does not apply.[17] This is, however, outside the scope of this study and therefore left for future research.

[15] Note that we implemented the blocks bootstrap using a block length of 60 months to account for unknown dependency structures in the data. For instance, Asness, Moskowitz, and Pedersen (2013) document return dependencies up to five years.

[16] The descriptive statistics are reported in Tables A.8 in the appendix.

[17] For instance, Cont and Tankov (2004) discuss some data generating processes that can be parametrized to either have or not have a finite first moment (α-stable processes).

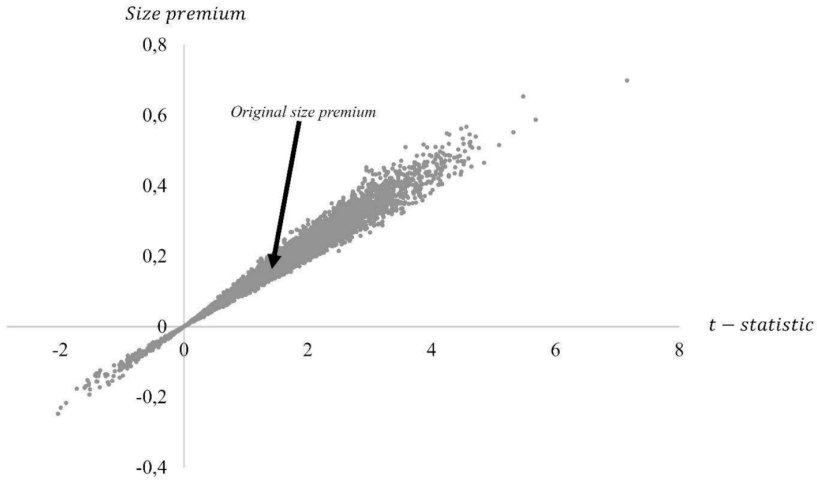

Fig. 5. Synthetic size factor premiums and statistical significances. We retrieve monthly data on Fama and French's size factor over the July 1963 to September 2022 sample from Kenneth French's website and employ blocks bootstraps, as described in section 3.3., to simulate 5,000 size factor evolutions. For each synthetic size factor, we compute the sample average and the corresponding t-statistic and plot these two metrics in this figure. The blocks bootstraps procedure uses an expected block length of 60 months.

4.4. Limitations and future research

It is important to note that our study has some limitations. First, in our study, we follow Merton (1980) in computing realized factor variances. Future studies are encouraged to explore this issue using other approaches to compute realized variances. Moreover, our study explicitly analysis monthly data. However, one could have the view that using intraday data would perhaps result in less noisy estimates for realized factor variances. On the other hand, it is important to note that high frequency data could be polluted with a substantially higher level of noise. Mandelbrot (2008) highlights that high frequency data may suffer from microstructure issues. In analyzing the market for foreign exchange rates, Mandelbrot (2008) argues that data having a higher frequency than two hours or a lower frequency that 180 days are subject to crossovers; that is, points where the mathematical relation does not appear to hold anymore. Also, Mandelbrot (2008) points out that crossovers are common for real as opposed to theoretical fractal data. While this study uses monthly data as monthly variance data is typically used in finance studies (e.g., Moreira and Muir, 2017; Grobys, Ruotsalainen and Äijö, 2018; Grobys and Vähämaa, 2020), future research could extend our study and explore higher or lower frequented factor variance data. Finally, we focus explicitly on exploring the realized variance risk of the well-known Fama and French equity factors. While these factors correspond only to a small subset of potential factor candidates (Feng, Giglio, & Xiu, 2020), one could investigate this issue for other long-short portfolios. This is, however, outside the scope of this paper and therefore left for future research.

5. Conclusion

While earlier studies focused on examining the premiums of equity market factors, this is the first study that explored the realized variances of the well-established Fama and French equity factors. Specifically, we investigated whether realized factor variances would exhibit power-law behavior and if so, to which degree. Contrary to the earlier literature our findings indicated that power laws more appropriately describe the data-generating processes than the time-honored lognormal distribution. The estimated power-law exponents were relatively low; that is, $\alpha \approx 2$ for all realized factor variances. We tested the infinite theoretical mean hypothesis and found that we could not reject this hypothesis. Again, this finding was contrary to the pertinent literature arguing that the theoretical variance exists for most financial assets. In line with earlier literature, we argued that realized variances more reliably reflect the dynamics of asset variance processes and therefore our results would deviate from results retrieved form estimated based on plain returns.

Furthermore, a multifractal model of asset invariances is capable of describing the evolution of the observed realized variances which appear to be generated by the same underlying mechanism manifested in a common power-law exponent $\alpha \approx 2$ resulting from a multifractal cascade based on binominal bending with $p = 0.70$. Our findings indicated that risk diversification is limited because extreme events occurring in the cross section of realized variances tend to coincide. Our findings are important for at least two reasons: From a practical point of view, our findings call for re-evaluations of financial reward-to-risk measures often used for evaluating investment strategies following factor investing which may suffer from sample-specificity. Knowing that often-used metrics such as

Sharpe ratios are misleading in the presence of infinite variances, new methodologies are required. Given that this industry has about $5 trillion asset under management this is not a trivial issue. From a theoretical point of view, our findings provide evidence that standard econometric methods often used for decision making provide a poor foundation for decision making because they deliver very misleading results due to sample-specificity. In this regard, our results call for a careful re-evaluation of econometric methods used for financial decision making.

Extremely fat-tailed data—which is something that we observe for the Fama and French factor variances—contain extreme observations that can significantly influence the results of statistical analysis. However, traditional parametric methods, which typically assume a normal distribution, may be highly sensitive to these extreme events, eventually resulting in biased estimates or misleading inferences (Fama, 1963). Taleb (2010) has sparked growing interest in non-parametric methods which do not rely on specific distributional assumptions. Such non-parametric methods are less affected by outliers since they do not depend on particular distributional premises (Wasserman, 2006). Results derived from statistical tests using those methodologies provide robustness even when confronted with extreme observations (Huber, 1981). Additionally, given that non-parametric methods are not bounded by distributional assumptions, the resulting estimates and inferences are often more intuitive and straightforward (Efron and Tibshirani, 1994). Moreover, Mandelbrot (2008) underscores the importance of alternative models that can capture the complexities of financial data more robustly. Further, Mandelbrot (1982) introduced the concept of fractal geometry, explaining how financial data often exhibits fractal-like patterns—an inspiration that has encouraged many researchers to seek alternative models for financial data representation. Though the potential of these novel methodologies is evident, a detailed investigation of their applicability and implications lies beyond the scope of this study.

CRediT authorship contribution statement

Masoumeh Fathi: Methodology, Investigation, Writing – review & editing. **Klaus Grobys:** Writing – review & editing, Writing – original draft, Validation, Supervision, Methodology, Investigation, Formal analysis. **Janne Äijö:** Writing – review & editing, Visualization, Supervision.

Appendix

Table A1
Results derived from one-sigma tests.

	$\widehat{\sigma}^2_{MKT}$	$\widehat{\sigma}^2_{SMB}$	$\widehat{\sigma}^2_{HML}$	$\widehat{\sigma}^2_{RMW}$	$\widehat{\sigma}^2_{CMA}$	$\widehat{\sigma}^2_{MOM}$
$prob(x \leq \pm 1\sigma)$	0.9498	0.9550	0.9247	0.9325	0.9241	0.9305
$\widehat{\lambda}$	133.44	144.25	87.17	61.87	53.06	96.58
(p-value)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)

This table reports the results of one-sigma tests as detailed in section 3.5.1. Note that a fraction of 0.8189 is within one standard deviation from the mean for a standardized lognormal distribution (LGN). The reference distribution for the estimated test statistic $\widehat{\lambda}$ is under the null hypothesis distributed as $\chi^2(1)$.

Table A2
Pricing realized factor variances using a multifractal model of asset invariances.

	$\widehat{\sigma}^2_{MKT}$	$\widehat{\sigma}^2_{SMB}$	$\widehat{\sigma}^2_{HML}$	$\widehat{\sigma}^2_{RMW}$	$\widehat{\sigma}^2_{CMA}$	$\widehat{\sigma}^2_{MOM}$
$\widehat{H}$	1.0770	1.1174	1.1838	1.1785	1.1273	1.0912
$\widehat{\sigma}_H$	0.0192	0.0156	0.0214	0.0296	0.0216	0.0245
$\widehat{\alpha}$	2.25	2.22	2.01	2.23	2.49	2.21

This table reports the estimated Hurst exponents using detrended fluctuation analysis. The standard errors are obtained from log-log regressions as detailed in section 3.5.2.

Table A3
Co-fractality matrix for realized factor variances.

	$\widehat{\sigma}^2_{MKT}$	$\widehat{\sigma}^2_{SMB}$	$\widehat{\sigma}^2_{HML}$	$\widehat{\sigma}^2_{RMW}$	$\widehat{\sigma}^2_{CMA}$	$\widehat{\sigma}^2_{MOM}$
#	77	393	272	456	204	149
$\widehat{\sigma}^2_{MKT}$	1	0.96***	0.94***	0.99***	0.66***	0.79***
$\widehat{\sigma}^2_{SMB}$		1	0.82***	0.82***	0.82***	0.91***
$\widehat{\sigma}^2_{HML}$			1	0.94***	0.79***	0.86***

(*continued on next page*)

Table A3 (*continued*)

	$\widehat{\sigma}^2_{MKT}$	$\widehat{\sigma}^2_{SMB}$	$\widehat{\sigma}^2_{HML}$	$\widehat{\sigma}^2_{RMW}$	$\widehat{\sigma}^2_{CMA}$	$\widehat{\sigma}^2_{MOM}$
$\widehat{\sigma}^2_{RMW}$				1	0.92***	0.98***
$\widehat{\sigma}^2_{CMA}$					1	0.60***
$\widehat{\sigma}^2_{MOM}$						1

*** Statistically significant on a 1 % level.

This table reports the estimated co-fractality matrix using the concept of weak co-fractality as detailed in section 3.5.3. The sample is from July 1963 to September 2022 comprising 711 observations. The symbol # denotes the count of sample observations which are governed by a power-law process.

Table A4

Summary statistics for blocks bootstrapped power-law exponents using sample split tests.

Panel A. Descriptive statistics for the bootstrapped tail exponents for the earlier subsample.						
	$\widehat{\alpha}^{boot}_{MKT}$	$\widehat{\alpha}^{boot}_{SMB}$	$\widehat{\alpha}^{boot}_{HML}$	$\widehat{\alpha}^{boot}_{RMW}$	$\widehat{\alpha}^{boot}_{CMA}$	$\widehat{\alpha}^{boot}_{MOM}$
Mean	1.96	1.84	2.02	3.30	3.00	2.21
Median	1.92	1.81	1.98	3.28	2.97	2.02
Maximum	3.42	2.63	4.57	4.49	5.87	7.22
Minimum	1.54	1.48	1.53	2.53	1.86	1.57
Std. Dev.	0.23	0.16	0.27	0.28	0.52	0.52
Skewness	1.78	1.00	2.69	0.50	0.66	2.19
Kurtosis	9.36	4.71	18.75	3.52	4.28	12.42
Jarque-Bera	2210.40	289.57	1,153.40	53.54	141.32	4,501.60
Probability	0.00	0.00	0.00	0.00	0.00	0.00
Observations	1,000	1,000	1,000	1,000	1,000	1,000
Panel B. Descriptive statistics for the bootstrapped tail exponents for the later subsample.						
	$\widehat{\alpha}^{boot}_{MKT}$	$\widehat{\alpha}^{boot}_{SMB}$	$\widehat{\alpha}^{boot}_{HML}$	$\widehat{\alpha}^{boot}_{RMW}$	$\widehat{\alpha}^{boot}_{CMA}$	$\widehat{\alpha}^{boot}_{MOM}$
Mean	2.39	2.71	2.01	2.51	2.36	2.14
Median	2.38	2.67	2.00	2.42	2.24	2.13
Maximum	3.41	3.99	2.90	6.41	4.87	3.24
Minimum	1.83	2.15	1.61	1.65	1.62	1.60
Std. Dev.	0.20	0.25	0.17	0.53	0.45	0.22
Skewness	0.96	0.96	0.92	1.73	1.60	0.88
Kurtosis	6.45	4.51	5.58	9.26	6.34	5.17
Jarque-Bera	649.96	247.89	416.82	2,127.90	892.48	326.41
Probability	0.00	0.00	0.00	0.00	0.00	0.00
Observations	1,000	1,000	1,000	1,000	1,000	1,000
Panel C. Estimated z-tests.						
$(\widehat{\alpha}_1 - \widehat{\alpha}_2)$	-0.43*	-0.87***	0.01	0.79*	0.64	0.07
$\widehat{\sigma}_{(\widehat{\alpha}_1 - \widehat{\alpha}_2)}$	0.22	0.21	0.23	0.42	0.49	0.40
z-statistic	-1.95	-4.14	0.04	1.88	1.31	0.18

*** Statistically significant on a 1 % level; * statistically significant on a 10 % level.

This table reports the results from a sample split analysis of blocks bootstrapped power law exponents. The first sub-sample spans from July 1926 to October 1972 for the realized MKT, SMB, and HML realized variances, July 1963 to February1993 for the realized RMW and CMA realized variances, and November 1926 to September 1974 for the realized MOM factor variance. The results are reported in Panel A. The second subsample spans from November 1972 to September 2022 for the realized MKT, SMB, and HML realized variances, March 1993 to September 2022 for the realized RMW and CMA realized variances, and October 1974 to September 2022 for the realized MOM factor variance. The results are reported in Panel B. The blocks bootstrap procedure is detailed in section 3.3.

Table A5

Summary statistics for quarterly realized Fama and French factor variance data.

	$\widehat{\sigma}^2_{MKT}$	$\widehat{\sigma}^2_{SMB}$	$\widehat{\sigma}^2_{HML}$	$\widehat{\sigma}^2_{RMW}$	$\widehat{\sigma}^2_{CMA}$	$\widehat{\sigma}^2_{MOM}$
Mean	75.72	22.67	24.99	10.00	8.98	39.27
Median	36.29	11.32	10.11	4.39	4.96	16.48
Maximum	927.75	323.58	325.97	165.27	103.37	573.50

(*continued on next page*)

Table A5 (*continued*)

	$\widehat{\sigma}^2_{MKT}$	$\widehat{\sigma}^2_{SMB}$	$\widehat{\sigma}^2_{HML}$	$\widehat{\sigma}^2_{RMW}$	$\widehat{\sigma}^2_{CMA}$	$\widehat{\sigma}^2_{MOM}$
Minimum	3.62	1.58	1.73	1.10	1.01	1.17
Std. Dev	122.68	40.10	44.96	17.60	12.41	68.74
Skewness	4.24	4.39	3.91	5.47	4.37	4.05
Kurtosis	21.08	21.89	17.12	39.62	24.95	19.93
Observations	384	384	384	226	226	382

This table reports the descriptive statistics for the quarterly realized Fama and French factor variances. Data used in this table are collected from the Data Library of Kenneth R. French and covers the excess market return (MKT), size factor (SMB), value factor (HML), profitability factor (RMW), investment factor (CMA), and momentum factor (MOM). The first three factors are the original Fama and French three factors model (Fama and French, 1993), while RMW and CMA are from the original Fama and French five-factor model (Fama and French, 2015). The MOM factor represents the momentum factor as proposed by Carhart (1997). The sample period for the realized variances of MKT, SMB, and HML spans the period July 1926 to September 2022 (384 quarterly observations). The realized variances for RMW and CMA cover the period from July 1963 to September 2022 (226 quarterly observations), whereas the realized variance of MOM encompasses the period from November 1926 to September 2022 (382 quarterly observations).

Table A6
Estimated power-law exponents for quarterly realized variance data.

Panel A. Estimated power-law **exponents for quarterly realized variance data.**						
	$\widehat{\sigma}^2_{MKT}$	$\widehat{\sigma}^2_{SMB}$	$\widehat{\sigma}^2_{HML}$	$\widehat{\sigma}^2_{RMW}$	$\widehat{\sigma}^2_{CMA}$	$\widehat{\sigma}^2_{MOM}$
$\widehat{\alpha}$	2.12	2.24	2.02	2.05	2.37	2.24
$\widehat{\sigma}$	0.07	0.09	0.07	0.09	0.14	0.13
x_{MIN}	31.20	12.67	10.51	3.47	5.74	43.00
D	0.05	0.04	0.05	0.04	0.04	0.06
% observations for which $x \geq x_{MIN}$	59.64 %	45.31 %	49.22 %	63.72 %	45.57 %	22.51 %
Panel B. 95 % confidence intervals for point estimates derived from monthly data.						
95 % Confidence Interval	(1.94,2.88)	(1.81,2.51)	(1.65,2.27)	(1.35,2.65)	(1.76,3.02)	(1,45,2.75)

This table reports the results of estimating power-law exponents for quarterly realized variances of the following factors: excess market factor (MKT), size factor (SMB), value factor (HML), profitability factor (RMW), investment factor (CMA), and momentum factor (MOM). The $\widehat{\alpha}$ represents the estimated tail exponent and $\widehat{\sigma}$ is the estimated standard deviation. The lower threshold x_{MIN} is estimated using the optimized Kolmogorov-Smirnov distance (D) distance, as detilaed in Clauset et al. (2009). The estimated x_{MIN} is the value that corresponds to the optimal distance D. Additionally, for each realized factor variance, the fraction of observations governed by power-law process is reported. The sample period for MKT, SMB, and HML realized variances spans the period July 1926 to September 2022 (384 quarterly observations). The RMW and CMA realized variances cover the period from July 1963 to September 2022 (226 quarterly observations), whereas the MOM realized variance encompasses the period from November 1926 to September 2022 (382 quarterly observations). Panel A reports the estimated power-law exponents, estimated standard deviations, cutoffs and percentages of observations governed by a power law. Panel B reports the 95 % confidence intervals for point estimates derived from monthly data (Table 2) where using bootstrapped standard deviation (Table 4).

Table A7
Goodness-of-fit tests for quarterly realized factor variances.

Realized factor variance	Power law	Log-normal	Exponential
$\widehat{\sigma}^2_{MKT}$	0.95	0.00	0.00
$\widehat{\sigma}^2_{SMB}$	0.63	0.08	0.00
$\widehat{\sigma}^2_{HML}$	0.59	0.02	0.00
$\widehat{\sigma}^2_{RMW}$	0.88	0.00	0.01
$\widehat{\sigma}^2_{CMA}$	0.74	0.01	0.01
$\widehat{\sigma}^2_{MOM}$	0.24	0.07	0.00

This table reports in the first column the results from goodness-of-fit tests derived from optimal Kolmogorov Smirnoff distances to examine if our empirical data and the data-generating process coincide with the hypothesized power law distribution parametrized via x_{min} and $\widehat{\alpha}$. Also, this table reports the goodness-of-fit tests for log- normal and exponential distributions in the third and fourth columns respectively. The null hypotheses of these tests assert that the empirical data and the data-generating process coincide with the log- normal or exponential distributions, respectively. To implement the tests, we make use of 2,500 synthetic data sets for each distribution under test.

Table A8
Summary statistics for blocks bootstrapped size factor premiums.

	$\overline{RET}^{SMB}_{BOOT}$	*t*-statistic
Mean	0.19	1.62
Median	0.19	1.58
Maximum	0.70	7.16
Minimum	−0.25	−2.03
Std. Dev.	0.12	1.06
Skewness	0.03	0.07
Kurtosis	3.05	3.22
Jarque-Bera	1.14	13.87
Probability	0.57	0.00
Observations	5,000	5,000

We retrieve the size factor over the July 1963 to September 2022 sample from Kenneth French's website and employ blocks bootstraps as described in section 3.3. to simulate 5,000 size factor evolutions. For each synthetic size factor, we compute the sample average and the corresponding *t*-statistic and plot these two metrics in this Figure. The blocks bootstraps procedure uses an expected block length of 60 months. This table reports the descriptive statistics for the simulated size premium obtained via blocks bootstraps ($\overline{RET}^{SMB}_{BOOT}$) and the *t*-statistics.

References

Andersen, T. G., Bollerslev, T., Diebold, F. X., & Ebens, H. (2001). The distribution of realized stock return volatility. *Journal of Financial Economics, 61*, 43–76.
Andersen, T. G., Bollerslev, T., Diebold, F. X., & Labys, P. (2001a). Modeling and forecasting realized volatility. *Econometrica, 71*, 579–625.
Andersen, T. G., Bollerslev, T., Diebold, F. X., & Labys, P. (2001b). The distribution of realized exchange rate volatility. *Journal of the American Statistical Association, 96*, 42–55.
Andersen, T. G., Bollerslev, T., Diebold, F. X., & Labys, P. (2003). Modeling and forecasting realized volatility. *Econometrica, 71*, 579–625.
Andersen, T. G., Bollerslev, T., & Meddahi, N. (2004). Analytical evaluation of volatility forecasts. *International Economic Review, 45*, 1079–1110.
Asness, C. S., Moskowitz, T. J., & Pedersen, L. H. (2013). Value and momentum everywhere. *Journal of Finance, 68*, 929–985.
Barndorff-Nielsen, O. E., & Shephard, N. (2002). Econometric analysis of realized volatility and its use in estimating stochastic volatility models. *Journal of the Royal Statistical Society: Series B (Statistical Methodology), 64*(2), 253–280.
Barroso, P., & Detzel, A. (2021). Do limits to arbitrage explain the benefits of volatility-managed portfolios? *Journal of Financial Economics, 140*(3), 744–767.
Bubák, V., Kočenda, E., & Žikeš, F. (2011). Volatility transmission in emerging European foreign exchange markets. *Journal of Banking and Finance, 35*, 2829–2841.
Cakici, N., Fieberg, C., Metko, D., & Zaremba, A. (2023). Do anomalies really predict market returns? New data and new evidence. *Review of Finance*. forthcoming.
Calvet, L., & Fisher, A. (2004). Regime-switching and the estimation of multifractal processes. *Journal of Financial Econometrics, 2*, 44–83.
Carhart, M. M. (1997). On Persistence in Mutual Fund Performance. *Journal of Finance, 52*(1), 57–82.
Cederburg, S., O'Doherty, M. S., Wang, F., & Yan, X. (2020). On the performance of volatility managed portfolios. *Journal of Financial Economics, 138*(1), 95–117.
Chan, L. K. C., Karceski, J., & Lakonishok, J. (2000). New paradigm or same old hype in equity investing? *Financial Analyst Journal, 56*, 23–36.
Chen, A., & Velikov, M. (2023). Zeroing in on the Expected Returns of Anomalies. *Journal of Financial and Quantitative Analysis, 58*(3), 968–1004.
Clauset, A., Shalizi, C. R., & Newman, M. E. J. (2009). Power law distributions in empirical data. *SIAM Review, 51*, 661–703.
Cont, R., & Tankov, P. (2004). *Financial Modelling With Jump Processes*. Boca Raton, FL: Chapman & Hall.
Dessaint, O., Olivier, J., Otto, C. A., & Thesmar, D. (2021). CAPM-Based Company (Mis)valuations. *Review of Financial Studies, 34*, 1–66.
Dichev, I. D. (1998). Is the risk of bankruptcy a systematic risk? *Journal of Finance, 53*, 1131–1147.
Dick-Nielsen, J., Feldhütter, P., Pedersen, L. H., C. Stolborg. *Corporate Bond Factors: Replication Failures and a New Framework* (September 28, 2023).
Dong, X., Li, Y., Rapach, D. E., & Zhou, G. (2022). Anomalies and the expected market return. *Journal of Finance, 77*, 639–681.
Efron, B., & Tibshirani, R. J. (1994). *An Introduction to the Bootstrap* (1st ed.). Chapman and Hall/CRC.
Fama, E. F. (1963). Mandelbrot and the stable Paretian hypothesis. *Journal of Business, 36*, 420–429.
Fama, E. F., & French, K. R. (1992). The cross-section of expected stock returns. *Journal of Finance, 47*, 427–465.
Fama, E. F., & French, K. R. (1993). Common factors in the returns on stocks and bonds. *Journal of Financial Economics, 33*, 3–56.
Fama, E. F., & French, K. R. (2015). A five-factor asset pricing model. *Journal of Financial Economics, 116*, 1–22.
Feng, G., Giglio, S., & Xiu, D. (2020). Taming the factor zoo: a test of new factors. *Journal of Finance, 75*, 1327–1370.
Godfrey, L. (2009). *Bootstrap Tests for Regression Models*. Palgrave Mcmillan, New York: Palgrave Texts in Econometrics.
Grobys, K. (2021). What do we know about the second moment of financial markets? *International Review of Financial Analysis, 78*, Article 101891.
Grobys, K. (2023a). Correlation versus co-fractality: Evidence from foreign exchange rate variances. *International Review of Financial Analysis, 86*, Article 102531.
Grobys, K. (2023b). A Multifractal Model of Asset (In)Variances. *Journal of International Financial Markets, Institutions and Money, 85*, Article 101767.
Grobys, K. (2023c). A fractal and comparative view of the memory of bitcoin and S&P 500 returns. *Research in International Business and Finance, 66*, Article 102021.
Grobys, K., & Juntilla, J.-P. (2021). Speculation and Lottery-Like Demand in Cryptocurrency Markets. *Journal of International Financial Markets, Institutions and Money, 71*, Article 101289.
Grobys, K., & Kolari, J. W. (2022). Choosing factors: the international evidence. *Applied Economics, 54*, 633–647.
Grobys, K., Ruotsalainen, J., & Äijö, J. (2018). Risk-managed industry momentum and momentum crashes. *Quantitative Finance, 18*, 1715–1733.
Grobys, K., & Vähämaa, S. (2020). Another look at value and momentum: volatility spillovers. *Review of Quantitative Finance and Accounting, 55*, 1459–1479.
Huber, P. J. (1981). *Robust statistics* (p. 308). New York: John Wiley & Sons.
Hirshleifer, D. (2001). Investor psychology and asset pricing. *Journal of Finance, 56*, 1533–1597.
Hou, K., Xue, C., & Zhang, L. (2020). Replicating Anomalies. *Review of Financial Studies, 33*(5), 2019–2133.
Huang, D., Li, J., Wang, L., & Zhou, G. (2020). Time series momentum: Is it there? *Journal of Financial Economics, 135*(3), 774–794.
Lintner, J. (1965a). The valuation of risk assets and the selection of risky investments in stock portfolios and capital budgets. *Review of Economics and Statistics, 47*, 13–37.

Lintner, J. (1965b). Security prices, risk and maximal gains from diversification. *Journal of Finance, 20*, 587–615.
Lux, T., & Alfarano, S. (2016). Financial power laws: Empirical evidence, models, and mechanisms. *Chaos, Solitons and Fractals, 88*, 3–18.
Lux, T., Morales-Arias, L., & Sattarhoff, C. (2014). A Markov-switching multifractal approach to forecasting realized volatility. *Journal of Forecasting, 33*, 532–541.
Mandelbrot, B. (1963a). New methods in statistical economics. *Journal of Political Economy, 71*, 421–440.
Mandelbrot, B. (1963b). The variation of certain speculative prices. *Journal of Business, 36*, 394–419.
Mandelbrot, B. (1982). *The Fractal Geometry of Nature*. W.H Freeman and Company.
Mandelbrot, B. (2008). The (Mis)behavior of Markets. A Fractal View of Risk, Ruin and Reward (*Profile Books, London*).
Merton, R. C. (1980). On estimating the expected return on the market: An exploratory investigation. *Journal of Financial Economics, 8*, 323–361.
Moreira, A., & Muir, T. (2017). Volatility-managed portfolios. *Journal of Finance, 72*, 1611–1643.
Mossin, J. (1966). Equilibrium in a capital asset market. *Econometrica, 35*, 768–783.
Pearson, K. F. R. S. (1900). On the criterion that a given system of deviations from the probable in the case of a correlated system of variables is such that it can be reasonably supposed to have arisen from random sampling. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science, 50*(302), 157–175.
Schwert, G. W. (2003). Anomalies and market efficiency. Constantinides G.M., Harris M., Stulz R.M. (Eds.), *Handbook of the Economics of Finance* (2003). *Amsterdam, North Holland*.
Sharpe, W. F. (1964). Capital Asset Prices: A Theory of Market Equilibrium Under Conditions of Risk. *Journal of Finance, 19*, 425–442.
Soares, J. O., Coutinho, M. C., & Martins, C. V. (2007). Forecasting errors in capital budgeting: A multi-firm post-audit study. *The Engineering Economist, 52*(1), 21–39.
Taleb, N. N. (2010). *The Black Swan*. New York, NY: Random House.
Smith, S.C., A. Timmermann (2022). Have risk premia vanished? *Journal of Financial Economics, 145*(2), Part B, 553–576.
Taleb, N. N. (2020). Statistical Consequences of Fat Tails: Real World Preasymptotics, Epistemology, and Applications. Papers and Commentary (STEM Academic Press).
Treynor, J. L. (1962). Toward a Theory of Market Value of Risky Assets. Unpublished manusscript. Final version in Asset Pricing and Portfolio Performance, 1999, Robert A. Korajczyk, ed., London: *Risk Books*, 15–22.
Wasserman, L. (2006). *All of nonparametric statistics*. New York: Springer.
White, E., Enquist, B., & Green, J. L. (2008). On estimating the exponent of power law frequency distributions. *Ecology, 89*, 905–912.

Journal of
Risk and Financial Management

MDPI

Article

On the Realized Risk of Foreign Exchange Rates: A Fractal Perspective

Masoumeh Fathi [1], Klaus Grobys [1,2] and James W. Kolari [3,*]

[1] Finance Research Group, School of Accounting and Finance, University of Vaasa, Wolffintie 34, 65200 Vaasa, Finland; masoumeh.fathi@uwasa.fi (M.F.); klaus.grobys@uwasa.fi (K.G.)
[2] Innovation and Entrepreneurship (InnoLab), University of Vaasa, Wolffintie 34, 65200 Vaasa, Finland
[3] Department of Finance, Mays Business School, Texas A&M University, College Station, TX 77843-4218, USA
* Correspondence: jkolari@mays.tamu.edu

Abstract: While well-established literature argues that realized variances are close to a lognormal distribution, this study follows Benoit Mandelbrot by taking a fractal perspective. Using power laws to model realized foreign exchange rate variances, our findings indicate that power laws offer an alternative to the lognormal in terms of goodness-of-fit tests. Further, our analysis shows that estimated power law exponents for seven out of nine realized FX variances are $\hat{\alpha} < 3$, which indicates that the variance of realized variance is statistically undefined. We conclude that the foreign exchange rate market is far riskier than earlier believed. By implication, documented research in an enormous body of literature that draws conclusions from variance analyses stands on shaky grounds.

Keywords: foreign exchange rates; Pareto distributions; power laws; second moment; variance; variance of variance

JEL Classification: C22; G12; G13; G14

Citation: Fathi, Masoumeh, Klaus Grobys, and James W. Kolari. 2024. On the Realized Risk of Foreign Exchange Rates: A Fractal Perspective. *Journal of Risk and Financial Management* 17: 79. https://doi.org/10.3390/jrfm17020079

Academic Editor: Thanasis Stengos

Received: 18 January 2024
Revised: 9 February 2024
Accepted: 12 February 2024
Published: 18 February 2024

> "There are strong pragmatic reasons to begin the study of economic distributions and time series by those that satisfy the law of Pareto."
>
> (Benoit Mandelbrot, mathematician known for the *Theory of Roughness and Fractals*)

1. Introduction

Uncertainty in the foreign exchange rate (FX) market has been the subject of intense study in the academic literature. This attention is not surprising in view of the fact that FX market capitalization is considerably larger than that in any other financial market. In 2022, the daily turnover of the global foreign exchange market with 39 different currencies reached 7.5 trillion per day.[1] Previous studies investigating FX risk have employed GARCH-type models (e.g., Jorion 1995; Baillie and Bollerslev 1991; Bollerslev and Melvin 1994).[2] Despite the widespread usage of GARCH-type models to analyze uncertainty in the FX market, work by Andersen et al. (2003) and Andersen et al. (2004) showed that simple reduced-form time series models for realized volatility (RV) outperform commonly-used GARCH-type models for forecasting future volatility. Unfortunately, many researchers incorporating RV to model FX risk have found there is a problem with realized variances (e.g., Wang and Yang 2009; Corsi 2004; Andersen et al. 2005; Bubák et al. 2011; Corsi et al. 2008). Specifically, Taleb (2020, p. 33) observed that, because the empirical distribution misrepresents the expected realizations in the tails, past data are not useful for predicting future maximums.

Power law functions offer a remedy to the aforementioned problem because the tail exponent of a power law function captures via extrapolation low-probability deviations not seen in the data. Taleb emphasized that such deviations play a disproportionately large share in determining the mean of the distribution. In advocating the usage of power law

functions in empirical finance research, he argued that "There are a lot of theories on why things should be power laws, as sort of exceptions to the way things work probabilistically. But it seems that the opposite idea is never presented: power laws should be the norm, and the Gaussian a special case . . ." Taleb (2020, p. 91).

This study uses power laws to explore the risk of foreign exchange rates in terms of weekly realized variances. Based on a sample of daily prices from 16 May 2006, to 29 December 2023, we compute weekly variances by summing the squared returns over five business days for each currency pair, leaving us with 920 weekly observations. Subsequently, as proposed by Clauset et al. (2009), we employ the maximum likelihood estimation approach to estimate power law exponents. Furthermore, we use the goodness-of-fit (GoF) test derived by these authors for testing various distributions, including lognormal, exponential, and power law. If the p-value for the power law model is higher than for the other competing distributions, the power law model is favored.

Our study contributes to the literature in several ways. First, it extends previous studies on power laws in financial economics. Calvet and Fisher (2004) and Lux et al. (2014) found that power law models usually outperform GARCH models in terms of forecasting future volatility. Some recent studies by Grobys (2021, 2023) used power law functions to model realized variances for foreign exchange rates, the S&P 500, and some commodities. Also, Grobys et al. (2021) employed power law functions to model the realized volatilities of some cryptocurrencies. The commonality among these studies is that the power law null model cannot be rejected for virtually all realized asset market variances. The study by Grobys (2023) found that estimated power law exponents vary between $\hat{\alpha} = 2.2515$ and $\hat{\alpha} = 2.780$, which indicates that the variance of daily realized FX variance does not exist. Further, using an application of Bayes' rule, he documented that, for the vast majority of realized FX rate variances, the lognormal performs worse in predicting likelihoods for the arrivals of extreme events; hence, the power law model should be favored. The present research adds to this literature by (a) using lower-frequented data; and (b) employing GoF tests to compare various potential distributions. The application of Bayes' rule in Grobys (2023) only provides evidence on the distribution that is associated with the *most extreme events*. Our approach is complementary in that we explicitly perform statistical tests using probability density functions. Furthermore, Mandelbrot (2008) argued that power law behavior is manifested in invariance of the power law exponent as time frequency changes. Hence, if realized FX variance exhibits fractal-like properties, we would expect that our estimates derived from weekly data should be close to the estimates documented for daily data. Moreover, due to the high rate of false discoveries in financial economics, Hou et al. (2020), Serra-Garcia and Gneezy (2021), and others have recommended scientific replications of reported research. Since we model realized foreign exchange rate variances using a different but related methodology to Grobys (2023), consistent with Hamermesh (2007), we provide replication findings to better understand the fractal-like nature of FX variances.

Relevant to our study, well-established literature has found that the distribution of realized volatility is typically very close to a lognormal distribution (e.g., Andersen et al. 2001a, 2001b; 2003). Renò and Rizza (2003) examined the unconditional volatility distribution of the Italian futures market. They concluded that the standard assumption of lognormal unconditional volatility is rejected and that a much better description is the Pareto distribution. Hence, there is no consensus in the literature on this issue. Is realized variance a lognormal distribution or closer to a power law? To our knowledge, no previous studies explicitly test power laws against other competing distributions to clarify this issue.

Using weekly realized FX variances over the sample from 16 May 2006, to 29 December 2023, our findings indicate that estimated power law exponents vary between $\hat{\alpha} = 2.4834$ and $\hat{\alpha} = 3.5968$. Strikingly, six out of nine realized FX variances have $\hat{\alpha} < 3$, which indicates that the variance of realized variance is statistically undefined. Since realized FX variances show fractal-like behavior, our results are in line with Grobys (2023) and corroborate Mandelbrot's (2008) arguments. Furthermore, using a significance level of

5%, GoF tests suggest that we cannot reject the power law model for seven out of nine realized FX variances as indicated by p-values > 5%. Interestingly, GoF tests suggest that the lognormal model cannot be rejected for eight out of nine realized FX variances. However, a comparison of the p-values shows that, for five realized FX variances, the p-values for the power law models are substantially higher than the corresponding p-values derived from the lognormal model. Since our findings are consistent with Grobys (2023), we conclude that power laws should be preferred for modeling realized FX risk.

The paper is organized as follows. The next section describes the data. Section 3 discusses the methodology. The last section provides discussion and concluding remarks.

2. Data

The data consist of daily observations downloaded from the global financial website investing.com (accessed on 1 January 2023) for Group of Ten (G10) currency pairs, including AUD/USD, EUR/USD, GBP/USD, NOK/USD, NZD/USD, USD/CAD, USD/CHF, USD/JPY, and USD/SEK. To ensure comparability with earlier documented results, we utilize a sample that is similar to earlier research. Retrieving data from yahoo.com, Grobys (2023) documented that data for the AUD/USD exchange rate are only publicly available from 16 May 2006. Therefore, our study employs data from 16 May 2006, to 29 December 2023. The analysis is restricted to daily data, wherein all FX rates are quoted on the same day, which results in 4599 daily observations. The daily return is computed as follows:

$$RET_{i,t} = 100 * \frac{(P_{i,t+1} - P_{i,t})}{P_{i,t}}, \quad (1)$$

where $RET_{i,t}$ denotes the return of FX i at time t and $P_{i,t}$ denotes the corresponding daily price. The weekly realized variance (WRV) of G10 currency pairs is computed as the sum of the squared returns over five business days for each currency pairs:

$$WRV_{i,j} = \sigma^2_{i,j} = \sum_{t \in j} (RET_{i,t})^2,$$

where j indicates the week and $t \in j$ indicates the corresponding trading days in the respective week. Table 1 reports the descriptive statistics of weekly realized variances for all nine currency pairs. The maximum and minimum values reveal a wide range of realized variances across most of the currency pairs, with USD/CHF showing the largest and EUR/USD the smallest fluctuations. The standard deviations represent the volatility of the realized FX rates variances, with USD/CHF and EUR/USD exhibiting the highest and lowest values, respectively. All currency pairs have high kurtosis due to the presence of fat tails in the distribution of weekly realized variances. Furthermore, among all FX weekly realized variances, skewness is most pronounced for realized USD/CHF variance.

Table 1. Summary statistics for the weekly realized variances of G10 currencies.

	AUD/USD	EUR/USD	GBP/USD	NOK/USD	NZD/USD	USD/CAD	USD/CHF	USD/JPY	USD/SEK
Mean	3.3170	1.6138	1.8063	3.2683	3.2357	1.5839	2.1195	1.9851	2.9710
Median	1.6427	1.0225	1.1017	1.9579	1.8758	0.9303	1.0874	1.0876	1.7309
Std.Dev.	8.7209	1.9674	3.3995	4.8487	4.8140	2.3482	9.2697	3.1163	4.0658
Kurtosis	187.8401	27.0404	289.4048	90.3521	91.2411	64.0870	618.1327	57.6324	25.8182
Skewness	12.3553	4.1499	13.9099	7.5595	7.2678	6.5452	23.5305	5.7512	4.3111
Minimum	0.0676	0.0333	0.0427	0.0398	0.0641	0.0230	0.0081	0.0155	0.0870
Maximum	161.8294	21.9666	78.6839	77.0604	82.8622	31.4039	255.0862	47.5975	39.5884
Obs	920	920	920	920	920	920	920	920	920

This table provides the descriptive statistics of the weekly realized variances for G10 currency pairs, including EUR/USD, AUD/USD, GBP/USD, NOK/USD, NZD/USD, USD/CAD, USD/CHF, USD/JPY, and USD/SEK. The weekly realized variance is com-

puted using daily data from investing.com, spanning from 16 May 2006, to 29 December 2023, and includes 920 observations.

Next, we conduct a visual analysis of the time series plots for each currency pair's realized variance to better understand how they evolved over time. Figure 1 provides a time series plot of the realized FX variances of various currency pairs. The key feature that stands out in these plots is that there are more extreme values or sudden shifts (spikes) than would be expected under normal or lognormal distributions. Some realized FX variances, such as the AUD/USD realized variance and the GBP/USD realized variance, have fewer spikes overall compared to other data series such as the EUR/USD realized variance and USD/JPY realized variance. These results suggest that the former pairs experienced fewer periods of extreme volatility during the sample period. The distribution of spikes over time is not uniform. Some realized FX variances, like the USD/CAD realized variance, show a concentration of extreme events around certain years. For example, there is a noticeable spike around 2008–2009 for the AUD/USD realized variance, which coincides with the global financial crisis. During this period, the Australian dollar had significant volatility arising from its close ties with global commodity prices and risk sentiment. Multiple spikes can be observed for the EUR/USD realized variance from 2010 to 2012, which is likely associated with the European sovereign debt crisis. The sharp spike for the GBP/USD realized variance around 2016 appears to be related to the Brexit referendum in which the United Kingdom voted to leave the European Union. This event caused significant uncertainty and led to a sharp depreciation of the British pound against the US dollar. Additionally, there is some asymmetry in the direction of the spikes. For example, the NOK/USD realized variance and the NZD/USD realized variance show shifts both upward and downward, whereas the EUR/USD realized variance shows predominantly upward shifts. These patterns indicate that certain currencies are more prone to strengthening or weakening against the US dollar under specific conditions. In sum, the presence of fat tails in the graphs indicates that events thought to be very rare may actually occur more frequently than expected. This finding has significant implications for financial risk management, as models based on normal or lognormal distributions may well underestimate the probability of extreme movements and thereby underestimate risk.

Figure 2 shows the difference between the theoretical lognormal distribution (red curve) and the distribution followed by each currency pair's weekly realized variance (black curve). In some cases, especially for currency pairs such as the USD/CHF realized variance and the AUD/USD realized variance, the empirical distribution has a higher peak (kurtosis) and heavier tails as shown by the empirical line extending further along the x-axis than the lognormal line. The presence of heavy tails in the empirical distribution indicates that extreme events—for instance, arising from financial crises, political events, or sudden economic shifts, etc.—leading to large variances in exchange rates are not as rare as the normal or lognormal distribution would suggest. As already mentioned, financial risk management and derivative pricing often wrongly assume lognormality in their models. Additionally, the realized variances of different currency pairs show different levels of variance and different fits to the lognormal distribution. For instance, the EUR/USD realized variance and GBP/USD realized variance have a lower spread, whereas the USD/CHF realized variance shows a much wider spread. These patterns reflect different levels of market volatility and risk associated with each underlying currency pair. Overall, the plots suggest that lognormal distribution is not an adequate model for the distribution of weekly realized variances of these currency pairs.

Figure 3 shows Lognormal quantile–quantile (QQ) plots. They offer a graphical tool to evaluate if a dataset follows lognormal distribution. It is evident that, even though all of the weekly realized variances of the currency pairs show a central tendency that follows a lognormal distribution, they exhibit heavy-tailed behavior as illustrated by the deviations from the theoretical line in the tails. This pattern again is indicative of a fat-tailed distribution, implying a higher frequency of extreme events than what would be expected by a lognormal distribution.

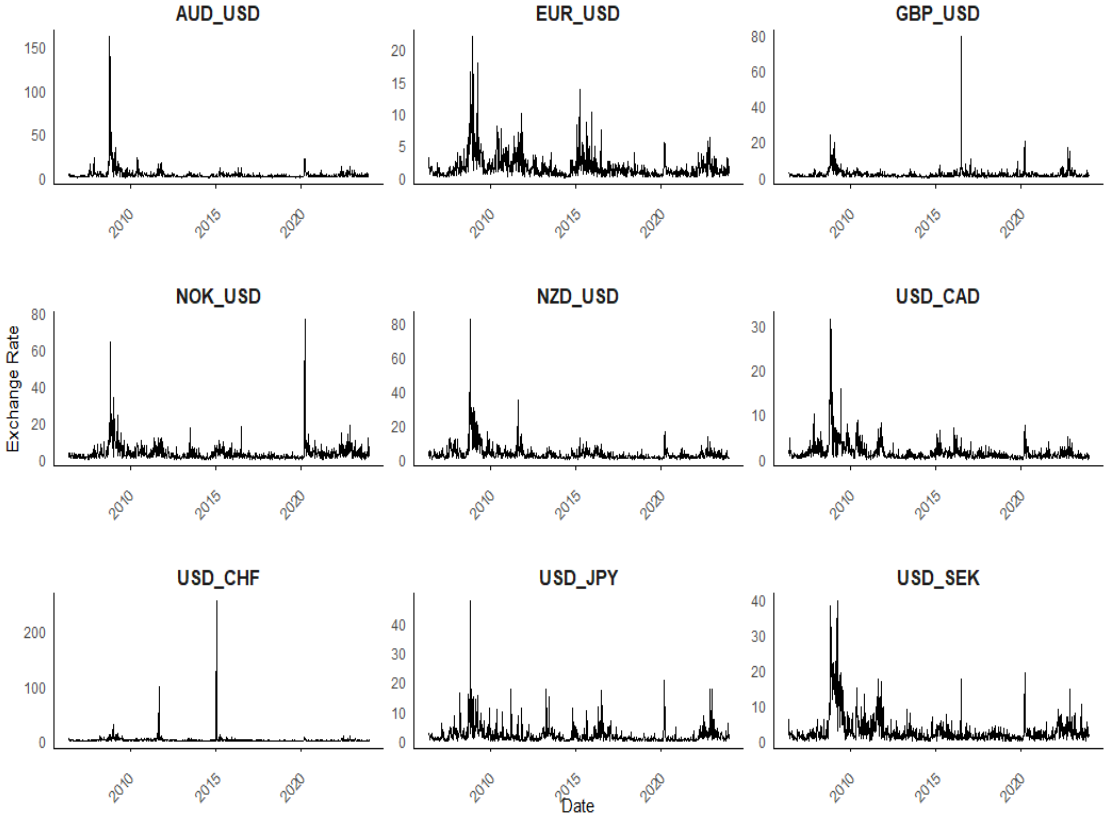

Figure 1. Time series plot.

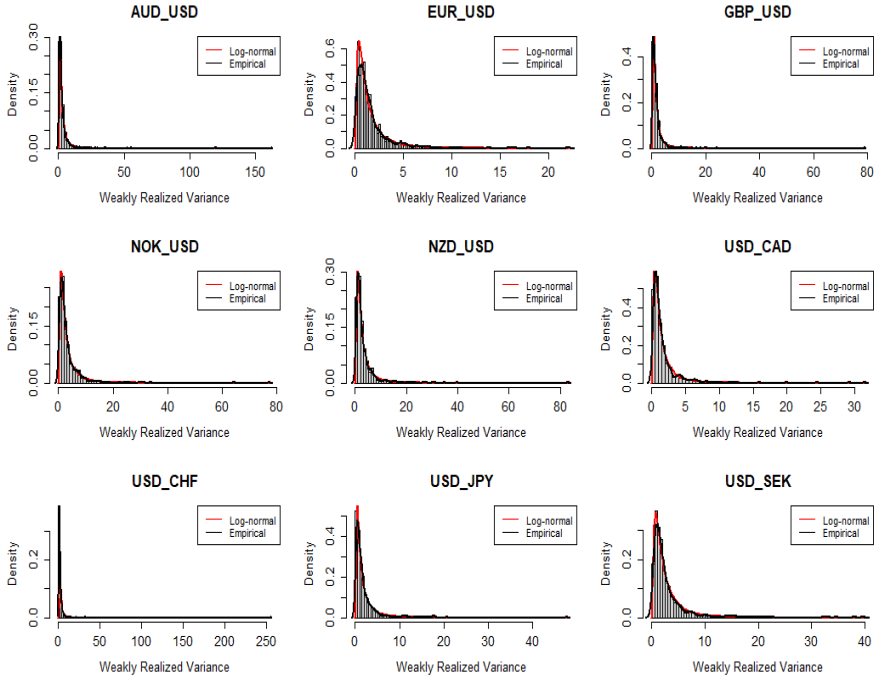

Figure 2. Histogram plot.

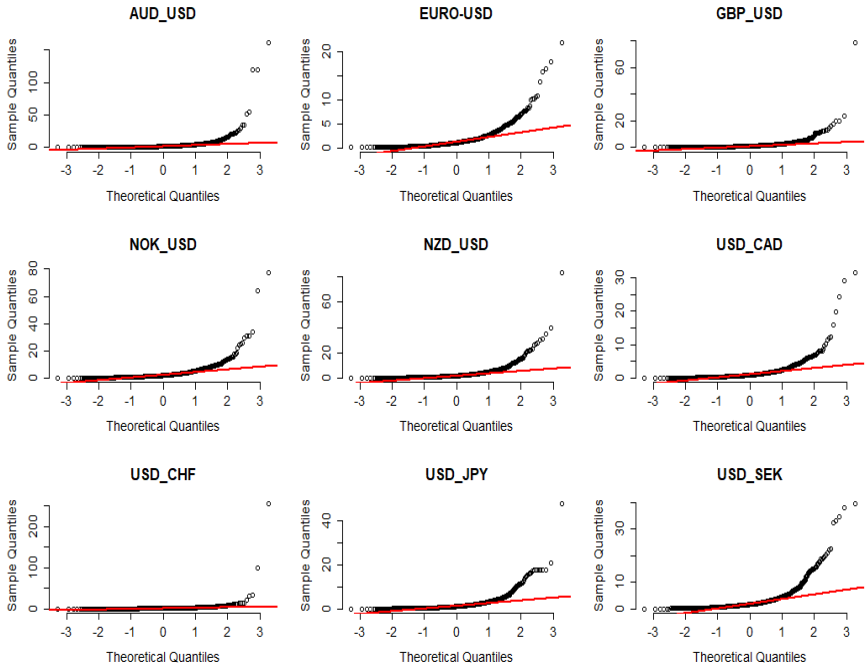

Figure 3. Lognormal QQ plot.

Below are time series graphs of monthly realized variance for weekly realized variances of G10 currency pairs, including EUR/USD, AUD/USD, GBP/USD, NOK/USD, NZD/USD, USD/CAD, USD/CHF, USD/JPY, and USD/SEK, using data from 16 May 2006, to 29 December 2023.

This figure illustrates theoretical lognormal distribution (red curve) versus the empirical distribution (black curve) for the weekly realized variances of G10 currencies, accounting for the following currencies: AUD/USD, EUR/USD, GBP/USD, NOK/USD, NZD/USD, USD/CAD, USD/CHF, USD/JPY, and USD/SEK. The sample is from 16 May 2006, to 29 December 2023.

The figure illustrates the theoretical log-normal quantile (red line) versus the quantile of weekly realized variance for G10 currency pairs including AUD/USD, EUR/USD, GBP/USD, NOK/USD, NZD/USD, USD/CAD, USD/CHF, USD/JPY, and USD/SEK. The sample is from 16 May 2006 to 29 December 2023.

In Figure 4, we employ the mean excess function (ME) plot to evaluate potential power law behaviors for our data. For a given threshold u, the mean excess of a random variable X is defined as:

$$E(X - u|X > u). \tag{2}$$

The presence of a linear upward trend in the ME plot points is consonant with power law behavior, where the slope of the line is directly related to the tail index parameter α. A greater value for α typically denotes a steeper slope on the ME plot. As depicted in Figure 4, the linear trend in the initial part of the plots is suggestive of a distribution with heavy tails, such as a power law distribution.

This figure displays the sorted data on the x-axis against the corresponding values of the mean excess function of weekly realized variances for G10 currency pairs, including AUD/USD, EUR/USD, GBP/USD, NOK/USD, NZD/USD, USD/CAD, USD/CHF, USD/JPY, and USD/SEK, using the whole dataset from 16 May 2006, to 29 December 2023. For a given threshold u of a random variable X, the mean excess is defined as $E(X - u|X > u)$.

Finally, we plot in Figure 5 the cumulative distribution function (CDF) of weekly realized variances for G10 currency pairs. Here, we compare the empirical CDF of each cur-

rency pair's realized variance with the power law (black), lognormal (red), and exponential (green) distributions. According to Figure 5, the power law and lognormal distributions align closely for AUD/USD and EUR/USD. However, power law provides a superior fit for GBP/USD, NOK/USD, USD/CAD, and USD/CHF, whereas exponential distribution is most suitable for USD/JPY and USD/SEK.

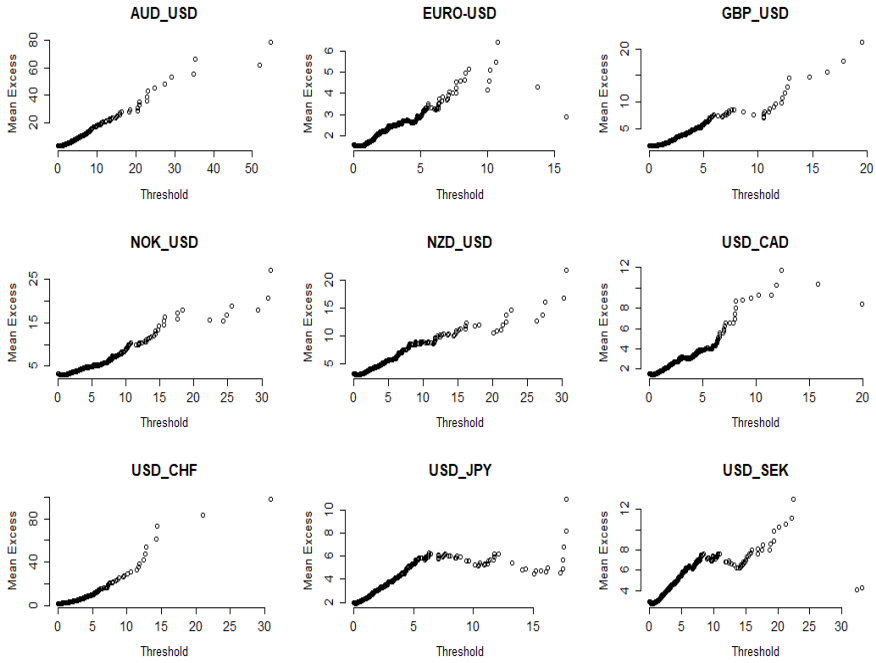

Figure 4. Mean excess plot.

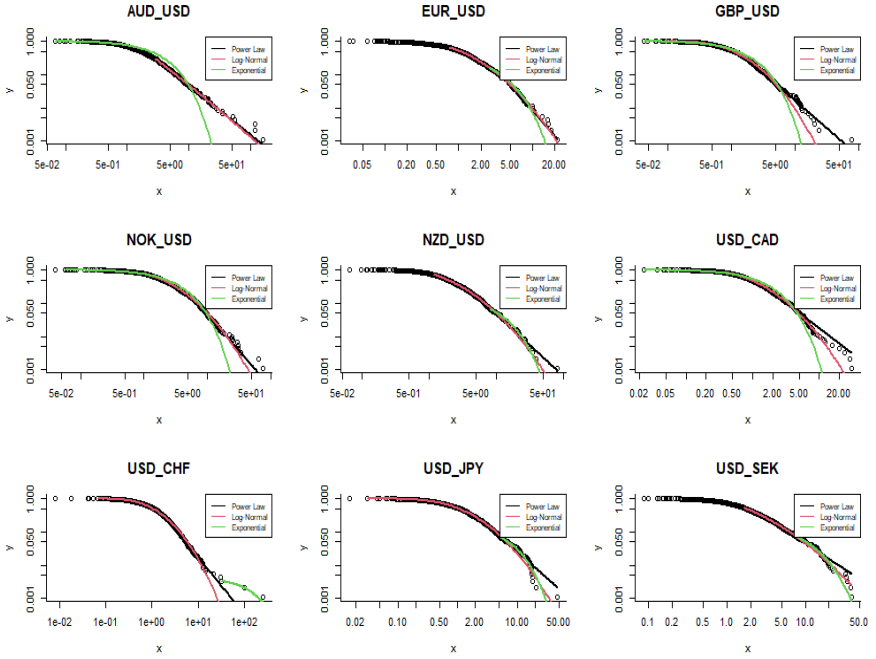

Figure 5. Distribution comparison.

This figure illustrates the cumulative density functions for the weekly realized variances of G10 currency pairs, including AUD/USD, EUR/USD, GBP/USD, NOK/USD, NZD/USD, USD/CAD, USD/CHF, USD/JPY, and USD/SEK, using data from 16 May 2006, to 29 December 2023. The fitted power law (black line), lognormal (red line), and exponential (green line) distributions are also given.

3. Methodology

3.1. Estimating Power Law Exponents Using Maximum Likelihood Estimation

In line with Grobys (2021, 2023), we utilize the power law function to model the variances of foreign exchange rates as:

$$p(x) = Cx^{-\alpha}, \tag{3}$$

where $C = (\alpha - 1)x_{MIN}^{\alpha-1}$ with $\alpha \in \{\mathbb{R}_+ | \alpha > 1\}$, $x = \sigma_i^2$ denotes the respective realized weekly FX variance for currency pair i, provided $x \in \{\mathbb{R}_+ | x_{MIN} \leq x < \infty\}$, x_{MIN} is the minimum value of realized FX variance that is obtained by a power law process, and α is the magnitude of the specific tail exponent.[3] The expected value of the realized weekly FX variance, given that $X > x_{MIN}$, is defined as $E[X|X > x_{MIN}]$ and computed by:

$$E[X|X > x_{MIN}] = \int_{x_{MIN}}^{\infty} xp(x)dx = \frac{(\alpha - 1)}{(\alpha - 2)}x_{MIN}. \tag{4}$$

The second moment, which is the variance of the realized variance, given that $X > x_{MIN}$, is defined as $E\left[X^2|X > x_{MIN}\right]$ and computed by:

$$E\left[X^2|X > x_{MIN}\right] = \int_{x_{MIN}}^{\infty} x^2p(x)dx = \frac{(\alpha - 1)}{(\alpha - 3)}x_{MIN}^2. \tag{5}$$

Finally, higher moments of order 'k' are computed as:

$$E\left[X^k|X > x_{MIN}\right] = \frac{(\alpha - 1)}{(\alpha - 1 - k)}x_{MIN}^k \tag{6}$$

Based on Equations (4) and (5), it is evident that the theoretical mean of realized variances is only defined for $\alpha > 2$, whereas the variance of the realized variance requires $\alpha > 3$.

Following White et al. (2008) and Clauset et al. (2009), who argued that maximum likelihood estimation (MLE) is the optimal method for calculating power law exponents, the tail exponents $\hat{\alpha}$ are computed as follows:

$$\hat{\alpha} = 1 + N\left(\sum_{i=1}^{N} \ln\left(\frac{x_i}{x_{MIN}}\right)\right)^{-1}, \tag{7}$$

where $\hat{\alpha}$ denotes the MLE estimator, $N \leq T$ is the number of observations greater than x_{MIN}, and the rest of notation is as earlier. According to Equation (7), the MLE estimator depends on the chosen x_{MIN}. Like Clauset et al. (2009), we select x_{MIN} by minimizing the distance between the power law model and the empirical data defined as Kolmogorov–Smirnov distance (D) or maximum distance between the cumulative density functions (CDFs) of the data and the fitted model:

$$D = MAX_{x \geq x_{MIN}}|S(x) - P(x)| \tag{8}$$

where $S(x)$ represents the cumulative distribution function (CDF) of the data for observations with values greater than or equal x_{MIN}, and $P(x)$ is the CDF for the power law model that provides the best fit to the data in the range of $x \geq x_{MIN}$. The estimate $\hat{x}_{MIN}$ is

determined as the value of x_{MIN} that minimizes D. Moreover, the standard deviation for estimated power law exponents is defined as:

$$\sigma = \frac{\hat{\alpha} - 1}{\sqrt{N}} + O(\frac{1}{N}). \quad (9)$$

Summary statistics for the estimated power law exponents are presented in Table 2. The estimated tail exponent for the realized AUD/USD variance is $\hat{\alpha} = 2.4834$, with corresponding standard deviation $\hat{\sigma} = 0.0911$. It follows that $\hat{\alpha}$ is 5.67 standard deviations below a hypothesized tail exponent corresponding to $\alpha = 3$. Therefore, α for the realized AUD/USD variance is statistically significant $\alpha < 3$ indicating that the variance of realized variance is statistically undefined. The same argument holds for the realized variances of USD/CAD, USD/GBP, USD/CHF, USD/JPY, and USD/SEK; however, the estimated tail exponent for the realized NZD/USD variance (e.g., $\hat{\alpha} = 2.8648$) is statistically not different from $\alpha = 3$. Furthermore, we observe from Table 2 for the lower bounds of the 95% confidence intervals, $\alpha < 3$ holds for all realized variance series, which provides further evidence that the infinite variance hypothesis cannot be rejected. Overall, we infer from these findings that the theoretical variances for the distributions of at least six out of nine realized FX variances are undefined, consistent with Grobys' (2023) findings derived from daily data (viz., estimated power law exponents for daily range-based FX variances varied between $\hat{\alpha} = 2.2515$ and $\hat{\alpha} = 2.7807$).

Table 2. Estimated power law functions.

Distribution	$\hat{\alpha}$	$\hat{\sigma}$	x_{MIN}	D	N_{PL}	95% CI
AUD/USD	2.4834	0.0911	2.9452	0.0242	28.80%	[2.3048, 2.6620]
EUR/USD	3.5968	0.3126	4.3242	0.0497	7.50%	[2.9841, 4.2095]
GBP/USD	2.6704	0.1092	1.9530	0.0319	25.43%	[2.4563, 2.8844]
NOK/USD	3.1845	0.1978	6.1729	0.0389	13.26%	[2.7968, 3.5721]
NZD/USD	2.8648	0.1617	5.4697	0.0509	14.46%	[2.5479, 3.1818]
USD/CAD	2.5048	0.0869	1.4408	0.0535	32.61%	[2.3345, 2.6751]
USD/CHF	2.7312	0.1256	2.4059	0.0275	20.65%	[2.4851, 2.9774]
USD/JPY	2.5888	0.1276	3.0011	0.0485	16.85%	[2.3387, 2.8390]
USD/SEK	2.5553	0.1019	3.2868	0.0462	25.33%	[2.3556, 2.7550]

This table reports the results of estimating power law exponents for weekly realized variances for G10 currencies, including AUD/USD, EUR/USD, GBP/USD, NOK/USD, NZD/USD, USD/CAD, USD/CHF, USD/JPY, and USD/SEK, using data from 16 May 2006, to 29 December 2023. The estimate $\hat{\alpha}$ represents the estimated tail exponent and $\hat{\sigma}$ is the estimated standard deviation. The lower threshold x_{MIN} is estimated using the optimized Kolmogorov–Smirnov distance (D) distance as proposed by Clauset et al. (2009). The estimated x_{MIN} is the value that corresponds to optimal distance D. N_{PL} denotes the percentage of sample observations governed by a power law process. The last column reports the 95% confidence interval (CI) intervals for the estimated power law exponents.

3.2. Testing the Power Law Model

According to Clauset et al. (2009), a tail exponent within the interval $2 < \alpha < 3$ does not necessarily indicate that the power law is a plausible fit to the data. Therefore, we need to verify whether the observed dataset genuinely adheres to a power law distribution. This test employs the Kolmogorov–Smirnov distance (D) as outlined in Equation (8). The objective is to determine whether the empirical data and synthetic data generated from the power law distribution, with specific x_{MIN} and α values, originate from the same distribution. The null hypothesis for the goodness-of-fit test hypothesizes that the empirical data and the synthetic data generated from the power law distribution, with specified x_{MIN} and α values, share the same distribution. If so, power law distribution is a plausible fit for the empirical data. Conversely, the alternative hypothesis posits that the

empirical data and the synthetic data do not originate from the same distribution, which indicates that power law distribution is not a suitable fit for the empirical data. In this case, alternative distribution models should be explored to better characterize the underlying data distribution.

To implement this analysis, we generate large synthetic datasets followed by the power law distribution utilizing the estimated parameters x_{MIN} and α derived from the previous section for each FX currency variances. Specifically, we compute the p-value of this test by fraction $\#(D_S > D)/K$, where D_S represents the estimated Kolmogorov–Smirnov distance for a synthetic dataset with specified x_{MIN} and α values mirroring those estimated for the original FX variances, D denotes the estimated Kolmogorov–Smirnov distance for the original dataset, and K is the number of synthetic datasets generated for testing purposes.

We set the significance level at 5%, which means that we do not reject the power law model if $\#\frac{D_S > D}{K} > 0.05$. This test aims to determine whether the empirical data and the generated dataset share the same distribution. We use $K = 500$ synthetic data series to compute the p-values for the goodness-of-fit tests. The second column of Table 3 reports the p-values corresponding to the null hypothesis of the power law model. It is evident that, for all realized FX variances, the p-values are above 5%, with the exception of USD/CAD and USD/SEK. These findings strongly support the notion that power laws are indeed plausible distributions for modeling realized FX variances.

Table 3. Goodness-of-fit test.

	Power Law	Lognormal	Exponential
EUR/USD	0.96	0.05	0.00
AUD/USD	0.76	0.40	0.08
GBP/USD	0.64	0.38	0.14
NOK/USD	0.86	0.60	0.00
NZD/USD	0.30	0.90	0.10
USD/CAD	0.00	0.46	0.00
USD/CHF	0.98	0.26	0.44
USD/JPY	0.16	0.64	0.32
USD/SEK	0.04	0.86	0.40

This table reports the p-values from goodness-of-fit tests derived from the Kolmogorov–Smirnoff distance to examine if our empirical data and the generated data from the power law distribution with particular x_{min} and α belong to the same distribution. The corresponding p-values are reported in the second column. The table gives the p-values for goodness-of-fit tests for exponential and lognormal distributions in the third and fourth columns, respectively. We use $K = 500$ synthetic data series to compute the p-values.

3.3. An Empirical Comparison of Distributions

Although our analysis in the previous section suggests that power law distributions provide a reasonable fit for realized FX variances, it should be noted that they may not be the optimal model, and other distributions could potentially offer a superior fit. Consequently, we focus on comparing two primary candidates—that is, the lognormal and exponential distributions against the power law to determine the appropriate fitting model for our data. It is important to note that prior studies have documented that realized volatility tends to follow a distribution that is nearly lognormal (e.g., Andersen et al. 2001a; 2001b, 2003). In our study, we investigate the suitability of the lognormal and exponential distributions for realized variance of G10 currency pairs. Employing a goodness-of-fit test similar to the one described by Clauset et al. in the previous section, we test the hypothesis that the lognormal (exponential) distribution can be considered a plausible fit. Thus, the lognormal (exponential) distribution is used as the null hypothesis against which the currency pair variances are tested at the 5 percent significance level as before. Columns 3 and 4 in Table 3 display the p-values of the goodness-of-fit test for lognormal and exponential distributions,

respectively. The results in column 3 indicate that eight out of nine p-values exceed the significance level of 5%, implying that lognormal distribution is a plausible fit to the data, with the exception of the EUR/USD realized variance. In addition, the results in column 4 indicate that exponential distribution is at least a plausible fit for the realized variances of the following currency pairs: AUD/USD, GBP/USD, NZD/USD, USD/CHF, USD/JPY, and USD/SEK.

According to Clauset et al., using p-values for power law models alongside other potential alternative distributions (e.g., lognormal and exponential) allows us to construct a persuasive argument regarding the suitability of the power law model for all currency pairs. Specifically, they observed that a high p-value for the power law model in combination with lower p-values for other distributions supports ruling out alternative distributions. We find in Table 3 that power law models for the realized variances of EUR/USD, AUD/USD, GBP/USD, NOK/USD, and USD/CHF exhibit considerably higher p-values than the alternative distributions. Interestingly, Grobys (2023) employed an application of Bayes' rule to examine the conditional probability that the underlying distribution is lognormally distributed, given that distribution-specific extreme events occurred. In doing so, he explored the cross-section of extreme events for daily realized range-based FX variances. His findings indicated that, when lognormal distribution is assumed and power law processes are equally likely, for six out of nine daily realized FX rate variances, the probability that the underlying distribution is lognormal given the arrival of the corresponding maximums is <25%. Even when it was assumed that lognormal distribution is 70% likely, for six out of nine daily realized FX rate variances, the probability that the underlying distribution is lognormal given the arrival of the corresponding maximums is <40%. Overall, Grobys (2023) concluded that, for the vast majority of realized FX rate variances, we can rule out a lognormal distribution as the corresponding underlying data-generating process for realized FX variances. The results of the present study in association with the results documented in Grobys (2023) provide evidence that power law models describe the distribution of realized FX variances more accurately. This inference is in line with Renò and Rizza (2003) but contradicts some well-established literature (Andersen et al. 2001a; 2001b; 2003).

3.4. Robustness Checks

Table A1 in the Appendix A reports the results from a sample-split analysis of estimating power law exponents. In this analysis, we divided the sample into two non-overlapping subsamples of equal length. Panel A of Table A1 reports the estimation results for the first subsample from 16 May 2006, to 3 March 2015, whereas Panel B documents the estimation results for the second subsample from 10 March 2015, to 29 December 2023. Each subsample includes 460 observations. From Table A1, it is evident that the point estimates $\hat{\alpha}_{1i}$ and $\hat{\alpha}_{2i}$ are fairly similar between the two subsamples. However, it is noteworthy that there are some differences between the two subsamples. For instance, subsample 1—covering the period up to 3 March 2015—includes the aftermath of the 2008 financial crisis and the European debt crisis, among other events. These severe economic events are likely to have influenced market volatility and, in turn, the behavior of currency variances. On the other hand, subsample 2 captures a different economic period, including the gradual recovery from previous crises, shifts in monetary policy across major economies, and the impact of the COVID-19 pandemic. Overall, from Panels A and B of Table A1 for the lower bounds of the 95% confidence intervals, it holds that $\alpha < 3$ for all realized variance series, which is similar to our previously documented evidence.

4. Discussion and Concluding Remarks

The present study applied a fractal perspective to investigate the second moment of foreign exchange rates. We examined G10 currencies that account for approximately 70% of overall FX market capitalization. Instead of relying on parametric GARCH-type models, we utilized realized variances derived from daily price data. Realized variances are employed in MLE estimations, which provide reliable estimation even in the presence

of extremely fat-tailed data. Well-established literature on realized volatility contends that the underlying distribution of realized asset uncertainty is close to normal (e.g., Andersen et al. 2001a; 2001b; 2003). Our findings contradict this view as power law models provide better data fits and more reliable probabilities for the arrivals of extreme events. We infer that the conventional lognormal distribution for modeling risk may severely underestimate tail risks.

Further findings indicated that the variance of variance does not exist for most of the foreign exchange rate variances under study. This evidence is in line with Grobys (2023), who fitted power laws to daily range-based FX variances. Mandelbrot (2008) argued that power law behavior is manifested in invariance of the power law exponent across time frequencies. In his seminal study on cotton price changes, Mandelbrot (1963) was the first to show that the power law exponent for cotton price changes does not change across various time frequency—a property suggesting fractal-like behavior. Consistent with his findings, weekly data on realized FX variances behave very similar to daily data.

An important implication of our empirical results is that standard statistical methodologies can lead to invalid inferences. In this regard, Fama (1963) observed that OLS regression assuming finite variances may not be appropriate in empirical analyses in financial economics due to infinite variances. Using a novel methodology that explicitly focuses on the statistical information residing in the tails, our findings suggested that the variances of the majority of exchange rate variance are infinite. As noted by Mandelbrot (2008), the reliance on traditional finance models gives a comforting impression of precision and competence:

"It is false confidence, of course. The problem lies at the roots of the standard model, in its assumption that the best way to think about stock markets is as a grand game of coin-tossing. If you are going to use probability to model financial markets, then you had better use the right kind of probability. Real markets are wild. Their price fluctuations can be hair-raising–far greater and more damaging than the mild variations of orthodox finance." (Mandelbrot 2008, p. 105).

This study supports Mandelbrot's (2008) viewpoint that price fluctuations in the currency market are wild—indeed so uncertain that the second moment of realized variances does not converge.

While we focused on G10 currencies in the FX market, future research is recommended to explore this issue in other financial markets, such as the market for cryptocurrencies which is exposed to greater risks than established financial markets. Also, it is possible that, as documented in Ibragimov et al. (2013), power law behavior could be more pronounced in emerging than advanced economies. Hence, future research is warranted to investigate this issue for emerging economies.

Finally, Gabaix and Ibragimov (2011) proposed various approaches for estimating the optimal tail exponent. Several studies have concluded that inferences based on the tail index using a maximum likelihood estimator suffer from shortcomings, including sensitivity to dependence and small sample sizes (e.g., Embrechts et al. 1997). Even though our sample consists of sufficient numbers of realized variance observations, dependency could exist in our data. We chose the MLE approach in line with Clauset et al. (2009) to (1) ensure that our results are comparable to previous studies and (2) implement goodness-of-fit tests to discriminate between competing distributions. However, following Gabaix and Ibragimov, future studies are encouraged to use the log–log rank size OLS regression approach rather than the MLE approach.

Author Contributions: Conceptualization: M.F., K.G. and J.W.K. Methodology: M.F. and K.G. Software: M.F. Validation: M.F., K.G. and J.W.K. Formal analysis: M.F. Investigation: M.F., K.G. and J.W.K. Resources: M.F., K.G. and J.W.K. Data curation: M.F. Writing—original draft preparation: M.F., K.G. and J.W.K. Writing—review and editing: J.W.K. Visualization: M.F. and J.W.K. Supervision: K.G. and J.W.K. Project administration: M.F., K.G. and J.W.K. Funding acquisition: N/A. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Data Availability Statement: Data used in this study was downloaded for free from the data base investing.com.

Conflicts of Interest: The authors declare no conflict of interest.

Appendix A

Table A1. Sample-split analysis.

Panel A						
Distribution	$\hat{\alpha}$	$\hat{\sigma}$	x_{MIN}	D	N_{PL}	**95% CI**
AUD/USD	2.4139	0.1666	6.2002	0.0545	15.65%	[2.0873, 2.7405]
EUR/USD	2.4325	0.0934	1.3364	0.0691	51.09%	[2.2493, 2.6156]
GBP/USD	2.5637	0.1196	1.4831	0.0416	37.17%	[2.3293, 2.7981]
NOK/USD	3.2158	0.2432	6.0219	0.0553	18.04%	[2.7391, 3.6926]
NZD/USD	2.4391	0.1085	3.4412	0.0582	38.26%	[2.2265, 2.6518]
USD/CAD	3.0643	0.2522	3.8276	0.0632	14.57%	[2.5700, 3.5586]
USD/CHF	2.7203	0.1640	2.9686	0.0462	23.91%	[2.3989, 3.0418]
USD/JPY	2.4901	0.1268	2.3796	0.0452	30.00%	[2.2415, 2.7387]
USD/SEK	2.3217	0.1058	3.2904	0.0736	33.91%	[2.1143, 2.5291]
Panel B						
Distribution	$\hat{\alpha}$	$\hat{\sigma}$	x_{MIN}	D	N_{PL}	**95% CI**
AUD/USD	3.3773	0.2389	3.0557	0.0453	21.52%	[2.9090, 3.8456]
EUR/USD	2.7133	0.1342	1.1669	0.0481	35.43%	[2.4503, 2.9764]
GBP/USD	2.7723	0.1682	2.0690	0.0476	24.13%	[2.4426, 3.1020]
NOK/USD	2.8486	0.1591	3.1635	0.0389	29.34%	[2.5367, 3.1604]
NZD/USD	3.1287	0.2139	2.9721	0.0732	21.52%	[2.7094, 3.5481]
USD/CAD	3.0508	0.1772	1.2762	0.0650	29.13%	[2.7036, 3.3981]
USD/CHF	3.0085	0.1804	1.3947	0.0474	26.95%	[2.6550, 3.3620]
USD/JPY	2.8305	0.2538	3.2205	0.0528	11.30%	[2.3330, 3.3281]
USD/SEK	3.4147	0.2491	3.0119	0.0689	20.43%	[2.9266, 3.9029]

This table presents the results from a sample-split analysis for estimating power law exponents. We divide the sample period into two equal-length subsamples. Panel A details the first subsample from 16 May 2006, to 3 March 2015, whereas Panel B covers the second subsample from 10 March 2015, to 29 December 2023. Each subsample includes 460 observations. $\hat{\alpha}$ represents the estimated tail exponent, and $\hat{\sigma}$ is the estimated standard deviation. The lower threshold x_{MIN} is estimated using the optimized Kolmogorov–Smirnov distance (D) distance as proposed by Clauset et al. (2009). The estimated x_{MIN} is the value that corresponds to optimal distance D. N_{PL} denotes the percentage of sample observations governed by a power law process. The last column reports the 95% confidence interval (CI) intervals for the estimated power law exponents.

Notes

1 See https://www.statista.com/statistics/247328/activity-per-trading-day-on-the-global-currency-market/ (accessed on 1 January 2023).

2 Other related studies are Alexander (1995) and Bauwens et al. (2005), who used (G)ARCH-type models to explore common volatility in the foreign exchange rate market as well as the impact of scheduled and unscheduled news announcements on foreign exchange rate return volatility.

3 Note that for the sake of readability, we skip in our notation here the index j denoting the week.

References

Alexander, Carol O. 1995. Common volatility in the foreign exchange market. *Applied Financial Economics* 5: 1–10. [CrossRef]

Andersen, Torben G., Tim Bollerslev, and Francis X. Diebold. 2005. *Roughing It Up: Including Jump Components in the Measurement, Modelling and Forecasting of Return Volatility*. NBER Working Paper No. 11775. Cambridge: National Bureau of Economic Research.

Andersen, Torben G., Tim Bollerslev, and Nour Meddahi. 2004. Analytical evaluation of volatility forecasts. *International Economic Review* 45: 1079–110. [CrossRef]

Andersen, Torben G., Tim Bollerslev, Francis X. Diebold, and Heiko Ebens. 2001a. The distribution of realized stock return volatility. *Journal of Financial Economics* 61: 43–76. [CrossRef]
Andersen, Torben G., Tim Bollerslev, Francis X. Diebold, and Paul Labys. 2001b. The distribution of realized exchange rate volatility. *Journal of the American Statistical Association* 96: 42–55. [CrossRef]
Andersen, Torben G., Tim Bollerslev, Francis X. Diebold, and Paul Labys. 2003. Modeling and forecasting realized volatility. *Econometrica* 71: 579–625. [CrossRef]
Baillie, Richard T., and Tim Bollerslev. 1991. Intra-day and inter-market volatility in foreign exchange rates. *Review of Economic Studies* 58: 565–85. [CrossRef]
Bauwens, Luc, Walid Ben Omrane, and Pierre Giot. 2005. News announcements, market activity and volatility in the euro/dollar foreign exchange market. *Journal of International Money and Finance* 24: 1108–25. [CrossRef]
Bollerslev, Tim, and Michael Melvin. 1994. Bid—Ask spreads and volatility in the foreign exchange market: An empirical analysis. *Journal of International Economics* 36: 355–72. [CrossRef]
Bubák, Vít, Evžen Kočenda, and Filip Žikeš. 2011. Volatility transmission in emerging European foreign exchange markets. *Journal of Banking and Finance* 35: 2829–41. [CrossRef]
Calvet, Laurent E., and Adlai J. Fisher. 2004. Regime-switching and the estimation of multifractal processes. *Journal of Financial Econometrics* 2: 44–83. [CrossRef]
Clauset, Aaron, Cosma Rohilla Shalizi, and Mark E. J. Newman. 2009. Power law distributions in empirical data. *SIAM Review* 51: 661–703. [CrossRef]
Corsi, Fulvio. 2004. *A Simple Long Memory Model of Realized Volatility*. Working Paper. Lugano: Institute of Finance, University of Lugano.
Corsi, Fulvio, Stefan Mittnik, Christian Pigorsch, and Uta Pigorsch. 2008. The volatility of realized volatility. *Econometric Reviews* 27: 46–78. [CrossRef]
Embrechts, Paul, Claudia Klüppelberg, and Thomas Mikosch. 1997. *Modelling Extremal Events for Insurance and Finance*. New York: Springer.
Fama, Eugene F. 1963. Mandelbrot and the stable Paretian hypothesis. *Journal of Business* 36: 420–29. [CrossRef]
Gabaix, Xavier, and Rustam Ibragimov. 2011. Rank—1/2: A simple way to improve the OLS estimation of tail exponents. *Journal of Business and Economic Statistics* 29: 24–39. [CrossRef]
Grobys, Klaus. 2021. What do we know about the second moment of financial markets? *International Review of Financial Analysis* 78: 101891. [CrossRef]
Grobys, Klaus. 2023. Correlation versus co-fractality: Evidence from foreign-exchange-rate variances. *International Review of Financial Analysis* 86: 102531. [CrossRef]
Grobys, Klaus, James W. Kolari, and Niranjan Sapkota. 2021. On the stability of stablecoins. *Journal of Empirical Finance* 64: 207–23. [CrossRef]
Hamermesh, Daniel S. 2007. Viewpoint: Replication in economics. *Canadian Journal of Economics* 40: 715–33. [CrossRef]
Hou, Kewei, Chen Xue, and Lu Zhang. 2020. Replicating anomalies. *Review of Financial Studies* 33: 2019–133. [CrossRef]
Ibragimov, Marat, Rustam Ibragimov, and Paul Kattuman. 2013. Emerging markets and heavy tails. *Journal of Banking and Finance* 37: 2546–59. [CrossRef]
Jorion, Philippe. 1995. Predicting volatility in the foreign exchange market. *Journal of Finance* 50: 507–28. [CrossRef]
Lux, Thomas, Leonardo Morales-Arias, and Cristina Sattarhoff. 2014. A Markov-switching multifractal approach to forecasting realized volatility. *Journal of Forecasting* 33: 532–41. [CrossRef]
Mandelbrot, Benoit B. 1963. The variation of certain speculative prices. *Journal of Business* 36: 394–419. [CrossRef]
Mandelbrot, Benoit B. 2008. *The (Mis)behavior of Markets. A Fractal View of Risk, Ruin and Reward*. London: Profile Books.
Renò, Roberto, and Rosario Rizza. 2003. Is volatility lognormal? Evidence from Italian futures. *Physica A: Statistical Mechanics and Its Applications* 322: 620–28. [CrossRef]
Serra-Garcia, Marta, and Uri Gneezy. 2021. Nonreplicable publications are cited more than replicable ones. *Science Advances* 7: 1705. [CrossRef]
Taleb, Nassim Nicholas. 2020. *Statistical Consequences of Fat Tails: Real World Preasymptotics, Epistemology, and Applications*. Papers and Commentary, STEM. Cambridge: Academic Press.
Wang, Jianxin, and Minxian Yang. 2009. Asymmetric volatility in the foreign exchange markets. *Journal of International Financial Markets, Institutions and Money* 19: 597–615. [CrossRef]
White, Ethan P., Brian J. Enquist, and Jessica L. Green. 2008. On estimating the exponent of power law frequency distributions. *Ecology* 89: 905–12. [CrossRef]

Quantitative Finance, 2025
https://doi.org/10.1080/14697688.2025.2529485

Modeling variance risk in financial markets using power-laws: new evidence from the Garman-Klass variance estimator

MASOUMEH FATHI†* and KLAUS GROBYS†‡**

†Finance Research Group, School of Accounting and Finance, University of Vaasa, Wolffintie 34, Vaasa 65200, Finland
‡Innovation and Entrepreneurship (InnoLab), University of Vaasa, Wolffintie 34, Vaasa 65200, Finland

(*Received 2 January 2025; accepted 26 June 2025*)

This study examines the range-based variance risk of five key financial asset markets—S&P 500, gold, crude oil, the USD/GBP exchange rate, and Bitcoin—using the noise-efficient Garman-Klass variance estimator. Our findings corroborate previous research by demonstrating that range-based asset variances adhere to power-law behavior generating variance behavior that is effectively infinite in practical terms. Furthermore, we provide novel evidence that the widely accepted log-normal model is unequivocally rejected for all range-based asset variances, underscoring its inadequacy in capturing the statistical properties of financial asset variances. A key contribution of this study is the discovery that a power-law function with $\alpha \approx 2.8$ represents a universal law governing the cross-sectional variances of otherwise unrelated asset markets. These findings have significant implications for risk management frameworks that rely on models developed within the mean-variance space, as they highlight the limitations of traditional approaches in assessing and managing financial risks.

Keywords: Garman-Klass estimator; Power-laws; Range-based variance; Second moment; Variance of variance

JEL Classification: C18, G10, G12, G14

1. Introduction

The seminal study by Markowitz (1952) laid the groundwork for portfolio analysis of financial assets, introducing the mean-variance framework as the cornerstone for optimizing portfolios of risky assets. Building on this foundation, Treynor (1962), Sharpe (1964), Lintner (1965), and Mossin (1966) developed the Capital Asset Pricing Model (CAPM), which remains a foundational concept in asset pricing literature. However, empirical failures of the CAPM have inspired the development of various factor models, culminating in what Feng *et al*. (2020) term a 'factor zoo.' Among these, the Fama-French factor models (Fama and French 1992, 1993, 2015, 2017, 2018, 2020) are widely recognized as benchmark models in empirical investigations aimed at addressing CAPM's limitations.

It is well understood that 'true' volatility is unobservable, and the literature offers various proxies to estimate it. These include the sample variance, conditional variance modeled through GARCH or stochastic volatility frameworks, and realized and range-based volatility derived from high-frequency or intraday price data. Range-based estimators, such as those proposed by Parkinson (1980) and Garman and Klass (1980), have been shown to be particularly efficient and robust, especially in the presence of market microstructure noise, and are widely used as nonparametric measures of realized variance (Chou *et al*. 2010).

While these factor models offer an appealing theoretical framework, their reliance on the mean-variance space presents critical challenges if the variance is statistically undefined. In his commentary on Mandelbrot's (1963) 'infinite variance hypothesis,' which profoundly disrupted traditional finance theory, Fama (1963) cautioned that statistical tools assuming finite variance—such as least-squares regression—could

*Corresponding author. Email: masoumeh.fathi@uwasa.fi

**Current address: Department of Monetary Economics and International Finance, Christian-Albrechts University (CAU) of Kiel, Wilhelm-Seelig Platz 1, 24118 Kiel, Germany

yield misleading results if variance is infinite. Subsequent evidence suggests that many proposed asset pricing factors may suffer from sample-specificity (Chan *et al*. 2000, Hirshleifer 2001, Schwert 2003). Extending this discourse, Grobys (2021) posited that replication failures in financial economics could stem from infinite variances or infinite variance of variances. Importantly, the qualitative implications of these phenomena are effectively identical.

This study aims to address a fundamental question: does the population mean of range-based variances in financial markets exist? Here, we focus on range-based realized variances, which serve as ex-post estimators of volatility and approximate the latent conditional variance of returns. Unlike model-based residual variance, which reflects unexplained variation after regression and depends on model assumptions, range-based variance is directly observable and model-free. If the distribution of these range-based variances exhibits extremely heavy tails, this implies that variance estimates themselves are unstable and potentially without finite moments. Consequently, the residual variance in linear regression models (e.g. CAPM or Fama-French), which requires the theoretical variance of the inputs to be statistically well-defined, may itself become undefined, thereby invalidating classical inference tools such as standard errors, t-statistics, and R^2. This makes the investigation of tail behavior in, for example, range-based variance distributions essential for assessing the reliability of asset pricing models built on OLS regression.

To this end, we investigate range-based variances in five major financial asset markets—S&P 500, gold, crude oil, the USD/GBP exchange rate, and Bitcoin—using intraday data. Range-based variances are computed using the Garman-Klass range-based variance estimator. We also fit power-law models to the data using maximum likelihood estimation (MLE) to examine the statistical properties of extreme variance realizations, with a specific focus on the tail behavior of the distribution beyond the range of typical observations. To assess model suitability, we conduct comprehensive goodness-of-fit tests, comparing the power-law hypothesis against other plausible distributions, such as log-normal and exponential models. Given the dependency structures in financial variance data (e.g. volatility clustering), we orthogonalize the range-based variance data through autoregressive modeling and reapply goodness-of-fit tests to the resulting innovation processes. Furthermore, we test for commonalities in the range-based variances of unrelated asset markets by examining whether power-law behavior is statistically invariant across these markets. Finally, we conduct an extensive set of robustness checks to validate our findings, including sample-split tests, the use of alternative volatility measures (e.g. realized variance), the examination of additional asset classes (e.g. bonds), the application of Bayesian inference to estimate tail exponents and cutoffs, model comparisons using Akaike Information Criterion (AIC) alongside the Kolmogorov–Smirnov test, rolling-window estimations, and the fitting of tempered stable and Generalized Beta of the Second Kind (GB2) distributions to assess the stability and generality of the observed power-law behavior.

This study makes several significant contributions to the literature, addressing critical methodological, theoretical, and empirical gaps. These contributions advance our understanding of range-based variances in financial markets and their implications for asset pricing and risk management. First, previous research, such as Grobys (2021), relied on the Parkinson variance estimator (Parkinson 1980) to compute range-based variances. While this approach has merits, Molnár (2012) highlighted that range-based variance estimators, such as those proposed by Garman and Klass (1980), outperform alternatives when accounting for noise in high-frequency data. Chou *et al*. (2010) also emphasized the informational advantage of intraday price ranges over closing prices. By adopting the Garman-Klass variance estimator, this study mitigates potential biases inherent in the use of inflated variance estimators, as in the study of Grobys (2021). Our methodological adjustment ensures greater precision in capturing range-based variances, offering a robust foundation for further analysis. This refinement is essential to derive reliable insights, particularly when assessing heavy-tailed distributions like power laws, and aligns with calls from Molnár (2012) for a more rigorous treatment of variance estimation methods.

Second, realized volatility has long been theorized as log-normal, with Andersen *et al*. (2001) providing seminal empirical support for this assumption. However, evidence challenging this perspective has been mounting. Renó and Rizza (2003), for instance, demonstrated that volatility in Italian futures markets aligns more closely with a Pareto distribution, necessitating the rejection of log-normality. Fathi *et al*. (2024) further confirmed power law distributions as superior in capturing realized variances in G10 currencies. Our study expands on this work by rigorously comparing the log-normal, exponential, and power-law models using comprehensive goodness-of-fit tests. We provide robust evidence that power laws offer the most plausible framework for modeling daily range-based variances too. By reconciling theoretical assumptions with observed data, this study extends Grobys' (2021) findings, which focused exclusively on power-law validity.

Third, Grobys (2021) used a relatively narrow sample, such as range-based variances for gold from August 30, 2000, to March 31, 2021, missing critical periods such as the gold price bubbles of the 1980s. Data omissions of this nature can lead to upward-biased estimates of power-law exponents, particularly if extreme events are underrepresented. To address this limitation, we extend the dataset for gold to cover February 7, 1980, to July 31, 2024, more than doubling the observation count to 11 196. This comprehensive data expansion ensures that periods of significant market stress are incorporated, mitigating biases and strengthening the reliability of our findings. By filling this empirical gap, the study builds a robust foundation for analyzing power-law behavior of range-based asset market variances.

Fourth, studies have explored common power-law behavior of range-based variances within related asset categories, such as equity factors (Fathi *et al*. 2025) and foreign exchange markets (Grobys 2024). These works documented universal exponents (e.g. $\alpha \approx 2.6$ for range-based variances of G10 currencies) that imply risk homogeneity within specific asset classes. However, no prior research has tested for such commonality across unrelated markets. This study pioneers cross-sectional risk analysis by examining whether the range-based variances

of diverse financial markets, including equities, commodities, exchange rates, and cryptocurrencies are governed by a common power law. The existence of such a commonality would suggest a shared underlying risk dynamic, likely driven by similar systemic or behavioral forces, such as herding behavior. The identification of a cross-market universal exponent would underscore the need for unified risk modeling approaches and contributes a novel dimension to the literature on power-law behavior.

Finally, robustness is a critical aspect of econometric analysis, especially when dealing with heavy-tailed distributions. Unlike Grobys (2021), which relied primarily on raw data, we employ autoregressive modeling to account for volatility clustering, a well-documented phenomenon in financial markets. Analyzing the resulting innovation processes confirms that power-law behavior persists even after accounting for dynamic dependencies. Furthermore, our sample-split tests validate the stability of power-law exponents across time periods and subsamples. These methodological advancements align with calls for more robust analytical frameworks in the financial econometrics literature (Mandelbrot 1963, Grobys 2021, Fathi *et al*. 2024) and enhance the generalizability of our findings.

Through these contributions, the study directly engages with and builds upon prior literature, addressing gaps in methodologies (Molnár 2012), theoretical assumptions (Andersen *et al*. 2001, Renó and Rizza 2003), and empirical scope (Grobys 2021, Fathi *et al*. 2024). By leveraging improved variance estimators, expanding datasets, and incorporating cross-sectional and dynamic robustness analyses, the research advances our understanding of range-based variances across financial markets. It also underscores the limitations of traditional models, advocating for the integration of power-law frameworks to better capture the realities of extreme market risks. These findings hold critical implications for both academic research and practical applications in risk management, portfolio optimization, and asset pricing.

The results reveal that Garman-Klass range-based variances exhibit heavy tails across all five markets, strongly supporting the power-law hypothesis. The MLE analysis of tail exponents confirms that values consistently fall within $2 < \alpha < 3$, implying infinite variance. This corroborates earlier findings (e.g. Grobys 2021, 2024, Fathi *et al*. 2024, 2025) and raises critical concerns about traditional econometric techniques relying on finite-variance assumptions. Goodness-of-fit tests firmly reject alternative distributions, including log-normal and exponential models, highlighting the systematic underestimation of tail risks by finite-moment frameworks. Moreover, striking evidence supports the existence of a universal cross-sectional power-law exponent ($\alpha \approx 2.8$) governing range-based variances across unrelated markets. Robustness checks using autoregressive modeling and sample-split tests confirm these findings, underscoring their consistency across time periods and market conditions.

This study provides critical insights into the statistical properties of range-based variances across key financial markets, offering important implications for finance theory and practice. First, the evidence of practically infinite variance challenges traditional econometric methods, such as point estimation and hypothesis testing, which assume finite variance, necessitating the development of alternative methodologies. Second, the rejection of the log-normal model in favor of power-law distributions highlights the systematic underestimation of tail risks by finite-moment frameworks, emphasizing the need for power-law-based approaches in assessing and managing extreme financial risks. Third, the identification of a universal power-law exponent across unrelated markets implies limited diversification benefits, as common dynamics across asset variances may amplify systemic risks, urging the adoption of unified risk modeling strategies. Fourth, the robustness of power-law exponents across time periods and subsamples, even after accounting for volatility clustering, confirms the generality and consistency of these findings, supporting their applicability to a wide range of financial contexts. Collectively, these results challenge prevailing assumptions underpinning modern portfolio theory and traditional risk management frameworks, reinforcing the necessity for alternative approaches that better capture the realities of extreme financial risks. By addressing methodological, theoretical, and empirical gaps, this study paves the way for more accurate risk modeling, portfolio optimization, and the development of financial instruments resilient to market extremes.

This study is organized as follows: Section 2 reviews the relevant literature; Section 3 describes the data; Section 4 outlines the methodology; Section 5 presents the empirical results and key findings; Section 6 discusses their implications; and Section 7 concludes the study.

2. Literature review

The modeling of volatility in financial markets has evolved into a rich and multifaceted area of research, resulting in a variety of theoretical and empirical frameworks. These models are designed to capture different aspects of volatility behavior, ranging from short-term clustering to long-memory and heavy-tailed properties. This literature can be grouped into six major research streams: GARCH-type models, stochastic volatility models, realized volatility models, range-based volatility models, multifractal and scaling models, and power-law (namely, tail-index) approaches.

2.1. Generalized Autoregressive Conditional Heteroskedasticity (GARCH) models

The Generalized Autoregressive Conditional Heteroskedasticity (GARCH) model, introduced by Bollerslev (1986), and its variants have become foundational in volatility modeling. These models describe volatility as a conditional process that evolves based on past squared returns and past variances. Key extensions include the Exponential GARCH (EGARCH) model by Nelson (1991) and the Threshold GARCH (TGARCH) model by Zakoian (1994), which allow for asymmetry in volatility responses. Although highly flexible, these models rely on parametric assumptions and are typically calibrated on daily or lower-frequency data. They perform well in capturing volatility clustering but struggle with representing extreme tail behavior unless augmented

with fat-tailed distributions for example, Student's t (Hansen 1994). A key limitation of GARCH-type models is that their parameter estimates can change substantially when underlying distributional assumptions are relaxed—for example, when the innovation process is altered from a normal to a Student's t distribution, or vice versa. In this regard, Mandelbrot (2008) criticized the tendency of researchers to 'fix' GARCH models when they failed to perform adequately, instead of abandoning them altogether.

2.2. Stochastic volatility (SV) models

In contrast to GARCH models, stochastic volatility (SV) models treat volatility as a latent stochastic process, independent from return innovations (Taylor 1982, Shephard 2005). These models provide greater flexibility in capturing volatility persistence and allow for richer dynamic structures. Recent developments include score-driven SV models (Creal *et al*. 2013, Harvey and Palumbo 2019), which use the score of the conditional likelihood to update latent volatility components. SV models are often computationally intensive but can better accommodate long-memory and structural breaks.

2.3. Realized volatility models

Realized volatility (RV) models utilize high-frequency intraday returns to construct ex-post volatility measures, offering a data-driven alternative to conditionally specified models. Andersen *et al*. (2001) played a pioneering role by formalizing the concept of realized volatility using the sum of squared intraday returns. This approach avoids strong distributional assumptions and allows researchers to observe and model actual volatility over fixed intervals, for example, daily or weekly periods. Subsequent studies have expanded upon this framework to examine properties for example, the distributional characteristics of RV (Barndorff-Nielsen and Shephard 2002), the impact of microstructure noise and asynchronous trading (Zhang *et al*. 2005), and the decomposition of volatility into continuous and jump components (Barndorff-Nielsen and Shephard 2004). Realized volatility is now a central tool in empirical finance, particularly in measuring and forecasting short-term volatility with high precision.

2.4. Range-based volatility models

Range-based models use daily high, low, open, and close prices to construct volatility estimators. The Parkinson (1980) and Garman and Klass (1980) estimators are seminal examples. These methods are shown to be more efficient than squared return-based measures, especially in the presence of microstructure noise and thin trading. Molnár (2012) and Chou *et al*. (2010) further highlight the robustness of range-based estimators, making them suitable for low-frequency but high-accuracy volatility measurement.

2.5. Multifractal and scaling models

Multifractal models attempt to capture the self-similar, long-memory, and scaling properties observed in financial volatility. The Markov-Switching Multifractal (MSM) model developed by Calvet and Fisher (2004) models volatility as a product of volatility components switching between different regimes at multiple frequencies. Earlier work by Mandelbrot *et al*. (1997) and Müller *et al*. (1997) emphasizes the hierarchical and fractal nature of market volatility, linking it to turbulence-like behavior in financial time series.

2.6. Power-law and tail-index approaches

A distinct stream focuses on the tail behavior of the volatility distribution. Rather than modeling conditional volatility, these approaches examine whether realized volatility or variance follows a power-law distribution with infinite mean or variance. Clauset *et al*. (2009) offer robust tools for tail-index estimation and goodness-of-fit testing. Recent empirical work (e.g. Grobys 2021, 2024) applies these methods to range-based variance data, exploring the implications of heavy-tailed volatility for the validity of classical econometric models for example, OLS regression.

While all these models aim to capture features of financial market volatility, they differ significantly in their assumptions and goals. GARCH and SV models are designed for conditional volatility forecasting, relying on parametric or semi-parametric methods. Realized and range-based models use empirical data to estimate volatility more accurately but often without strong assumptions about underlying distributions. Multifractal models and power-law approaches focus on structural properties like scaling laws and tail behavior. Notably, most traditional models do not explicitly address extremely fat-tailed distributions or long-range dependence in volatility, unless such features are specifically built into the model. The power-law approach is unique in explicitly testing for the existence (or lack) of moments, which has profound implications for econometric modeling and asset pricing.

2.7. Alternative approaches to address heavy tails and long memory in volatility

Notably, the literature on financial volatility offers a range of alternative approaches that contribute to understanding heavy-tailed behavior beyond traditional power-law modeling. For instance, Corsi (2009) introduced the Heterogeneous Autoregressive (HAR) model, which captures long-memory effects in realized volatility by aggregating volatility components across different time horizons. This approach has proven effective in modeling the persistent nature of volatility without requiring high-frequency data. Similarly, Harvey and Palumbo (2019) proposed a score-driven model for realized volatility that dynamically adapts to changes in the underlying data-generating process using a state-space framework. These methods highlight the growing emphasis on flexible and robust volatility modeling frameworks. Furthermore, Barndorff-Nielsen and Shephard (2002) introduced stochastic volatility models driven by Lévy processes, allowing for jumps and capturing the non-Gaussian nature of returns. Similarly, Müller *et al*. (1997) proposed a time-deformation model using multiplicative cascades, revealing scaling laws and volatility components at different frequencies. Engle and

Rangel (2008) addressed long-term volatility trends through a spline-GARCH model that combines smooth structural components with short-term GARCH effects. Calvet and Fisher (2004) developed the MSM model, which generates multifractal volatility patterns and long memory via regime-switching cascades. Bollerslev and Todorov (2011) emphasized jump-tail risk by integrating extreme events into asset price and volatility dynamics, accounting for the role of rare but impactful shocks.

Collectively, these models offer alternative perspectives on how volatility evolves over time under non-Gaussian conditions, and they share common objectives: capturing fat-tailed returns, long-range dependence, volatility clustering, and asymmetric responses to market shocks. While traditional GARCH-type models, including their extensions, are typically parametric and rely on specific distributional assumptions, several of the above frameworks (for example, MSM and multiplicative cascades) adopt semi- or non-parametric approaches, allowing for greater flexibility in capturing the empirical irregularities of financial time series. Moreover, models based on realized measures, for example, range-based or high-frequency volatility, circumvent strong distributional assumptions by leveraging more granular data. We have now referenced these studies in the revised manuscript to better position our research within this broader academic discourse and to clarify how our focus on power-law behavior and extreme variance dynamics complements, instead of contradicts, the insights offered by these alternative models.

While the existing literature offers valuable insights into the dynamics and persistence of volatility, our study diverges by explicitly focusing on the tail behavior of the volatility distribution—measured in terms of range-based variances—itself. Rather than modeling the full distribution or volatility path, this study investigates whether the extreme realizations of the data generating process—often dismissed as noise—may actually be governed by a power-law distribution, suggesting that the tail, instead of the center, contains the meaningful signal. This tail-centric perspective enables us to characterize the statistical properties of extreme variance events, a topic that remains underexplored in much of the existing volatility modeling literature.

3. Data

The data for this study were collected from Yahoo Finance. We retrieved daily observations including open, high, low, and closing prices on S&P 500, gold, crude oil, USD/GBP exchange rate, and Bitcoin. Due to data availability, the data sample for the S&P 500 spans from February 19, 1980, to July 31, 2024, providing 11 207 observations. This extensive period captures over four decades of market activity, including significant events such as the 1987 market crash, the dot-com bubble, the 2008 financial crisis, and the COVID-19 pandemic. For gold, the dataset covers a similarly long timeframe, from February 7, 1980, to July 31, 2024, resulting in 11 196 observations. This period reflects both periods of stability and heightened volatility, such as 2008 financial crisis. The crude oil dataset begins on January 2, 1991, and runs until July 31, 2024, including 8556 observations. Specifically, in our study we use data on West Texas Intermediate (WTI) which is one of the main global benchmark prices for crude oil. This period includes major fluctuations crude oil prices due to events such as the Gulf War, the shale revolution, and the COVID-19-related demand shock. The USD/GBP exchange rate data is available from July 25, 2000, to July 31, 2024, yielding 6263 observations. This timeframe captures significant currency market events, including the global financial crisis and Brexit. Finally, Bitcoin data is available from September 17, 2014, to July 31, 2024, with 3603 observations. This relatively short period reflects the cryptocurrency's promising stage and its rapid growth as a speculative and alternative asset, marked by extreme volatility.

To compute the Garman-Klass variance estimator, which combines open and close prices along with daily highs and lows, we follow the methodology proposed by Garman and Klass (1980) and compute daily range-based variances as:

$$\hat{\sigma}^2_{GK} = 0.5\left[\ln\left(\frac{H_t}{L_t}\right)\right]^2 - [2ln2 - 1]\left[\ln\left(\frac{C_t}{O_t}\right)\right]^2, \quad (1)$$

where H represents the daily high price, L denotes the daily low price, O is the daily opening price, and C refers the daily closing price of each financial asset.

GARCH models typically estimate conditional variances based on closing prices and thus may fail to fully capture the intraday dynamics of price movements (e.g. Andersen *et al*. 2003, 2004, Bubák *et al*. 2011). In contrast, range-based variances—particularly those derived from range-based estimators such as the Garman-Klass method—leverage intraday data (high, low, open, and close prices) to provide a more efficient and information-rich measure of daily volatility. As shown by Garman and Klass (1980) and reinforced by Molnár (2012), range-based estimators outperform return-based estimators, especially in the presence of microstructure noise and market frictions. Given our focus on modeling tail behavior and extreme variance risk, range-based variances offer clear advantages: they better capture the heavy-tailed nature of financial markets and are less prone to distributional misspecification. This allows for a more robust empirical foundation when testing power-law behavior and modeling the variance of variance—a key objective of our study.

Table 1 reports the descriptive statistics for the annualized daily range-based variances of these assets. The statistics show significant variations in the mean, standard deviation, skewness, and kurtosis, indicating the existence of heavy tails in the distribution of range-based variances, which could indicate power-law behavior. The mean of Garman-Klass range-based variance ranges from 0.95 for the USD/GBP exchange rate to 36.18 for Bitcoin, indicating substantial variation in the average level of volatility. Bitcoin's high Garman-Klass range-based variance mean value reflects its nature as a speculative and emergent asset class, characterized by extreme price fluctuations. Similarly, crude oil exhibits a high mean variance of 14.46. Conversely, traditional asset classes like gold (2.79) and the S&P 500 (2.18) show relatively moderate average variances, aligning with their roles as more stable investment assets. The USD/GBP exchange rate, with a

mean of 0.95, suggests that currency markets involving major economies tend to exhibit lower volatility. In addition, all assets exhibit significant positive skewness and excess kurtosis, suggesting non-normality in their range-based variance distributions. Crude oil exhibits the highest skewness and kurtosis (skewness: 88.83, kurtosis: 8095.44), while Bitcoin has the lowest skewness and kurtosis (skewness: 9.50, kurtosis: 130.54) among the analyzed assets.

4. Methodology

4.1. Main analysis

4.1.1. Estimating power-law exponents via MLE and comparing model fit. Following earlier studies (Grobys 2021, 2023, 2024, Fathi *et al*. 2024, 2025), we model the Garman-Klass range-based variances of each asset in our sample using a power-law function, expressed as:

$$p(x) = Cx^{-\alpha}, \tag{2}$$

where $C = (\alpha - 1)x_{MIN}^{\alpha-1}$ with $\alpha \in \{\mathbb{R}_+ | \alpha > 1\}$, and $x = \sigma_{j,t}^2$ denotes the respective range-based Garman-Klass variance of financial asset j at time t. Note that $x \in \{\mathbb{R}_+ | x_{MIN} \leq x < \infty\}$, where x_{MIN} is the minimum threshold for variance values governed by the power-law distribution, and α represents the tail exponent.

For range-based variances exceeding x_{MIN} ($X > x_{MIN}$), the conditional expected first moment is defined as:

$$E[X|X > x_{MIN}] = \int_{x_{MIN}}^{\infty} xp(x)dx = \frac{(\alpha - 1)}{(\alpha - 2)}x_{MIN}. \tag{3}$$

The conditional expected second moment, or the variance of range-based variances for $X > x_{MIN}$, is expressed as:

$$E[X^2|X > x_{MIN}] = \int_{x_{MIN}}^{\infty} x^2p(x)dx = \frac{(\alpha - 1)}{(\alpha - 3)}x_{MIN}^2. \tag{4}$$

Consequently, higher conditional expected moments of order k follow the general form:

$$E[X^k|X > x_{MIN}] = \frac{(\alpha - 1)}{(\alpha - 1 - k)}x_{MIN}^k. \tag{5}$$

From Equations (3) and (4), it can be concluded that the theoretical mean of range-based Garman-Klass variances is defined only when $\alpha > 2$, while the existence of the variance for range-based Garman-Klass variances requires $\alpha > 3$.

Next, as suggested by White *et al*. (2008) and Clauset *et al*. (2009), we utilized MLE to estimate the power-law tail exponents. The MLE estimator is given by:

$$\hat{\alpha} = 1 + N\left(\sum_{i=1}^{N} \ln\left(\frac{x_i}{x_{MIN}}\right)\right)^{-1}, \tag{6}$$

where $\hat{\alpha}$ is the estimated tail exponent, N is the number of observations exceeding x_{MIN}, and other terms maintain their previously defined meanings. The selection of x_{MIN} is determined by minimizing the distance between the empirical data and the power-law model (Clauset *et al*. 2009). This is achieved using the Kolmogorov–Smirnov (KS) distance, D, which measures the maximum difference between the cumulative distribution functions (CDFs) of the observed data and the fitted model. The distance D is given by:

$$D = MAX_{x \geq x_{MIN}} |S(x) - P(x)|, \tag{7}$$

where $S(x)$ represents the CDF of the empirical data for observations $x \geq x_{MIN}$, and $P(x)$ is the CDF of the fitted power-law model. The optimalx_{MIN} is the value that minimizes D. Additionally, as described by Clauset *et al*. (2009), the standard deviation of the estimated power-law exponents is expressed as:

$$\sigma = \frac{\hat{\alpha} - 1}{\sqrt{N}} + O\left(\frac{1}{N}\right). \tag{8}$$

To verify the appropriateness of a power-law fit, we conducted a goodness-of-fit test based on the KS distance, D, following the methodology derived in Clauset *et al*. (2009) which compares the empirical data to synthetic data generated from a power-law distribution using specific x_{MIN} and $\hat{\alpha}$ values. The null hypothesis of goodness-of-fit test is that the empirical data and the synthetic data generated from a power-law distribution using the estimated x_{MIN} and $\hat{\alpha}$ values belong to the same distribution, suggesting the power-law model is a plausible fit. While the alternative hypothesis suggests that the empirical and synthetic data originate from different distributions, implying that other models should be considered.

Thus, we generated synthetic datasets using the estimated parameters x_{MIN} and $\hat{\alpha}$ in the prior analysis. For each synthetic dataset, the distance (D_S) was calculated and compared to the distance (D) from the original dataset. The p-value for the goodness-of-fit test is defined as the proportion of synthetic datasets where $D_S > D$, denoted as $\#(D_S > D)/K$, where K is the total number of synthetic datasets. A significance level of 5% was applied, meaning the power-law model is not rejected if $\#D_S > D/K > 0.05$. In addition, to ensure robust results, 1000 synthetic datasets were simulated for each financial asset variance.

However, non-rejected power-law null models ascertained by means of Clauset *et al*.'s (2009) goodness-of-fit test do not necessarily mean that power-law distributions are the best fit for the data, as alternative distributions could potentially offer a better match. In addition, as mentioned earlier, prior studies such as Andersen *et al*. (2001), Andersen *et al*. (2001a, 2001b) have suggested that realized volatility often aligns closely with a log-normal distribution. To address this, we compare the power-law distributions with the log-normal distributions for each range-based asset variance to determine the most suitable model for the data. As noted by Clauset *et al*. (2009), comparing p-values of power-law models with those of alternative distributions, such as log-normal or exponential, provides a robust basis for evaluating the appropriateness of the power-law model. Thus, following the methodology of Clauset *et al*. (2009), we use the same goodness-of-fit test approach applied to the power-law distribution for assessing the plausibility of other competing distributions (viz.,

Table 1. Descriptive statistics.

Statistic	S&*P* 500	Gold	Crude oil	USD/GBP	BTC
Mean	2.18	2.79	14.46	0.95	36.18
Median	0.95	1.42	6.91	0.60	13.66
Maximum	161.31	200.71	15901.67	135.29	1809.48
Minimum	0.00	0.00	0.00	0.00	0.00
Std. Dev.	4.81	5.54	174.13	2.32	90.06
Skewness	14.32	11.90	88.83	35.59	9.50
Kurtosis	342.75	265.07	8095.44	1864.48	130.54
Observations	11207	11196	8556	6263	3603
Period (MM/DD/YYYY)	02/19/1980	02/07/1980	01/02/1991	07/25/2000	09/17/2014
	07/31/2024	07/31/2024	07/31/2024	07/31/2024	07/31/2024

Note: This table presents the descriptive statistics of the annualized daily range-based variance for various assets, including the S&*P* 500, gold, crude oil, USD/GBP exchange rate, and Bitcoin. The annualized daily range-based variances for each asset calculated using the Garman-Klass methodology (Garman and Klass 1980). The statistics reported in this table are sample statistics, which summarize empirical properties observed in the data. While our analysis later demonstrates that the variance of variance is theoretically infinite due to heavy tails ($\alpha \in (2,3)$), these sample-based measures remain meaningful for descriptive purposes. They do not imply the existence of finite theoretical moments.

log-normal, exponential). In this case, the null hypothesis assumes that the alternative distribution is a plausible fit for the data.

4.1.2. Assessing robust standard errors and testing for a common power-law component. Motivated by recent research documenting that range-based foreign exchange rate variances are governed by a universal power-law exponent (Grobys 2024), we explore this issue for the range-based variances of our five asset markets. An important difference between Grobys' (2024) study and our research is that the underlying asset markets in our study are otherwise unrelated, whereas foreign exchange rates are not. Therefore, a common power-law exponent governing the range-based variances of otherwise unrelated asset markets would be a surprising finding. To investigate this issue, we employ the test procedure proposed in the study of Grobys (2024).

While the MLE proposed by Clauset *et al*. (2009) remains asymptotically unbiased even under dependent data, the presence of autocorrelation—manifested, for example, through volatility clustering—can lead to an underestimation of standard errors when standard i.i.d. assumptions are applied. This would result in overly narrow confidence intervals and potentially misleading inference. To address this issue, we implement a block bootstrap method, as recommended by Godfrey (2009) and employed by Grobys (2024), which preserves temporal dependencies within blocks. This adjustment produces more realistic, wider standard errors and thereby ensures that hypothesis testing related to tail exponents is appropriately conservative; that is, our block bootstrap approach mitigates the issue of downward-biased standard errors due to serial dependence. Hence, to address this issue, we first estimate robust standard errors for each range-based asset variance. This is achieved by employing the blocks bootstrap methodology, as outlined in the study of Grobys (2024).

Specifically, a block bootstrap procedure is implemented with the block length m, where $E[m] = T^{1/3}$ following the recommendation by Godfrey (2009) for variance estimation. From a given data vector of range-based variances, $\boldsymbol{x}$, blocks of size m are randomly selected, where m follows a geometric distribution, $m \sim GEO(p)$ with an expected value $E[m] = (1-p)/p$. Thus, in the overall data sample, we use $p = 0.0435, p = 0.0434, p = 0.0476, p = 0.0526$ and $p = 0.0625$ for the range-based variances of the S&P 500, gold, crude oil, USD/GBP exchange rate, and Bitcoin, respectively. Using this bootstrap method, blocks drawn from the data vector $\boldsymbol{x}$ have varying lengths. These randomly selected blocks m are then stacked into a new vector $\boldsymbol{x}^b$ as follows:

$$\boldsymbol{x}_i^b = \begin{bmatrix} m_1 \\ m_2 \\ m_3 \\ \vdots \end{bmatrix}. \tag{9}$$

Again, it is important to note that the lengths of the blocks $m_1, m_2, \ldots$ vary. The process continues until the length of the constructed vector $\boldsymbol{x}^b$ exceeds T. Then, we cut off any observations exceeding T, ensuring that the artificial data vector $\boldsymbol{x}^b$ matches the length of the original data vector $\boldsymbol{x}$. By applying this blocks bootstrap, 1000 artificial data vectors $\begin{bmatrix} \boldsymbol{x}^1 & \boldsymbol{x}^2 & \cdots & \boldsymbol{x}^B \end{bmatrix}$, are generated for each given original data vector. Next, the point estimates $\hat{\alpha}$ for each bootstrap data vector $\begin{bmatrix} \boldsymbol{x}^1 & \boldsymbol{x}^2 & \cdots & \boldsymbol{x}^B \end{bmatrix}$, are calculated and stored in a vector $\begin{bmatrix} \hat{\alpha}^1 & \hat{\alpha}^2 & \cdots & \hat{\alpha}^B \end{bmatrix}$ following the method described by Clauset *et al*. (2009). As a final step, the bootstrapped standard error $\hat{\sigma}_{BOOT}$ is calculated from $\begin{bmatrix} \hat{\alpha}^1 & \hat{\alpha}^2 & \cdots & \hat{\alpha}^B \end{bmatrix}$ for each given data vector of range-based asset variances as:

$$\hat{\sigma}_{BOOT} = \sqrt{\frac{1}{B}\sum_{b=1}^{B}\left(\hat{\alpha}^b - \left(\frac{1}{B}\sum_{b=1}^{B}\hat{\alpha}^b\right)\right)^2}. \tag{10}$$

After obtaining the robust standard errors, we can conduct the joint test to examine whether the Garman-Klass range-based variances of different assets share a common component, represented by a common power-law exponent.

Identifying such a shared exponent would suggest a unified risk of these assets. It is worth noting that our bootstrapping method ensures $COV(\hat{\alpha}_i, \hat{\alpha}_j) = 0 \forall i \neq j$. Based on this, the covariance matrix for estimated power law exponents is defined as:

$$\hat{\boldsymbol{\Sigma}}_{\hat{\boldsymbol{\alpha}}} = \begin{pmatrix} \hat{\sigma}^2_{\hat{\alpha}_{boot},S\&P500} & 0 & 0 & 0 & 0 \\ 0 & \hat{\sigma}^2_{\hat{\alpha}_{boot},Gold} & 0 & 0 & 0 \\ 0 & 0 & \hat{\sigma}^2_{\hat{\alpha}_{boot},Crudeoil} & 0 & 0 \\ 0 & 0 & 0 & \hat{\sigma}^2_{\hat{\alpha}_{boot},USD/GBP} & 0 \\ 0 & 0 & 0 & 0 & \hat{\sigma}^2_{\hat{\alpha}_{boot},BTC} \end{pmatrix}, \tag{11}$$

where, $\hat{\sigma}^2_{\hat{\alpha}_{boot},S\&P500}, \hat{\sigma}^2_{\hat{\alpha}_{boot},Gold}, \hat{\sigma}^2_{\hat{\alpha}_{boot},Crudeoil}$, $\hat{\sigma}^2_{\hat{\alpha}_{boot},USD/GBP}$ and $\hat{\sigma}^2_{\hat{\alpha}_{boot},BTC}$ represent the corresponding bootstrapped variances (viz., $\hat{\sigma}^2_{BOOT}$). Then, consistent with Grobys (2024), we define the test statistic $\hat{\lambda}$ as follows:

$$\hat{\lambda} = (\hat{\boldsymbol{\alpha}} - q\mathbf{1})' \hat{\boldsymbol{\Sigma}}^{-1}_{\hat{\boldsymbol{\alpha}}} (\hat{\boldsymbol{\alpha}} - q\mathbf{1}), \tag{12}$$

where the covariance matrix $\hat{\boldsymbol{\Sigma}}_{\hat{\boldsymbol{\alpha}}}$ is 5×5, $\hat{\boldsymbol{\alpha}}$ represents a 5×1 vector of estimated power-law exponents, $\mathbf{1}$ is a 5×1 vector of ones, and q denoted the hypothesized common power-law exponent. The estimated test statistic $\hat{\lambda}$ in Eq. (10) follows a $\chi^2(5)$ distribution under the null hypothesis, with the corresponding critical value at the 5% significance level being $\chi^2_{0.95}(5) = 11.07$. To evaluate whether the Garman-Klass variances of our five asset markets share a common power-law exponent, the test statistic $\hat{\lambda}$ is iteratively computed across the economically significant range of q values, spanning from 1.2 to 3.2 in increments of 0.1. This iterative computation ensures that the analysis may capture the potential existence of a common power-law exponent governing variance risks across the selected assets.

4.2. Additional analysis

4.2.1. Investigating the impact of autocorrelation on tail index estimates in range-based variances. A potential concern could be whether the observed power-law behavior in the Garman-Klass range-based variance distributions of our assets is genuine or a statistical effect resulting from autocorrelation. It is well-recognized that time-varying volatility can create patterns in the tails of return distributions that resemble power-laws. This phenomenon could lead to inaccurate conclusions, as standard statistical techniques for estimating power laws might incorrectly attribute thick tails to a power-law structure when they are actually produced by a combination of exponential-tail distributions rather than an actual power law. This raises the important question as to whether similar issues affect the estimation of power-law behavior in the range-based variances of our assets?

To examine this, we fit autoregressive models of order p to each asset's Garman-Klass range-based variance. The models are defined as:

$$\hat{\sigma}^2_{j,t} = \beta_{0,j} + \beta_{1,j}\hat{\sigma}^2_{j,t-1} + \beta_{2,j}\hat{\sigma}^2_{j,t-2} + \ldots + \beta_{p,j}\hat{\sigma}^2_{j,t-p} + \varepsilon_{j,t}, \tag{13}$$

where $\hat{\sigma}^2_{j,t}$ defines the Garman-Klass range-based variance of asset j at time t, $\varepsilon_{j,t}$, represents the innovation process, and p is the lag-order. The optimal lag order p is determined by analyzing the partial autocorrelation function. This approach is similar to ARCH-type models, where the conditional variance is modeled using the lagged values of range-based variances directly. Next, we use the absolute values of the estimated innovation processes of each asset, $|\hat{\boldsymbol{\varepsilon}}_{S\&P500}|$, $|\hat{\boldsymbol{\varepsilon}}_{Gold}|$, $|\hat{\boldsymbol{\varepsilon}}_{Crudeoil}|$, $|\hat{\boldsymbol{\varepsilon}}_{USD/GBP}|$, $|\hat{\boldsymbol{\varepsilon}}_{BTC}|$, to re-evaluate the power-law exponents for each asset's orthogonalized range-based variance process. Then, we carried out Clauset *et al.*'s (2009) goodness-of-fit test, following the methodology outlined earlier for the absolute values of the innovation processes for each asset variance. This procedure helps to evaluate whether the observed power-law behavior persists after accounting for autocorrelation, providing a robust check on the validity of the power-law model for our range-based variances.

4.2.2. Evaluating the reliability of power-law exponents: evidence from sample-split tests. To further validate the robustness of the observed power-law behavior in the innovation processes of the range-based variances, we perform a sample-split test. We divide the residual series for each asset, obtained from the autoregressive models (viz., $|\hat{\boldsymbol{\varepsilon}}_{S\&P500}|$, $|\hat{\boldsymbol{\varepsilon}}_{Gold}|$, $|\hat{\boldsymbol{\varepsilon}}_{Crudeoil}|$, $|\hat{\boldsymbol{\varepsilon}}_{USD/GBP}|$, $|\hat{\boldsymbol{\varepsilon}}_{BTC}|$), into two non-overlapping subsamples of equal length. By estimating the power-law exponents for the absolute values of the residuals in each subsample separately, we aim to assess the consistency of the power-law characteristics across different sample episodes. This approach allows us to examine whether the power-law behavior persists independently within each subsample, mitigating concerns about potential sample-specificity. Additionally, we applied the goodness-of-fit tests to the absolute values of the residuals in each subsample to evaluate the validity of the power-law model in both cases.

4.2.3. Alternative measures of asset market variance: A realized variance approach. While the main analysis in this paper is based on the Garman-Klass range-based estimator, concerns can be raised about the potential limitations of relying on a single volatility measure. Specifically, it has been noted that different volatility estimators may capture different aspects of return dynamics, and even the most efficient estimators remain noisy approximations of true volatility (Shu and Zhang 2006, Molnár 2012, Będowska-Sójka and Kliber 2021). To address this concern and assess the robustness of our findings, we extend our analysis by computing weekly realized variances based on the sum of squared daily returns over five-day intervals. This return-based volatility measure offers a conventional benchmark and helps assess whether the heavy-tailed behavior observed under the Garman-Klass estimator persists under a different volatility estimation framework. The daily return of asset market i at time t is computed as follows:

$$RET_{i,t} = 100\frac{(P_{i,t+1} - P_{i,t})}{P_{i,t}}, \tag{14}$$

where $P_{i,t}$ denotes the corresponding daily closing price. Next, we compute the weekly realized variance of each five assets as the sum of the squared returns over five business days according to the following Equation:

$$RV_{j,t} = \sigma_{i,j}^2 = \sum_{t \in j} (RET_{i,t})^2, \tag{15}$$

where j indicates the week and $t \in j$ indicates the corresponding trading days in the respective week. We then estimate power-law exponents for the tail behavior of the weekly realized variances of each asset as described in section 4.1.1. Finally, to evaluate the appropriateness of the power-law model relative to alternative specifications, we compare the empirical data with synthetic data generated from power-law, log-normal, and exponential distributions for each realized asset variance to determine the most suitable model for the weekly data as described in section 4.1.2.

4.2.4. Analysis of additional asset classes: evidence from bonds. To test the generalizability of our findings across asset classes, we extend the analysis to U.S. government bond yields. Specifically, we include range-based bond yield variance data for the 2-year, 5-year, and 10-year Treasury yields, computed using the Garman-Klass estimator on daily high, low, open, and close prices. This addition allows us to explore whether power-law behavior is also present in fixed income markets, which are often considered less volatile than equities or cryptocurrencies. Following the methodology outlined in Section 4.1.1, we estimate power-law exponents for each bond series and conduct comparative goodness-of-fit tests using synthetic data from power-law, log-normal, and exponential distributions. This enables a robust evaluation of the tail behavior in bond market volatility.

4.2.5. Alternative estimation of power-law exponents and cutoffs: A Bayesian inference approach. While our main results rely on the widely used Kolmogorov–Smirnov-based method for estimating power-law parameters, one might be concerned that the selection of the threshold x_{min}—a crucial modeling choice—could influence the inferred tail behavior. To address this potential sensitivity and enhance the credibility of our findings, we perform a robustness check using an alternative, Bayesian estimation framework. This approach follows the methodology introduced by Virkar and Clauset (2014) and Grigaityte and Atwal (2019), which allows for the simultaneous estimation of the power-law exponent $\hat{\alpha}$ and the lower-bound threshold x_{min}. Unlike frequentist methods that rely on point estimates and deterministic threshold selection, the Bayesian approach yields posterior distributions for both parameters. Using Markov Chain Monte Carlo (MCMC) sampling, we generate estimates of $\hat{\alpha}$ and x_{min}. This method not only enhances transparency around threshold selection but also enables a more informative inference by incorporating parameter uncertainty. To formally describe the Bayesian approach, we consider a sample of continuous observations $x_1, x_2, \ldots, x_N$, each satisfying $x_i \geq x_{min}$. The continuous power-law distribution is defined by the following probability density function:

$$p(x|\alpha, x_{min}) = \frac{\alpha - 1}{x_{min}} \left(\frac{x}{x_{min}} \right)^{-\alpha}, \text{for } x \geq x_{min}, \alpha > 1. \tag{16}$$

Given N observations above a threshold x_{min}, the likelihood ($\mathcal{L}$) of the data is:

$$\mathcal{L}(\alpha \mid x, x_{min}) = \prod_{i=1}^{N} \frac{\alpha - 1}{x_{min}} \left(\frac{x_i}{x_{min}} \right)^{-\alpha}. \tag{17}$$

Taking the log, the log-likelihood becomes:

$$log\ \mathcal{L}(\alpha \mid x, x_{min}) = Nlog(\alpha - 1) - Nlogx_{min} - \alpha \sum_{i=1}^{N} log \left(\frac{x_i}{x_{min}} \right). \tag{18}$$

To complete the Bayesian model, we specify prior distributions where for the exponent α, we use a weakly informative uniform *prior:*

$$\alpha \sim Uniform(1.01, 5). \tag{19}$$

In addition, for the threshold x_{min}, we assume a discrete uniform prior over a range of candidate values:

$$x_{min} \in \{x^{(1)}, x^{(2)}, \ldots, x^{(k)}\}, P(x_{min}) = \frac{1}{k}, \tag{20}$$

where, $P(x_{min})$ denotes the prior probability assigned to each candidate threshold, reflecting an assumption of equal plausibility before observing the data. In this setting, the joint posterior distribution over α and x_{min} is given by:

$$P(\alpha, x_{min}|x) \propto \mathcal{L}(\alpha|x, x_{min}) \cdot P(\alpha) \cdot P(x_{min}), \tag{21}$$

where $P(\alpha)$ is the prior distribution over the power-law exponent α. Because x_{min} is discrete and α is continuous, we marginalize over the possible values of x_{min} by computing:

(1) The posterior distribution $P(\alpha|x, x_{min})$ for each candidate x_{min}.
(2) The marginal likelihood (also known as the model evidence):

$$\mathcal{Z}(x_{min}) = \int \mathcal{L}(\alpha|x, x_{min}) \cdot P(\alpha) d\alpha. \tag{22}$$

(3) The posterior over x_{min} using Bayes' rule:

$$P(x_{min}|x) = \frac{\mathcal{Z}(x_{min})}{\sum_{k=1}^{K} \mathcal{Z}(x^k)}. \tag{23}$$

The overall posterior for α can be computed as a weighted mixture of posteriors conditioned on each x_{min}, using the posterior probabilities $P(x_{min}, x)$ as weights. Finally, the power-law exponent and threshold are being estimated by performing MCMC sampling for each candidate ${x_{min}}^{(k)}$, compute their marginal likelihoods, derive posterior probabilities for each threshold, and calculate the weighted posterior mean.

4.2.6. An alternative model comparison metric: the Akaike Information Criterion. In addition to the previously reported goodness-of-fit tests based on Clauset *et al.* (2009), we incorporated a formal model comparison framework to strengthen the case for power-law behavior relative to alternative models. Specifically, we compare the power-law, log-normal, and exponential distributions using the AIC. Unlike the standalone goodness-of-fit tests, AIC provides a systematic approach that balances model fit and complexity through the MLE. For each candidate distribution, model parameters are estimated via MLE. The power-law parameters—the scaling exponent $\widehat{\alpha}$ and the lower bound $\hat{x}_{min}$—are determined following Clauset *et al.* (2009). The log-likelihood for each model is computed based on the corresponding probability density function and the AIC is calculated as:

$$AIC = 2k - 2log(\hat{L}), \tag{24}$$

where k represents the number of parameters (with $k = 1$ for the power-law and exponential models, and $k = 2$ for the log-normal model) and $log(\hat{L})$, is the maximum log-likelihood of the fitted model.

4.2.7. Rolling window estimations for retrieving time-Varying power-law exponents. One could argue that financial volatility may be subject to structural changes over time, potentially challenging the validity of our results from the sample-split tests. In this context, implementing rolling window estimations offers a complementary approach to assess the stability of the estimated power-law exponents. To address this issue, we performed rolling window estimations as follows: for each series of range-based variances, we applied a rolling window of 500 observations and estimated the MLE-based power-law exponent as outlined in Section 4.1.1. The window advances in steps of 50 observations, starting from $t = 500$ until $t = T$.

4.2.8. Estimation of tempered stable distributions. A finding of a shared power-law exponent governing all range-based asset variances would represent a significant theoretical insight. While it may be argued that certain financial assets exhibit distinct distributional characteristics—particularly in response to varying market conditions—such evidence would motivate deeper investigation. Therefore, in this robustness check, we assess whether the tempered stable distribution can offer a more nuanced understanding of the common dynamics underlying range-based variances. A tempered stable distribution generalizes classical stable distributions (e.g. Lévy) by introducing exponential tempering to the heavy tails. This modification allows for finite moments while preserving power-law behavior over intermediate ranges. Formally, the tail decays as

$$P(X > x) \sim x^{-\alpha} e^{-\lambda x}, \tag{25}$$

where $\alpha \in (0, 2)$ controls the tail heaviness and $\lambda > 0$ governs the rate of tempering. This construction retains the scaling properties of stable laws while ensuring statistical regularity in the extreme tails. Tempered stable distributions are particularly useful for modeling empirical phenomena that exhibit approximate power-law scaling but experience faster-than-power decay in the tails, as frequently observed in financial, physical, and biological systems. In applications involving multiple datasets that share a common power-law exponent yet display distinct truncation or decay rates, the tempered stable framework provides a coherent modeling strategy. By fixing a common α across data series and allowing data series-specific tempering parameters λ_i, one can jointly capture universal scaling behavior and dataset-specific tail dynamics. This approach is particularly suited for comparative studies or hierarchical modeling contexts such as volatility dynamics. To estimate the parameters via maximum likelihood, we numerically invert the characteristic function, as the probability density function lacks a closed-form expression. Parameters are then estimated by maximizing the log-likelihood using numerical optimization techniques (BFGS).

4.2.9. Comparison between power-law distributions and the generalized beta of the second kind (GB2) family. One might argue that the comparison of power-law distributions with the exponential and log-normal distributions introduced in Section 4.1.1 does not provide an informative benchmark for evaluating the appropriateness of the power-law model. As a consequence, we consider the Generalized Beta of the Second Kind (GB2) family as a more flexible alternative, as it accommodates both fat tails and finite second moments across a wide range of parameterizations. The GB2 is a highly flexible, four-parameter family of continuous probability distributions that can represent a wide variety of shapes, including both heavy-tailed and light-tailed forms. It is frequently employed in the analysis of skewed or heavy-tailed data. The GB2 distribution is defined as:

$$f(x; a; b; p; q) = \frac{a(x/b)^{ap-1}}{bB(p,q)[1 + (x/b)^a]^{p+q}}, \tag{26}$$

where $B(p, q)$ denotes the Beta function, $a > 0$ is the tail parameter controlling the tail heaviness, $b > 0$ is a scale parameter, and $p > 0$, $q > 0$ determine the shape of the distribution, including skewness. In Table A.5 of the appendix, we list several important subfamilies of the GB2 distribution, highlighting its role as a 'distributional umbrella' that encompasses many well-known distributions. Asymptotically, the GB2 exhibits power-law decay under certain conditions:

$$P(X > x) \sim x^{-aq}, \tag{27}$$

as $x \to \infty$, indicating that the product aq serves as an effective tail exponent, thereby linking the GB2 to power-law modeling. For a valid comparison between the GB2 and power-law models, both must be estimated using maximum likelihood, and the likelihoods must be evaluated over the same data range—specifically, the range above the power-law cutoff x_{MIN}. Accordingly, we restrict the GB2 model to the same data ranges used for the power-law estimation, as reported in table 3. We then compute the log-likelihood values and the corresponding AIC for each model.

5. Results

5.1. Main results

This section presents the findings from our analysis, applying the methodological approach outlined in Section 4.1 to the selected financial assets: S&P 500, gold, crude oil, USD/GBP exchange rate, and Bitcoin.

The analysis begins by modeling the Garman-Klass range-based variances using the power-law function described in Section 4.1.1. Table 2 highlights the share of cumulative total variance contributed by the top 1% and top 20% of observations for each asset, revealing a significant concentration. A small proportion of extreme observations accounts for a disproportionate share of total variance, consistent with the heavy-tailed behavior expected under a power-law distribution. For instance, among the top 1% of observations, crude oil exhibits the highest concentration at 27.23%, followed by Bitcoin (20.03%), S&P 500 (16.10%), gold (15.05%), and USD/GBP exchange rate (14.39%). Similarly, for the top 20%, Bitcoin leads with 72.03%, followed by crude oil (66.15%), S&P 500 (65.91%), gold (63.00%), and USD/GBP exchange rate (55.45%). These results underscore the presence of extreme events driving range-based variances, particularly in crude oil and Bitcoin, aligning with the Pareto 80/20 rule.

Next, a power-law distribution is fitted to the annualized daily range-based variances of each asset using MLE, as detailed in Section 4.1.1 and proposed by White *et al.* (2008) and Clauset *et al.* (2009). Table 3 reports the tail exponents, ranging from $\hat{\alpha} = 2.62$ for crude oil to $\hat{\alpha} = 2.94$ for the S&P 500, indicating that variance of variance is infinite for all assets. Such heavy tails invalidate foundational assumptions underlying conventional statistical tools like *t*-tests, which depend on finite variances. Standard metrics derived from ordinary least squares regression, such as *t*-statistics, become unreliable due to their sample-specific dependency. The evidence confirms that heavy-tailed distributions render traditional econometric methods inadequate, consistent with critiques by Grobys (2021) and Fathi *et al.* (2025).

To further validate the plausibility of the power-law model, goodness-of-fit tests (Clauset *et al.* 2009) are applied to assess the plausibility of power-law and alternative distributions, including log-normal and exponential models. Table 4 reports *p*-values for these tests, demonstrating strong support for the power-law model, as all *p*-values for the null hypothesis of a power-law distribution exceed 5%. In contrast, the log-normal model is rejected outright ($p = 0.00$) for all assets except Bitcoin, where the power-law model ($p = 0.66$) outperforms log-normal ($p = 0.52$). The exponential model is conclusively ruled out, with *p*-values equating to zero for all cases. These results collectively confirm the superior fit of the power-law model to the data.

A joint test, described in Section 4.1.2, examines whether the Garman-Klass range-based variances exhibit a shared component characterized by a common power-law exponent. To account for dependencies such as volatility clustering, block bootstrap methods are employed, yielding robust standard errors. Table 5 illustrates the increased standard deviations derived from bootstrapping. For instance, the standard deviation for S&P 500 increases from 0.05 to 0.20, and for gold, from 0.06 to 0.14. These adjustments underscore the critical importance of incorporating dependency structures in estimating tail exponents.

Subsequently, the joint test evaluates whether the assets share a common power-law exponent. From table 6, the test statistic $\hat{\lambda}$, computed iteratively over a range of α-values

Table 2. Share of the top 1% and top 20%.

Metric	S&*P* 500	Gold	Crude oil	USD/GBP	BTC
Top 1% Share	16.10%	15.05%	27.23%	14.39%	20.03%
Top 20% Share	65.91%	63.00%	66.15%	55.45%	72.03%

Note: This table displays the share of the cumulative total represented by the top 1% and top 20% of the distribution for the annualized daily Garman-Klass range-based variance across the S&*P* 500, gold, crude oil, the USD/GBP exchange rate, and Bitcoin.

Table 3. Estimation of power-law exponents for annualized daily range-based variance.

Metric	S&*P* 500	Gold	Crude oil	USD/GBP	BTC
$\hat{\alpha}$	2.94	2.74	2.62	2.78	2.84
$\hat{x}_{MIN}$	4.59	6.62	23.97	1.09	187.71
D	0.02	0.02	0.02	0.01	0.04
N_{PL}	11.93%	8.33%	10.30%	22.77%	3.41%
Std. Dev	0.05	0.06	0.05	0.04	0.17
Observations	11207	11196	8556	6263	3603
Period (MM/DD/YYYY)	02/19/1980	02/07/1980	01/02/1991	07/25/2000	09/17/2014
	07/31/2024	07/31/2024	07/31/2024	07/31/2024	07/31/2024

Note: This table presents the estimation results for power-law exponents applied to the annualized daily range-based variances of the S&*P* 500, gold, crude oil, the USD/GBP exchange rate, and Bitcoin. The parameter $\hat{\alpha}$ signifies the estimated tail exponent, while $\hat{\sigma}$ denotes the estimated standard deviation. The lower threshold $\hat{x}_{MIN}$ is identified through the optimized Kolmogorov-Smirnov distance (D), following the procedure described by Clauset *et al.* (2009). The reported $\hat{x}_{MIN}$ corresponds to the optimal distance D. Additionally, the fraction of observations obtained by the power-law process, denoted as N_{PL} is provided for each asset.

Table 4. Goodness-of-fit tests for the annualized daily range-based variance.

Asset	Power-law	Log-normal	Exponential
S&*P* 500	0.25	0.00	0.00
Gold	0.39	0.00	0.00
Crude oil	0.40	0.00	0.00
USD/GBP	0.46	0.00	0.00
BTC	0.66	0.52	0.00

Note: This table presents the results of the goodness-of-fit tests conducted using the Kolmogorov-Smirnov (KS) method to assess whether the empirical data and synthetic data generated from a power-law distribution, specified by x_{min} and α, originate from the same underlying distribution (column 2). Additionally, the table includes the results of goodness-of-fit tests for log-normal and exponential distributions in columns 3 and 4, respectively. The null hypothesis for these tests states that the empirical data and the synthetic data generated from the specified distributions (log-normal and exponential) are consistent, suggesting that each distribution provides an acceptable fit to the observed data.

from 1.2 to 3.2, fails to reject the null hypothesis of a shared exponent within the economically important range of $2.5 < \alpha < 3.1$. This finding suggests a unified component influencing variances. Notably, our results suggest that $\hat{\alpha} \approx 2.8$ produces the highest p-value ($p = 0.87$), which indicates that the second moment of the common power law governing the variance risk of unrelated financial asset markets is, statistically, infinite.

5.2. Results from robustness checks

To ensure the observed power-law behavior is not an artifact of statistical effects such as autocorrelation, robustness checks are performed. As discussed in Section 4.2.1, autoregressive models of order p are fitted to each asset's Garman-Klass range-based variance. Table 7 reports the point estimates, confirming significant autocorrelation across all assets, with partial autocorrelation functions indicating orders of $p = 2$, 6, 1, 4 for S&P 500, gold, crude oil, USD/GBP exchange rate, and Bitcoin, respectively. Descriptive statistics for the resulting innovation processes, detailed in table 8, reveal extremely heavy tails. For instance, kurtosis values range from 133.83 for Bitcoin to 8053.8 for crude oil, reinforcing the persistence of extreme values even after autoregressive adjustments.

Table 9 presents power-law exponents derived from the absolute innovation processes, closely matching those in table 3, further substantiating the presence of genuine power-law behavior. The goodness-of-fit tests applied to the innovation processes (table 10) corroborate these findings, failing to reject the power-law null hypothesis for all assets except gold, where the evidence is slightly weaker. Log-normal and exponential distributions are once again rejected outright ($p = 0.00$) across all assets.

Next, sample-split tests are conducted to validate the robustness of power-law characteristics across subsamples. Table 11 provides power-law exponent estimates for non-overlapping subsamples, demonstrating consistent statistical

Table 5. Descriptive statistics for the power-law exponents estimated using block bootstrap methods.

Statistic	$\hat{\alpha}^{boot}_{\text{S\&P 500}}$	$\hat{\alpha}^{boot}_{Gold}$	$\hat{\alpha}^{boot}_{Crudeoil}$	$\hat{\alpha}^{boot}_{USD/GBP}$	$\hat{\alpha}^{boot}_{BTC}$
Mean	2.95	2.74	2.61	2.81	2.33
Median	2.95	2.74	2.59	2.79	2.27
Maximum	3.57	3.37	3.48	3.75	3.44
Minimum	2.38	2.37	2.26	2.38	1.91
Std. Dev	0.20	0.14	0.17	0.18	0.25
Skewness	0.01	0.31	1.01	0.88	1.63
Kurtosis	2.80	3.42	4.94	4.51	5.93
Jarque-Bera (JB)	1.70	23.08	326.81	225.13	802.29
(p-value) JB	0.41	0.00	0.00	0.00	0.00
Observations	11207	11196	8557	6263	3603

Note: This table reports the block bootstrap methodology test to the annualized daily range-based variance for the S&*P* 500, gold, crude oil, USD/GBP exchange rate, and Bitcoin. The block bootstrap procedure is implemented with the block length m, where $E[m] = T^{1/3}$ following the recommendation by Godfrey (2009) for variance estimation. From the given data vector $\mathbf{x}$, blocks of size m are randomly selected, where m follows a geometric distribution $m \sim GEO(p)$ with an expected value $E[m] = \frac{(1-p)}{p}$. Thus, in the overall data sample, we use $p = 0.0435$, $p = 0.0434$, $p = 0.0476$, $p = 0.0526$ and $p = 0.0625$ for the range-based variances of the S&*P* 500, gold, crude oil, USD/GBP exchange rate, and Bitcoin, respectively. Using this bootstrap method, blocks drawn from the data vector x have varying lengths. The randomly selected blocks, m, are then stacked into a new vector

$$\boldsymbol{x}_i^b = \begin{bmatrix} m_1 \\ m_2 \\ m_3 \\ \vdots \end{bmatrix}$$

where the lengths of the blocks $m_1, m_2, \ldots$ vary. The process continues until the length of the constructed vector $\mathbf{x}^b$ exceeds T. Then, we cut off any observations exceeding T, ensuring that the artificial data vector $\boldsymbol{x}^b$ matches the length of the original data vector $\boldsymbol{x}$. By applying this blocks bootstrap, 1000 artificial data vectors $\begin{bmatrix} \boldsymbol{x}^1 & \boldsymbol{x}^2 & \cdots & \boldsymbol{x}^B \end{bmatrix}$, are generated for each given original data vector. Next, the point estimates $\hat{\alpha}$ for each bootstrap data vector $\begin{bmatrix} \boldsymbol{x}^1 & \boldsymbol{x}^2 & \cdots & \boldsymbol{x}^B \end{bmatrix}$, are calculated and stored in a vector $\begin{bmatrix} \hat{\alpha}^1 & \hat{\alpha}^2 & \cdots & \hat{\alpha}^B \end{bmatrix}$ following the method described by Clauset *et al*. (2009). The block bootstrap method is described in detail in Section 3.1.2.

Table 6. Testing for a common power-law exponent governing from annualized daily range-based variance.

q	$\hat{\lambda}$	p-value
1.2	390.68	**0.00**
1.3	342.52	**0.00**
1.4	297.53	**0.00**
1.5	255.73	**0.00**
1.6	217.11	**0.00**
1.7	181.68	**0.00**
1.8	149.42	**0.00**
1.9	120.34	**0.00**
2	94.45	**0.00**
2.1	71.74	**0.00**
2.2	52.21	**0.00**
2.3	35.87	**0.00**
2.4	22.70	**0.00**
2.5	12.72	**0.03**
2.6	5.91	0.31
2.7	2.29	0.81
2.8	1.85	0.87
2.9	4.60	0.47
3	10.52	0.06
3.1	19.63	**0.00**
3.2	31.92	**0.00**

Note: This table presents the results of a joint test designed to assess whether the range-based variances of S&*P* 500, gold, crude oil, USD/GBP exchange rate, and Bitcoin share a common component. To investigate whether a common component underlies the power-law behavior of these assets, the following test statistic is utilized:

$$\hat{\lambda} = (\hat{\boldsymbol{\alpha}} - q\mathbf{1})' \hat{\boldsymbol{\Sigma}}_{\hat{\boldsymbol{\alpha}}}^{-1} (\hat{\boldsymbol{\alpha}} - q\mathbf{1}),$$

where $\hat{\boldsymbol{\Sigma}}_{\hat{\boldsymbol{\alpha}}}$ is the estimated 5×5 covariance matrix, $\hat{\boldsymbol{\alpha}}$ is a 5×1 vector of estimated power-law exponents, $\mathbf{1}$ is a 5×1 vector of ones, and q represents the hypothesized power-law exponent. The test statistic $\hat{\lambda}$ follows a $\chi^2(5)$ distribution under the null hypothesis. The statistic is iteratively calculated across the economically significant range of $q = (1.2, 1.3, \ldots, 3.1, 3.2)$. Bolded values in the results represents statistical significance at the 5% level.

properties across both panels. Exceptions include the S&P 500 in Panel A and gold in Panel B, where the goodness-of-fit test rejects the power-law hypothesis. Nevertheless, the log-normal and exponential models remain consistently rejected across all subsamples ($p = 0.00$). This consistency across tests and subsamples strengthens the conclusion that genuine power-law behavior governs the range-based variances of these financial assets.

Building on the robustness analysis outlined in Section 4.2.3, we report the results derived from the weekly realized variance estimator in this section. Table A.1 presents the descriptive statistics for the weekly realized variances across the five asset classes. Consistent with our earlier findings, the realized variance distributions exhibit substantial skewness and excess kurtosis—hallmarks of heavy-tailed behavior. Weekly realized variances for crude oil and Bitcoin, in particular, display considerable variation and extreme upper-tail values, further supporting the presence of volatility clustering and rare but impactful events. We then estimate power-law exponents for the tail behavior of the weekly realized variances using the MLE approach proposed by Clauset *et al*. (2009). As reported in Table A.2 the estimated exponents all lie within the critical range $2 < \hat{\alpha} < 3$, confirming that the variance of variance remains theoretically undefined for all asset markets. These results align closely with those obtained using the Garman-Klass estimator, indicating that our primary conclusions are not sensitive to the choice of volatility measure. To further validate the plausibility of the power-law model under this alternative volatility measure, we conduct KS goodness-of-fit tests for power-law, log-normal, and exponential distributions (Table A.3). Based on the results in Table A.3, the power-law model remains statistically plausible for all assets (p-values > 0.05), except gold, while the exponential model is decisively rejected across all the assets. Additionally, although the log-normal model is also accepted for all assets (p-values > 0.05), the power-law model generally provides a superior fit, as reflected in higher p-values—except in the case of gold. In summary, the persistence of power-law behavior under both range-based and return-based volatility estimators confirms the robustness of our core results. These findings also underscore the broader insight that volatility risk in financial markets may be governed by a common, heavy-tailed structure—regardless of the specific estimation method employed.

Additionally, as described in Section 4.2.4, we extended the analysis to incorporate government bond markets in order to examine whether the statistical properties of realized variances also apply to U.S. Treasury yields with 2-year, 5-year, and 10-year maturities. This extension allows us to assess the broader applicability of power-law behavior to fixed income markets. The results of this extended analysis are presented in tables 12–14. As shown in table 12, the descriptive statistics for the range-based variances of the 2-year, 5-year, and

Table 7. Point estimates for autoregressive models of for annualized daily range-based variance.

	$\hat{\beta}_{0,j}$	$\hat{\beta}_{1,j}$	$\hat{\beta}_{2,j}$	$\hat{\beta}_{3,j}$	$\hat{\beta}_{4,j}$	$\hat{\beta}_{5,j}$	$\hat{\beta}_{6,j}$	R^2
$\hat{\sigma}^2_{S\&P500,t}$	0.49***	0.42***	0.10	0.08***	0.18***			0.42
	(12.27)	(44.89)	(10.33)	(7.81)	(18.90)			
$\hat{\sigma}^2_{Gold,t}$	0.74***	0.36***	0.07**	0.03*	0.11***	0.08***	0.09***	0.30
	(7.33)	(8.71)	(2.53)	(1.69)	(2.64)	(2.95)	(3.58)	
$\hat{\sigma}^2_{Crudeoil,t}$	13.34***	0.08***						0.01
	(6.09)	(3.37)						
$\hat{\sigma}^2_{USD/GBP,t}$	0.48***	0.16***	0.13**	0.09***	0.11***			0.10
	(3.52)	(3.11)	(2.11)	(3.29)	(3.12)			
$\hat{\sigma}^2_{BTC,t}$	20.35***	0.44***						0.19
	(10.73)	(9.93)						

*, ** and *** indicate statistical significance at the 10%, 5% and 1% levels, respectively.
Note: This table presents the estimated coefficients from autoregressive models of order P fitted to the annualized daily range-based variance for the S&P 500, gold, crude oil, the USD/GBP exchange rate, and Bitcoin (BTC). The models are specified as

$$\hat{\sigma}^2_{j,t} = \beta_{0,j} + \beta_{1,j}\hat{\sigma}^2_{j,t-1} + \beta_{2,j}\hat{\sigma}^2_{j,t-2} + \ldots + \beta_{p,j}\hat{\sigma}^2_{j,t-p} + \varepsilon_{j,t},$$

where $\varepsilon_{j,t}$ is the innovation process, j refers to the respective asset. The optimal lag order P is determined by the partial autocorrelation function. The reported coefficients include the intercept $\hat{\beta}_{0,j}$ and autoregressive terms ($\beta_{0,j}$, $\beta_{1,j}$, $\beta_{2,j}$, … , $\beta_{p,j}$). The t-statistics are reported in parentheses.

Table 8. Descriptive statistics for the innovations process of annualized daily range-based variance.

Statistic	$\hat{\boldsymbol{\varepsilon}}_{S\&P500}$	$\hat{\boldsymbol{\varepsilon}}_{Gold}$	$\hat{\boldsymbol{\varepsilon}}_{Crudeoil}$	$\hat{\boldsymbol{\varepsilon}}_{USD/GBP}$	$\hat{\boldsymbol{\varepsilon}}_{BTC}$
Mean	0.00	0.00	0.00	0.00	0.00
Median	− 0.40	− 0.63	− 7.14	− 0.24	− 14.29
Maximum	139.02	163.67	15818.40	132.17	1750.00
Minimum	− 70.72	− 56.60	− 1048.3	− 21.49	− 676.04
Std. Dev	3.68	4.64	173.62	2.20	80.99
Skewness	11.83	10.18	88.46	38.32	8.33
Kurtosis	377.28	250.67	8053.8	2145.90	133.83
Observations	11203	11190	8555	6259	3602

Note: This table summarizes the descriptive statistics for the innovation process $\varepsilon_{i,t}$ of the annualized daily range-based variance for the S&P 500, gold, crude oil, the USD/GBP exchange rate, and Bitcoin (BTC), as reported in Table 3. The innovation process of annualized daily range-based variances for all five key assets is estimated by running autoregressive models of order p:

$$\hat{\sigma}^2_{j,t} = \beta_{0,j} + \beta_{1,j}\hat{\sigma}^2_{j,t-1} + \beta_{2,j}\hat{\sigma}^2_{j,t-2} + \ldots + \beta_{p,j}\hat{\sigma}^2_{j,t-p} + \varepsilon_{j,t},$$

where $\varepsilon_{j,t}$ is the innovation process, p is the lag-order, j refers to the respective asset. The optimal lag order p is determined through an analysis of the partial autocorrelation function.

10-year Treasury yields exhibit high levels of skewness and kurtosis, indicating heavy-tailed distributions similar to those observed in the five major financial asset markets. The power-law estimation results (table 13) show tail exponents of U.S. Treasury yields with 2-year, 5-year, and 10-year maturities, all falling within the critical interval $2 < \hat{\alpha} < 3$. This implies that the second moment of range-based variances for these bonds remains undefined, in line with our findings for major financial asset markets. In addition, goodness-of-fit test results in table 14 further confirm the plausibility of the power-law model for bond market variances. The null hypothesis of a power-law distribution is not rejected for any of the three bond series (p-values: 0.94, 0.54, and 0.58, respectively). In contrast, the log-normal and exponential models are decisively rejected for all three bond series. These results provide compelling evidence that power-law behavior is a persistent and robust characteristic of range-based variances across diverse financial markets—including government bond markets. This suggests that heavy-tailed variance risk may be a universal feature of modern financial systems, further reinforcing the need to reconsider traditional models based on finite-variance assumptions.

Next, to complement our primary estimation strategy and assess the robustness of the results, we apply a Bayesian inference approach to estimate the power-law parameters across all assets as described in section 4.2.5. Table 15 reports the Bayesian estimates of the power-law exponent $\hat{\alpha}$, the inferred lower-bound $\hat{x}_{min}$. The results confirm that all

Table 9. Estimating power-law exponents for the absolute amount of the innovation processes of the annualized daily range-based variance.

Metric	$\lvert\hat{\varepsilon}_{S\&P500}\rvert$	$\lvert\hat{\varepsilon}_{Gold}\rvert$	$\lvert\hat{\varepsilon}_{Crudeoil}\rvert$	$\lvert\hat{\varepsilon}_{USD/GBP}\rvert$	$\lvert\hat{\varepsilon}_{BTC}\rvert$
$\hat{\alpha}$	2.54	2.46	2.41	2.40	2.30
$\hat{x}_{MIN}$	3.71	2.81	40.90	1.12	56.55
D	0.02	0.02	0.03	0.03	0.04
N_{PL}	5.51%	14.59%	2.58%	6.65%	9.69%
Std. Dev	0.06	0.04	0.09	0.07	0.07
Observations	11203	11190	8555	6259	3602

Note: This table presents the estimation results for the power-law exponents applied to the absolute amount of the residual series of annualized daily range-based variances for the S&*P* 500, gold, crude oil, the exchange rate of the U.S. dollar and Bitcoin (BTC). The innovation process for each asset is obtained by using the following autoregressive model:

$$\hat{\sigma}^2_{j,t} = \beta_{0,j} + \beta_{1,j}\hat{\sigma}^2_{j,t-1} + \beta_{2,j}\hat{\sigma}^2_{j,t-2} + \ldots + \beta_{p,j}\hat{\sigma}^2_{j,t-p} + \varepsilon_{j,t}$$

where $\varepsilon_{j,t}$ is the innovation process, p is the lag-order, and j refers to the respective asset. The parameter $\hat{\alpha}$ represents the estimated tail exponent, while $\hat{\sigma}$ denotes the estimated standard deviation. The lower threshold $\hat{x}_{MIN}$ is determined using the optimized Kolmogorov-Smirnov distance (D) approach, as outlined by Clauset *et al*. (2009). The reported $\hat{x}_{MIN}$ corresponds to the optimal distance D. Additionally, the fraction of observations governed by the power-law process, denoted as N_{PL} is provided for each asset.

Table 10. Goodness-of-fit tests for the innovation processes of annualized daily range-based variance.

	Power-law	Log-normal	Exponential
$\lvert\hat{\varepsilon}_{S\&P500}\rvert$	0.82	0.00	0.00
$\lvert\hat{\varepsilon}_{Gold}\rvert$	0.02	0.00	0.00
$\lvert\hat{\varepsilon}_{Crudeoil}\rvert$	0.88	0.00	0.00
$\lvert\hat{\varepsilon}_{USD/GBP}\rvert$	0.22	0.00	0.00
$\lvert\hat{\varepsilon}_{BTC}\rvert$	0.07	0.00	0.00

Note: This table presents the results of the goodness-of-fit tests conducted using the Kolmogorov-Smirnov (KS) method to assess whether the empirical data and synthetic data generated from a power-law distribution, specified by x_{min} and α, originate from the same underlying distribution (column 2). Additionally, the table includes the results of goodness-of-fit tests for log-normal and exponential distributions in columns 3 and 4, respectively. The null hypothesis for these tests states that the empirical data and the synthetic data generated from the specified distributions (log-normal and exponential) are consistent, suggesting that each distribution provides an acceptable fit to the observed data.

estimated exponents fall within the range typically associated with heavy-tailed behavior $2 < \hat{\alpha} < 3$, consistent with our main findings based on Clauset *et al*.'s (2009) approach. Importantly, the 95% credible intervals for $\hat{\alpha}$ across all assets are relatively narrow, indicating precise and stable estimates of the power-law exponent under the Bayesian framework. In addition, figure 1 visualizes the posterior distribution of the power-law exponent $\hat{\alpha}$ for each asset, illustrating the concentration of posterior mass around the mean and the precision of the Bayesian estimates. This robustness reinforces the credibility of our earlier conclusions regarding tail behavior in financial range-based variances.

To complement the goodness-of-fit tests, we conducted a formal model comparison using the AIC, as described in Section 4.2.6. The AIC accounts for both model fit and complexity, with lower values indicating a more parsimonious and better-fitting model. Each distribution—power-law, log-normal, and exponential—was fitted to the empirical data using MLE. Table 16 reports the AIC values for each model across the five asset classes. In all cases, the power-law distribution produces the lowest AIC, indicating a superior fit to the extreme tails of the range-based variance distributions when compared to the log-normal and exponential alternatives. These findings provide consistent and robust evidence that power-law models are better suited to capture the heavy-tailed nature of range-based asset variance distributions. The results reinforce the conclusions drawn from the goodness-of-fit tests and underscore the relevance of power-law behavior in modeling extreme financial events.

Furthermore, the results of the rolling-window estimations outlined in Section 4.2.7 are presented in Table A.4 while Figure A.1 visualizes the evolution of the time-varying power-law exponents over the sample period. Notably, the S&P 500 dataset contains $T = 11\,207$ observations of range-based variance, allowing for the estimation of $N = 205$ power-law exponents. In contrast, for Bitcoin, only $N = 53$ exponents could be estimated due to the shorter sample. A comparison between the full-sample estimates (table 3) and the rolling window medians (Table A.4) reveals close alignment. For example, Table A.4 reports median power-law exponents of 2.73 and 2.77 for the S&P 500 and Gold, respectively, while the corresponding full-sample point estimates in table 3 are 2.94 and 2.74. Similar consistency is observed across other range-based variance series. Finally, a visual inspection of Figure A.1 suggests that the time-varying power-law exponents derived from the rolling window estimation appear stationary, reinforcing the conclusion that the exponent estimates are stable—that is, converging.

Next, we fit tempered stable distributions, as outlined in Section 4.2.8, to the range-based variances of the S&P 500, USD/GBP exchange rate, gold, crude oil, and Bitcoin, initially fixing the power-law exponent $\alpha_i = \alpha = 2.8$. The results

Table 11. Sample split analysis of power-law exponents for the innovation processes of annualized daily range-based variance.

Panel A: Estimation results for the first subsample of power-law exponents and related metrics.					
Metric	S&*P* 500	Gold	Crude oil	USD/GBP	BTC
$\hat{\alpha}$	2.18	2.48	2.29	2.30	2.44
$\hat{x}_{MIN}$	0.65	2.82	33.07	1.24	26.81
$(p\text{-value})_{\text{power-law}}$	0.00	0.27	0.28	0.96	0.10
D	0.02	0.02	0.04	0.03	0.04
N_{PL}	48.79	16.93	5.10	5.08	18.17
Std. Dev	0.02	0.05	0.09	0.10	0.08
$(p\text{-value})_{\text{log-normal}}$	0.00	0.00	0.00	0.00	0.00
$(p\text{-value})_{\text{exponential}}$	0.00	0.00	0.00	0.00	0.00
Observations	5601	5595	4277	3129	1801
Panel B: Estimation results for the second subsample of power-law exponents and related metrics.					
Metric	S&*P* 500	Gold	Crude oil	USD/GBP	BTC
$\hat{\alpha}$	2.76	2.31	2.49	2.47	2.21
$\hat{x}_{MIN}$	3.13	1.04	13.42	0.92	49.73
$(p\text{-value})_{\text{power-law}}$	0.85	0.00	0.39	0.12	0.40
D	0.02	0.02	0.04	0.04	0.03
N_{PL}	7.10	43.66	5.35	9.90	14.71
Std. Dev	0.09	0.03	0.10	0.08	0.07
$(p\text{-value})_{\text{log-normal}}$	0.00	0.00	0.00	0.00	0.00
$(p\text{-value})_{\text{exponential}}$	0.00	0.00	0.00	0.00	0.00
Observations	5602	5595	4278	3130	1801

Note: This table presents the sample-split test of power-law exponents applied to the residual series of annualized daily range-based variances, with results divided into two subsamples: Panel A reports the first subsample, and Panel B reports the second subsample. The parameter $\hat{\alpha}$ represents the estimated tail exponent, while $\hat{\sigma}$ denotes the estimated standard deviation. The lower threshold $\hat{x}_{MIN}$ is determined using the optimized Kolmogorov-Smirnov distance (D) approach, as outlined by Clauset *et al*. (2009). The reported $\hat{x}_{MIN}$ corresponds to the optimal distance D. Additionally, the fraction of observations obtained by the power-law process, denoted as N_{PL} is provided for each asset.

Table 12. Descriptive statistics of range-based U.S. Treasury yield variances.

Statistic	UST 2Y	UST 5Y	UST 10Y
Mean	51.41	35.47	21.70
Median	14.50	12.28	6.97
Maximum	3042.13	2975.99	5041.98
Minimum	0.00	0.00	0.00
Std. Dev.	157.46	114.40	120.87
Skewness	10.02	12.83	28.73
Kurtosis	133.63	233.29	1087.59
Observations	3010	3033	2818
Period (MM/DD/YYYY)	10/26/2012	10/26/2012	08/26/2013
	07/31/2024	07/31/2024	07/31/2024

Note: This table presents the descriptive statistics of the annualized daily range-based variances for U.S. Treasury yields with maturities of 2 years (UST 2Y), 5 years (UST 5Y), and 10 years (UST 10Y). The variances are calculated using the Garman-Klass estimator, based on daily high, low, open, and close yields.

reveal that λ_i could be estimated only for the range-based variance for the USD/GBP exchange rate, while the estimation failed to converge for the remaining series. To validate this finding, we repeated the estimation without constraining α. Again, both α_i and λ_i could be reliably estimated only for the range-based variance for the USD/GBP series. This pattern is highly informative. It suggests that, for most assets, the data do not exhibit statistically significant exponential truncation in the tails. Rather, the data appear to conform closely to a pure power law over the observed range, rendering the inclusion of a tempering parameter either superfluous or numerically unstable. This interpretation aligns with earlier empirical evidence of strong power-law decay in these series. The consistent failure to estimate λ_i even when α_i is unconstrained further supports the conclusion that the tail behavior is well captured by a standard power law without the need for exponential tempering.

Finally, restricting the data to $x \geq x_{MIN}$, we fit GB2 models to the range-based variances and computed the AIC values, as outlined in Section 4.2.9. Table 17 presents the

Table 13. Estimation of power-law exponents for range-based U.S. Treasury yield variances.

Metric	UST 2Y	UST 5Y	UST 10Y
$\hat{\alpha}$	2.30	2.31	2.32
$\hat{x}_{MIN}$	87.42	56.62	27.45
D	0.02	0.03	0.02
N_{PL}	12.82%	12.43%	14.05%
Std. Dev	0.07	0.07	0.07
Observations	3010	3033	2818
Period (MM/DD/YYYY)	02/19/1980	02/07/1980	01/02/1991
	07/31/2024	07/31/2024	07/31/2024

Note: This table presents the estimation results for power-law exponents applied to the annualized daily range-based variances of U.S. Treasury yields with 2-year (UST 2Y), 5-year (UST 5Y), and 10-year (UST 10Y) maturities. The parameter $\hat{\alpha}$ signifies the estimated tail exponent, while $\hat{\sigma}$ denotes the estimated standard deviation. The lower threshold $\hat{x}_{MIN}$ is identified through the optimized Kolmogorov-Smirnov distance (D), following the procedure described by Clauset *et al*. (2009). The reported $\hat{x}_{MIN}$ corresponds to the optimal distance D. Additionally, the fraction of observations obtained by the power-law process, denoted as N_{PL} is provided for each bond.

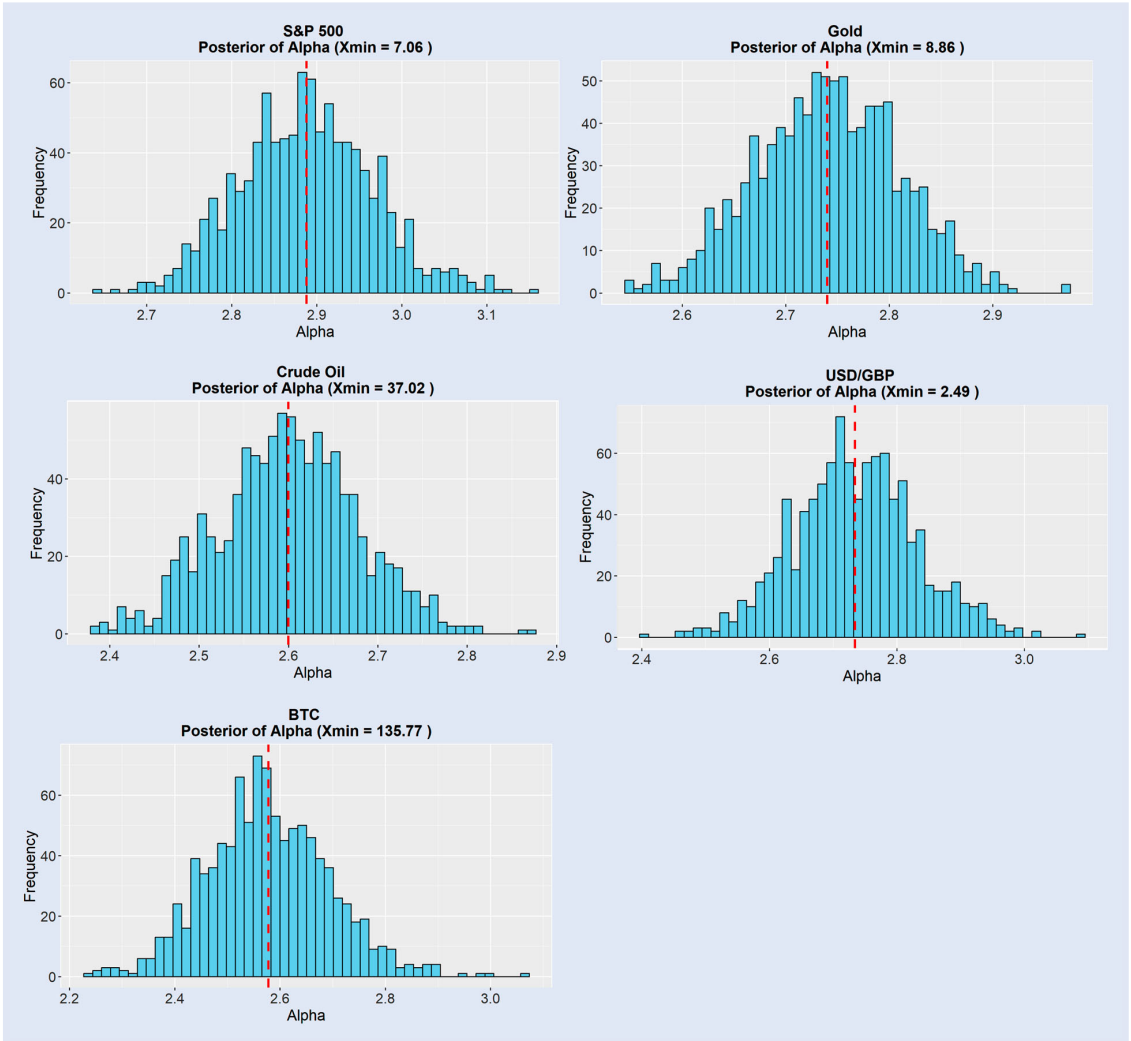

Figure 1. Posterior distribution of the power-law exponent $\hat{\alpha}$, estimated using Bayesian inference. Note: This figure shows the posterior distribution of the power-law exponent $\hat{\alpha}$, estimated using Bayesian inference for values greater than the threshold $\hat{x}_{min}$. The red dashed line indicates the posterior mean of $\hat{\alpha}$.

Table 14. Goodness-of-fit tests for range-based U.S. Treasury yield variances.

Asset	Power-law	Log-normal	Exponential
UST 2Y	0.94	0.04	0.00
UST 5Y	0.54	0.00	0.00
UST 10Y	0.58	0.00	0.00

Note: This table reports the results (viz., p-values) of the goodness-of-fit tests for the annualized daily range-based variances of U.S. Treasury yields with maturities of 2 years (UST 2Y), 5 years (UST 5Y), and 10 years (UST 10Y). The tests are conducted using the Kolmogorov-Smirnov (KS) method to assess whether the empirical data and the synthetic data—generated from a power-law distribution specified by x_{min} and α —originate from the same underlying distribution (column 2). Additionally, the table includes the results of goodness-of-fit tests for log-normal and exponential distributions in columns 3 and 4, respectively. The null hypothesis for these tests states that the empirical data and the synthetic data generated from the specified distributions (log-normal and exponential) are consistent, suggesting that each distribution provides an acceptable fit to the observed data.

log-likelihoods and AIC values for both the power-law models—using the parameterizations from table 3—and the GB2 models restricted to the equivalent data ranges. Parameter estimates for the GB2 models are provided in Table A.6 in the appendix. The results in table 17 indicate that the AIC values are consistently lower for the power-law models across all range-based variance series, regardless of the underlying asset class. This finding reinforces the results of our main analysis, further supporting the conclusion that power-law models are particularly well suited for capturing the dynamics of range-based variances in financial assets.

6. Discussion

In this study we explored whether the Garman-Klass range-based variances of five key financial assets—S&P 500, gold, crude oil, USD/GBP exchange rate, and Bitcoin—follow power-law distributions and examined the broader implications of these findings for financial modeling and risk analysis. By utilizing the Garman-Klass variance estimator and conducting robust statistical tests, we provided a comprehensive analysis of power-law behavior in range-based variances of otherwise unrelated asset markets. Unlike previous studies that primarily focused on absolute returns or single-asset analyses, as highlighted by Lux and Alfarano (2016), we employ the Garman-Klass range-based variances, which provide greater sensitivity to intraday price dynamics and mitigate estimation bias. This methodological improvement enables us to capture the complexity of volatility structures across different asset classes more effectively.

First, we confirmed the heavy-tailed nature of range-based variances across all five assets, consistent with power-law behavior. Data on range-based asset variances exhibited

Table 15. Bayesian estimation for power-law exponents.

Metric	S&P 500	Gold	Crude oil	USD/GBP	BTC
$\hat{\alpha}$	2.89	2.74	2.60	2.73	2.58
$\hat{x}_{MIN}$	7.06	8.86	37.02	2.49	135.77
Std. Dev	0.08	0.07	0.08	0.09	0.12
CI $(\hat{\alpha})_{95\%}$	(2.73, 3.05)	(2.60, 2.88)	(2.44, 2.76)	(2.55, 2.91)	(2.34, 2.82)
Observations	11207	11196	8556	6263	3603
Period (MM/DD/YYYY)	02/19/1980	02/07/1980	01/02/1991	07/25/2000	09/17/2014
	07/31/2024	07/31/2024	07/31/2024	07/31/2024	07/31/2024

Note: This table presents the estimation results for power-law exponents using Bayesian approach applied to the annualized daily range-based variances of the S&P 500, gold, crude oil, the USD/GBP exchange rate, and Bitcoin (BTC). The table includes the posterior means of the exponent $\hat{\alpha}$ the estimated lower-bound threshold $\hat{x}_{min}$,the standard deviations of the posterior distributions for $\hat{\alpha}$, and the 95% credible intervals for $\hat{\alpha}$. The number of observations and the corresponding sample period for each asset are also provided.

Table 16. AIC model comparison.

Asset	Power-law*	Log-normal	Exponential
S&*P* 500	**6346.79**	36394.16	39929.06
Gold	**5433.58**	42811.29	45326.74
Crude oil	**7592.74**	58142.31	62010.57
USD/GBP	**3044.78**	10295.22	11925.14
BTC	**1518.62**	31259.29	33017.41

*Bold values indicate the best-fitting model by AIC.
Note: This table presents the AIC values for power-law, log-normal, and exponential distributions fitted to the annualized range-based variances for each five assets, including the S&*P* 500, gold, crude oil, USD/GBP exchange rate, and Bitcoin (BTC). Each model is estimated using maximum likelihood estimation methods. The AIC is calculated as $AIC = 2k - 2log(\hat{L})$, where k is the number of estimated parameters and $log(\hat{L})$ is the maximum log-likelihood of the model. Lower AIC values indicate a better trade-off between model fit and complexity.

Table 17. Comparison between power-law and GB2 functions.

Range-based variance	Log-likelihood Power law (AIC)	Log-likelihood GB2 (AIC)
S&*P* 500	–1134.28 (2268.57)	–3176.18 (6360.37)
Gold	–953.78 (1907.56)	–2714.71 (5437.42)
Crude oil	–996.82 (1993.63)	–3805.39 (7618.79)
USD/GBP	–1397.88 (2795.76)	–1528.31 (3064.61)
Bitcoin	–114.42 (228.85)	–782.70 (1573.40)

Note: We compare the GB2 distribution with our power-law models by fitting the GB2 function to the data range $x \geq x_{min}$ for each respective data vector, as reported in Table 3. The values of the log-likelihood functions are computed, along with the corresponding AIC. This table reports the log-likelihood and AIC values for both the power-law models—using the parameterizations documented in Table 3—and the GB2 models, which are estimated over the same data ranges defined by the respective power-law cutoffs.

significant skewness and kurtosis, with extreme values concentrated in the upper tails of the variance distributions. For example, Bitcoin exhibited the highest mean value of Garman-Klass range-based variance, reflecting its speculative and volatile nature, while the USD/GBP exchange rate displayed relatively lower mean value of Garman-Klass range-based variance.

Second, power-law fitting through MLE verified that tail exponents for all assets fell within the range $2 < \alpha < 3$, implying that the variance-of-variance is infinite. Interestingly, these findings contrast with literature from the early 1990s, as discussed by Lux and Alfarano (2016), which focused on the tail behavior of distributions and suggested that financial asset returns exhibit power-law behavior with finite variances. This result has profound implications, as it challenges traditional econometric assumptions that rely on finite variances for hypothesis testing. The estimated power-law exponent α has direct implications for economic risk. Lower values of α indicate heavier tails in the return distribution, implying a higher probability of extreme losses. This suggests a more volatile and risk-prone market environment. For instance, when α lies between 2 and 3, the theoretical mean of the variance process is defined but higher-order moments such as variance or skewness are undefined, posing challenges for traditional risk models. When α approaches 2 or below, the theoretical mean of the variance process becomes infinite, highlighting severe tail risk. Hence, the power-law exponent serves as a quantifiable indicator of the intensity of tail events, informing asset pricing, risk management, and regulatory capital frameworks.

This finding also aligns with the previous literature (Grobys 2021, 2023, 2024, Fathi *et al.* 2024, 2025) which documented tail exponents for realized variances or range-based variances are within the range $2 < \alpha < 3$, indicating that the variance-of-variances is infinite. In the presence of infinite variance-of-variance, extreme observations can disproportionately influence the variance estimate, leading to results that may appear significant in one sample but fail to replicate in others. As highlighted by Fama (1963), if variances are undefined, results derived from traditional statistical methods based on the concept of correlation (i.e. *t*-statistics) become inherently sample-dependent. According to Grobys (2021), unstable second moments lead to *t*-statistics that are highly dependent on the specific sample used. Thus, this study, in line with Grobys (2021), argues that the replication failures observed in financial research may stem from the infinite variance of variances, which leads to traditional statistical methods becoming unreliable.

Third, our results derived from goodness-of-fit tests strongly supported power-law distributions as plausible models for the range-based variances of these assets. These findings align with Mandelbrot's (1963) early insights that financial asset returns may lack finite variances, and strengthens the notion that traditional statistical tools, such as OLS and other moment-dependent methods, may be inadequate for modeling financial data governed by power-law distributions.

Fourth, our analysis comprehensively compared power-law distributions with alternative models (viz., log-normal and exponential). The goodness-of-fit test strongly supported power-law models as the most plausible representation of the data for all assets, with alternative distributions being consistently rejected. This finding strengthens the case for power-law distributions as a fundamental characteristic of financial asset variances. The rejection of log-normal and exponential models aligns with Taleb's (2020) view on the necessity of adopting power-law frameworks for analyzing financial data, particularly in the context of extreme events and tail risks. Furthermore, these findings contrast with the literature suggesting that realized asset volatility follows a log-normal distribution (e.g. Andersen *et al.* 2001, 2001a, 2001b), but they align with studies proposing that a Pareto distribution better describes volatility in financial markets (Renò and Rizza 2003, Grobys 2021, 2023, 2024, Fathi *et al.* 2024, 2025).

The presence of power-law behavior in financial returns and volatility carries profound implications for risk measurement frameworks such as Value-at-Risk (VaR) and stress testing. Traditional models, including those based on log-normal

or Gaussian assumptions, underestimate the probability and magnitude of extreme events due to their thin-tailed nature. Empirical evidence from Bouchaud (2001) and Gencay *et al*. (2003) demonstrates that financial markets often exhibit heavy-tailed distributions, particularly in the tails of return and variance distributions. This tail heaviness implies that large losses occur far more frequently than predicted by normal models, which can severely compromise the reliability of VaR estimates. Gencay and Selçuk (2004) show that applying Extreme Value Theory (EVT) and the Generalized Pareto Distribution (GPD) provides a more robust framework for modeling the tails of financial distributions. These approaches allow for better estimation of risk at high quantiles and facilitate more realistic stress testing scenarios. Moreover, power-law scaling properties are indicative of long memory and clustered volatility, which further challenge the assumptions of independence and stationarity in standard risk models. As such, incorporating power-law features into risk modeling frameworks is essential for accurately capturing the true nature of financial market risk, particularly in the context of systemic shocks.

Fifth, to assess whether the range-based variances of the five assets share a common component, we implemented a joint test hypothesizing a common power-law exponent. Recognizing the limitations of standard errors derived under independence assumptions of Clauset *et al*. (2009), we employed block bootstrap techniques following Grobys (2024) to compute robust standard errors. In line with the study of Grobys (2024), the bootstrap procedure showed that robust standard errors were significantly larger than the ones derived in the study of Clauset *et al*. (2009). Using the robust standard error, we carried out conjoint tests, iteratively computed across a range of hypothesized exponents (q). Our findings suggest strong evidence for a shared power-law exponent governing the range-based variances of these assets. Specifically, the null hypothesis of a shared exponent could not be rejected for the range $2.5 < q < 3.1$, suggesting a unified component influencing variance risks across the S&P 500, gold, crude oil, USD/GBP exchange rate, and Bitcoin. This finding aligns with theoretical expectations of power-law behavior and underscores the interconnected nature of these assets within a broader market framework.

Next, to address the potential impact of autocorrelation on tail index estimates, we applied autoregressive models of order p to the Garman-Klass range-based variances of each asset. The results indicate significant autocorrelation patterns in the processes generating range-based variances. The descriptive statistics indicates that even after accounting for higher-order autocorrelation, the innovation processes of the annualized daily range-based variances for all the assets display extremely heavy tails.

We then used the absolute values of the estimated innovation processes for each asset to estimate the power-law exponents. Notably, the estimated exponents closely align with those observed for the original data, providing strong evidence for the presence of genuine power-law behavior. Next, we performed Clauset *et al*.'s (2009) goodness-of-fit test, for the absolute values of the innovation processes of each asset. This procedure helps to evaluate whether the observed power-law behavior persists after accounting for autocorrelation, providing a robust check on the validity of the power-law estimates for our range-based variances. The goodness-of-fit test indeed indicated that the power-law null hypothesis holds for all innovation processes except gold, providing additional confirmation of the consistent power-law characteristics in the range-based variances. Additionally, the goodness-of-fit test results, with zero p-values across all cases, reject log-normal and exponential distributions, favoring the power-law distribution instead. These results support recent findings of Fathi *et al*. (2025) documenting that realized variances of Fama-French factors are subject to power-law behavior even after controlling for autocorrelations.

Finally, to validate the robustness of the power-law exponents, we performed sample-split tests by dividing the residual series for each asset into two equal-length subsamples. The consistency of tail exponents across both subsamples demonstrated that power-law behavior persists independently of sample-specific effects. Additionally, the goodness-of-fit test supports the power-law model in both subsamples for most assets, further enhancing the reliability of our findings.

A reader could be concerned about the robustness of tail estimation due to limited data in the extreme upper tail. However, the tail sample sizes in our study consistently exceed the empirical threshold suggested by Clauset *et al*. (2009), who argue that reliable estimation of power-law parameters typically requires a minimum of approximately 50 tail observations. For instance, in the case of Bitcoin, the number of observations in the power-law tail constitutes at least 3.41% of the total sample size of 3603—yielding more than 120 observations. Thus, the sample size in the tail comfortably meets the minimum requirement for credible tail exponent estimation using maximum likelihood techniques.

Another potential concern relates to whether our range-based volatility estimates sufficiently address market micro structure noise. In this regard, we note that our analysis is based on daily OHLC data and employs the Garman–Klass estimator, which has been identified as one of the most efficient and least noisy among range-based volatility measures (Molnár 2012). Importantly, microstructure noise is particularly problematic in high-frequency data, whereas range-based estimators constructed from daily prices are considerably less susceptible to such distortions. Thus, while no volatility estimator is entirely free from noise, the estimator we use represents a methodologically sound and practically robust choice that mitigates the typical concerns associated with market microstructure effects.

Finally, the potential influence of structural market features, such as liquidity and trading hours, on tail behavior merits attention. However, our results suggest that markets with widely varying characteristics—such as Bitcoin (traded 24/7 with relatively lower trading volume) and the S&P 500 (traded during regular hours with much higher volume)—exhibit statistically indistinguishable power-law exponents in range-based volatility. This finding is robust across asset classes, including commodities and FX. Additionally, trading volume in the S&P 500 has shown a clear upward trend over the sample period (see Figure A.2). As such, any attempt to split the sample into low- and high-liquidity regimes would effectively mirror our existing time-based splits (as in table 11), which already reveal stable tail exponent estimates

across subperiods. These observations suggest that structural differences such as trading hours and volume have limited impact on the tail behavior we observe.

7. Conclusion

This study presents novel evidence on the variance risk of financial markets by modeling range-based variances using the Garman-Klass variance estimator and analyzing their properties under hypothesized power-law distributions. Focusing on five key financial assets—S&P 500, gold, crude oil, the USD/GBP exchange rate, and Bitcoin—our findings reveal that range-based variances exhibit heavy tails, aligning closely with power-law distributions. These findings carry profound implications for financial modeling and risk management.

First, the estimated tail exponents for range-based variances consistently fall within the range $2 < \alpha < 3$, indicating characteristics of infinite variance-of-variance. As noted by Mandelbrot (1967), such variance can be considered 'so large that it may in practice be assumed infinite.' This challenges the foundational assumption of finite variances in traditional econometric models and highlights the limitations of statistical tools like *t*-tests, which are derived from the ordinary least squares (OLS) regression framework. The evidence of sample specificity, as highlighted by Fama's (1963) reflections on Mandelbrot's infinite variance hypothesis, underscores the inadequacy of moment-dependent statistical methods for handling the heavy-tailed distributions common in financial markets.

Second, goodness-of-fit tests consistently reject alternative distributions, including the widely used log-normal model, in favor of power-law distributions, emphasizing the inadequacy of conventional models in capturing the tail risks that define financial markets. Additionally, evidence derived from blocks bootstrap underscores the importance of accounting for dependency structures when estimating the standard deviations of tail exponents. Our joint test for a common power-law exponent reveals statistical invariance in power-law behavior for the range $(2.5 < q < 3.1)$ across all five assets, suggesting a universal factor governing variance risk. Autoregressive modeling of range-based variances further corroborates the persistence of power-law behavior even after adjusting for autocorrelation. Lastly, sample-split tests confirm the stability of power-law exponents over different time periods, reinforcing the robustness of the findings.

The empirical finding that a common power-law exponent governs realized variances across otherwise unrelated asset markets finds theoretical support in Gabaix (2016). In his comprehensive review, Gabaix (2016) outlines mechanisms—such as proportional random growth processes and universality—that explain why power-law distributions emerge across diverse economic systems. These mechanisms suggest that similar tail behaviors can arise in structurally different markets due to shared statistical properties, rather than asset-specific dynamics. Gabaix (2016) also emphasizes the concept of universality, where systems governed by different microfoundations may exhibit identical scaling laws at the macro level. This perspective aligns with the observation in our study that assets as diverse as crude oil, gold, and Bitcoin exhibit similar tail exponents. Furthermore, his discussion of aggregation effects and macroeconomic granularity provides a theoretical foundation for interpreting power-law regularities as emergent from complex but generalizable dynamics. Hence, our finding is consistent with and theoretically grounded in the framework developed by Gabaix (2016).

Furthermore, our results underscore the necessity of adopting alternative statistical frameworks for analyzing financial data and managing extreme risks. Effectively infinite variances and the rejection of log-normality suggests that reliance on traditional tools can significantly underestimate tail risks, particularly in environments characterized by extreme market events. The heavy-tailed nature of range-based variances has substantial implications for portfolio optimization, asset pricing, and risk management, necessitating frameworks that effectively account for the inherent extremity in financial data. By advancing a robust foundation for understanding variance risks across diverse asset classes, this study contributes to a more accurate and unified approach to financial modeling, particularly under infinite variance phenomena.

Despite its significant contributions, this study has limitations that merit consideration. For example, the analysis is limited to five financial assets—S&P 500, gold, crude oil, USD/GBP exchange rate, and Bitcoin—which constrains the generalizability of the findings to other asset classes or broader market conditions. Expanding the analysis to a wider range of assets and markets would enhance the applicability and breadth of the results.

8. Open Scholarship

This article has earned the Center for Open Science badge for Open Data. The data are openly accessible at (https://zenodo.org/records/15707465 (DOI 10.5281/zenodo.15707464).

Acknowledgments

The authors are grateful to the three anonymous referees for their valuable comments.

Data availability statement

The dataset has been deposited in Zenodo and is publicly available at: https://zenodo.org/records/15707465 (DOI 10.5281/zenodo.15707464).

Disclosure statement

No potential conflict of interest was reported by the author(s).

References

Andersen, T. G., Bollerslev, T., Diebold, F. X. and Ebens, H., The distribution of realized stock return volatility. *J. financ. econ.*, 2001, **61**(1), 43–76.

Andersen, T. G., Bollerslev, T., Diebold, F. X. and Labys, P., Modeling and forecasting realized volatility. *Econometrica*, 2001a, **71**(2), 579–625.

Andersen, T. G., Bollerslev, T., Diebold, F. X. and Labys, P., The distribution of realized exchange rate volatility. *J. Am. Stat. Assoc.*, 2001b, **96**(453), 42–55.

Andersen, T. G., Bollerslev, T. and Meddahi, N., Analytical evaluation of volatility forecasts. *Int. Econ. Rev. (Philadelphia)*, 2004, **45**, 1079–1110.

Andersen, T. G., Bollerslev, T., Diebold, F. X. and Labys, P., Modeling and forecasting realized volatility. *Econometrica*, 2003, **71**, 579–625.

Barndorff-Nielsen, O. E. and Shephard, N., Econometric analysis of realized volatility and its use in estimating stochastic volatility models. *J Roy Statist Soc Series B: Stat Method*, 2002, **64**(2), 253–280.

Barndorff-Nielsen, O. E. and Shephard, N., Power and bipower variation with stochastic volatility and jumps. *Journal of Financial Econometrics*, 2004, **2**(1), 1–37.

Bedowska-Sojka, B. and Kliber, A., Information content of liquidity and volatility measures. *Physica A*, 2021, **563**, 125436.

Bollerslev, T., Generalized autoregressive conditional heteroskedasticity. *J. Econom.*, 1986, **31**(3), 307–327.

Bollerslev, T. and Todorov, V., Tails, fears, and risk premia. *J Fin*, 2011, **66**(6), 2165–2211.

Bouchaud, J. P., Power laws in economics and finance: Some ideas fromphysics. *Quant Fin*, 2001, **1**(1), 105.

Bubák, V., Kočenda, E. and Žikeš, F., Volatility transmission in emerging European foreign exchange markets. *J Bank Fin*, 2011, **35**, 2829–2841.

Calvet, L. E. and Fisher, A. J., How to forecast long-run volatility: Regime switching and the estimation of multifractal processes. *J Fin Econom*, 2004, **2**(1), 49–83.

Chan, L. K. C., Karceski, J. and Lakonishok, J., New paradigm or same old hype in equity investing? *Fin Anal J*, 2000, **56**, 23–36.

Chou, R. Y., Chou, H. C. and Liu, N., Range volatility models and their application in finance. In *Handbook of Quantitative Finance and Risk Management*, edited by C. F. Lee, A. C. Lee and J. Lee, pp. 1273–1281, 2010 (Springer: Boston, MA).

Clauset, A., Shalizi, C. R. and Newman, M. E. J., Power-law distributions in empirical data. *SIAM Rev.*, 2009, **51**, 661–703.

Corsi, F., A simple approximate long-memory model of realized volatility. *J Fin Econ*, 2009, **7**(2), 174–196.

Creal, D., Koopman, S. J. and Lucas, A., Generalized autoregressive score models with applications. *J. Appl. Econ.*, 2013, **28**(5), 777–795.

Engle, R. F. and Rangel, J. G., The spline-GARCH model for low-frequency volatility and its global macroeconomic causes. *Rev Fin Stud*, 2008, **21**(3), 1187–1222.

Fama, E. F., Mandelbrot and the stable Paretian hypothesis. *J. Bus.*, 1963, **36**(4), 420–429.

Fama, E. F. and French, K. R., The cross-section of expected stock returns. *J Fin*, 1992, **47**, 427–465.

Fama, E. F. and French, K. R., Common factors in the returns on stocks and bonds. *J. Financ. Econ.*, 1993, **33**, 3–56.

Fama, E. F. and French, K. R., International tests of a five-factor asset pricing model. *J. Financ. Econ.*, 2017, **123**(3), 441–463.

Fama, E. F. and French, K. R., Choosing factors. *J. Financ. Econ.*, 2018, **128**(2), 234–252.

Fama, E. F. and French, K. R. The value premium. *Fama-Miller Working Paper No. 20–01*, 2020.

Fama, E. F. and French, K. R., A five-factor asset pricing model. *J. Financ. Econ.*, 2015, **116**, 1–22.

Fathi, M., Grobys, K. and Äijö, J., A common component of Fama and French factor variances. *Nor Am J Econ Fin*, 2025, **75**, 102292.

Fathi, M., Grobys, K. and Kolari, J. W., On the realized risk of foreign exchange rates: A fractal perspective. *J Risk Fin Manag*, 2024, **17**(2), 79.

Feng, G., Giglio, S. and Xiu, D., Taming the factor zoo: A test of new factors. *J Fin*, 2020, **75**, 1327–1370.

Gabaix, X., Power laws in economics: An introduction. *J Econ Pers*, 2016, **30**(1), 185–206.

Garman, M. B. and Klass, M. J., On the estimation of security price volatilities from historical data. *J. Bus.*, 1980, **53**(1), 67–78.

Gencay, R. and Selçuk, F., Extreme value theory and value-at-risk: relative performance in emerging markets. *Int. J. Forecast.*, 2004, **20**(2), 287–303.

Gençay, R., Selçuk, F. and Ulugülyağci, A., High volatility, thick tails and extreme value theory in value-at-risk estimation. *Ins: Mathe Econ*, 2003, **33**(2), 337–356.

Godfrey, L., *Bootstrap Tests for Regression Models*, 2009 (Palgrave Macmillan: Basingstoke).

Grigaityte, K. and Atwal, G. *Bayesian inference of power law distributions. bioRxiv*, 2019.

Grobys, K., What do we know about the second moment of financial markets? *Int Rev Fin Anal*, 2021, **78**, 101891.

Grobys, K., Correlation versus co-fractality: Evidence from foreign exchange rate variances. *Int Rev Fin Analy*, 2023, **86**, 102531.

Grobys, K., A universal exponent governing foreign exchange rate risks. *Int Rev Fin Anal*, 2024, **95**(Part B), 103422.

Hansen, B. E., Autoregressive conditional density estimation. *Int. Econ. Rev. (Philadelphia)*, 1994, **35**(3), 705–730.

Harvey, A. and Palumbo, D. Score-driven models for realized volatility. *Cambridge Working Papers in Economics*: 1950, 2019.

Hirshleifer, D., Investor psychology and asset pricing. *J Fin*, 2001, **56**, 1533–1597.

Lintner, J., The valuation of risk assets and the selection of risky investments in stock portfolios and capital budgets. *Rev. Econ. Stat.*, 1965, **47**(1), 13–37.

Lux, T. and Alfarano, S., Financial power-laws: Empirical evidence, models, and mechanisms. *Chaos, Solitons Fractals*, 2016, **88**, 3–18.

Mandelbrot, B., The variation of certain speculative prices. *J. Bus.*, 1963, **36**, 394–419.

Mandelbrot, B., The variation of some other speculative prices. *J. Bus.*, 1967, **40**, 393–413.

Mandelbrot, B. B., *The (Mis)behavior of Markets: A Fractal View of Risk, Ruin, and Reward*., 2008 (Profile Books).

Mandelbrot, B. B., Fisher, A. and Calvet, L. A multifractal model of asset returns. *Cowles Foundation Discussion Paper No.* 1164, 1997.

Markowitz, H., Portfolio selection. *J. Finance.*, 1952, **7**(1), 77–91.

Molnár, P., Properties of range-based volatility estimators. *Int Rev Fin Anal*, 2012, **23**, 20–29.

Mossin, J., Equilibrium in a capital asset market. *Econometrica*, 1966, **35**, 768–783.

Müller, U. A., Dacorogna, M. M., Davé, R. D., Olsen, R. B., Pictet, O. V. and Von Weizsäcker, J. E., Volatilities of different time resolutions—analyzing the dynamics of market components. *J Emp Fin*, 1997, **4**(2-3), 213–239.

Nelson, D. B., Conditional heteroskedasticity in asset returns: A new approach. *Econometrica*, 1991, **59**(2), 347–370.

Parkinson, M., The extreme value method for estimating the variance of the rate of return. *J. Bus.*, 1980, **53**(1), 61–65.

Renò, R. and Rizza, R., Is volatility lognormal? Evidence from Italian futures. *Physica A*, 2003, **322**, 620–628.

Schwert, G. W., Anomalies and market efficiency. In *Handbook of the Economics of Finance*, edited by G. M. Constantinides, M. Harris and R. M. Stulz, Vol. 1, pp. 939–974, 2003 (Elsevier Science B.V.: Amsterdam).

Shu, J. and Zhang, J. E., Testing range estimators of historical volatility. *J Fut Mark*, 2006, **26**, 297–313.

Sharpe, W. F., Capital Asset Prices: A Theory of Market Equilibrium under Conditions of Risk. *J. Finance.*, 1964, **19**(3), 425–442.

Shephard, N., *Stochastic Volatility: Selected Readings*, 2005 (Oxford University Press: Oxford).

Taleb, N. N., *Statistical Consequences of Fat Tails: Real World Preasymptotics, Epistemology, and Applications. Papers and Commentary*, 2020 (STEM Academic Press: New York).

Taylor, S. J., Financial returns modelled by the product of two stochastic processes—A study of daily sugar prices 1961–79. In *Time Series Analysis: Theory and Practice*, edited by O. D. Anderson, Vol. 1, pp. 203–226, 1982 (North-Holland).

Treynor, J. L. Toward a theory of market value of risky assets. Available at SSRN: https://ssrn.com/abstract=628187, 1962.

Virkar, Y. S. and Clauset, A., Power-law distributions and binned empirical data. *Ann. Appl. Stat.*, 2014, **8**(1), 89–119.

White, E., Enquist, B. and Green, J. L., On estimating the exponent of power-law frequency distributions. *Ecology*, 2008, **89**, 905–912.

Zakoian, J. M., Threshold heteroskedastic models. *J Econ Dyn Cont*, 1994, **18**(5), 931–955.

Zhang, L., Mykland, P. A. and Aït-Sahalia, Y., A tale of two time scales: Determining integrated volatility with noisy high-frequency data. *J. Am. Stat. Assoc.*, 2005, **100**(472), 1394–1411.

Appendix

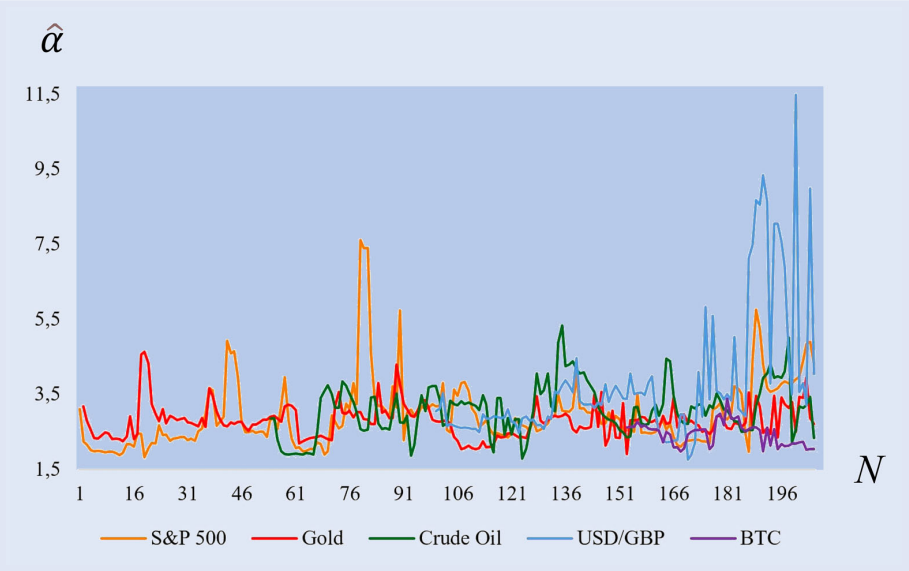

Figure A1. Time-varying estimated power-law exponents based on annualized daily range-based variances for various assets. Note: This figure plots the estimated power-law exponents derived from rolling window estimations based on annualized daily range-based variances for various assets, including the S&*P* 500, gold, crude oil, the USD/GBP exchange rate, and Bitcoin (BTC). The annualized daily range-based variances for each asset are calculated using the Garman-Klass methodology. The power-law exponents are estimated using the maximum likelihood estimation approach.

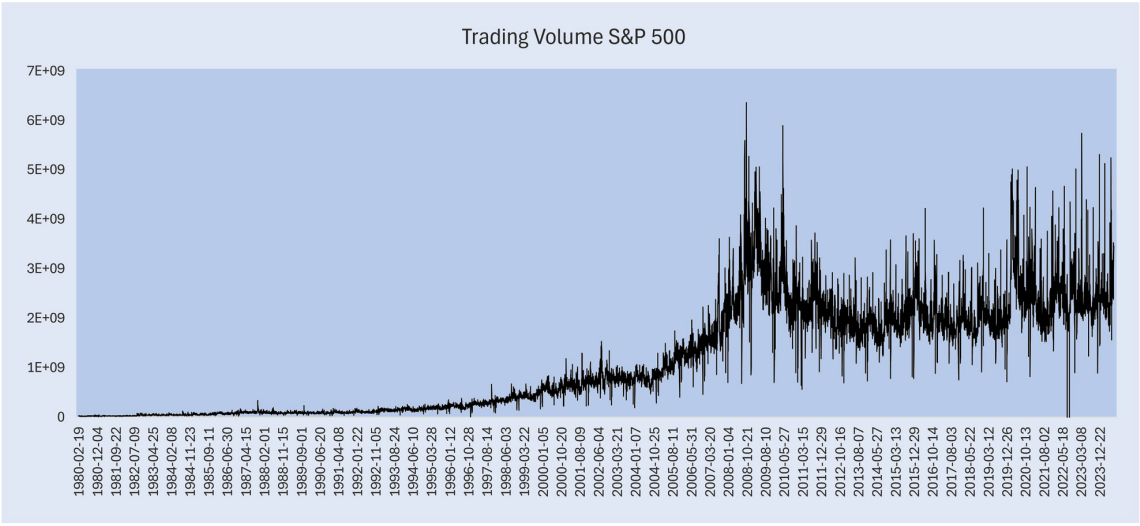

Figure A2. Trading volume of the S&*P* 500 over the sample period. Note: The figure displays the daily dollar trading volume over the sample period from February 19, 1980 to July 31, 2024.

Table A1. Descriptive statistics for weekly realized variance.

Statistic	S&*P* 500	Gold	Crude oil	USD/GBP	BTC
Mean	6.49	6.54	34.81	1.67	74.28
Median	2.99	3.35	15.36	1.09	35.46
Maximum	794.29	265.27	8330.27	78.36	4029.48
Minimum	0.02	0.02	0.14	0.01	0.00
Std. Dev.	22.28	12.33	220.41	3.01	189.70
Skewness	23.90	9.25	32.64	15.04	14.63
Kurtosis	758.30	142.91	1186.98	345.92	277.88
Observations	2242	2239	1711	1253	721
Period (MM/DD/YYYY)	02/19/1980	02/07/1980	01/02/1991	07/25/2000	09/17/2014
	07/31/2024	07/31/2024	07/31/2024	07/31/2024	07/31/2024

Note: This table presents the descriptive statistics of the weekly realized variance for various assets, including the S&*P* 500, gold, crude oil, USD/GBP exchange rate, and Bitcoin (BTC). The weekly realized variances for each asset calculated using the sum of the squared returns over five business days.

Table A2. Estimation of power-law exponents for weekly realized variance.

Metric	S&*P* 500	Gold	Crude oil	USD/GBP	BTC
$\hat{\alpha}$	2.51	2.24	2.43	2.84	2.93
$\hat{x}_{MIN}$	10.51	4.13	23.42	1.96	162.39
D	0.02	0.04	0.03	0.03	0.05
N_{PL}	12.53%	43.19%	33.26%	24.18%	10.40%
Std. Dev	0.09	0.04	0.06	0.11	0.22
Observations	2242	2239	1711	1253	721
Period (MM/DD/YYYY)	02/19/1980	02/07/1980	01/02/1991	07/25/2000	09/17/2014
	07/31/2024	07/31/2024	07/31/2024	07/31/2024	07/31/2024

Note: This table presents the estimation results for power-law exponents applied to the weekly realized variances of the S&*P* 500, gold, crude oil, the USD/GBP exchange rate, and Bitcoin. The weekly realized variances for each asset calculated using the sum of the squared returns over five business days. The parameter $\hat{\alpha}$ signifies the estimated tail exponent, while $\hat{\sigma}$ denotes the estimated standard deviation. The lower threshold $\hat{x}_{MIN}$ is identified through the optimized Kolmogorov-Smirnov distance (D), following the procedure described by Clauset *et al.* (2009). The reported $\hat{x}_{MIN}$ corresponds to the optimal distance D. Additionally, the fraction of observations obtained by the power-law process, denoted as N_{PL} is provided for each asset.

Table A3. Goodness-of-fit tests for the weekly realized variance.

Asset market	Power-law	Log-normal	Exponential
S&*P* 500	0.99	0.58	0.00
Gold	0.04	0.65	0.00
Crude oil	0.34	0.11	0.00
USD/GBP	0.67	0.31	0.00
BTC	0.73	0.31	0.00

Note: This table presents the results (viz., p-values) of the goodness-of-fit tests for weekly realized variance, conducted using the Kolmogorov-Smirnov (KS) method to assess whether the empirical data and the synthetic data—generated from a power-law distribution specified by x_{min} and α —originate from the same underlying distribution (column 2). Additionally, the table includes the results of goodness-of-fit tests for log-normal and exponential distributions in columns 3 and 4, respectively. The null hypothesis for these tests states that the empirical data and the synthetic data generated from the specified distributions (log-normal and exponential) are consistent, suggesting that each distribution provides an acceptable fit to the observed data. The weekly realized variances for each asset market are calculated using the sum of the squared returns over five business days.

Table A4. Descriptive statistics of estimated power-law exponents derived from rolling window estimations.

Statistic	S&*P* 500	Gold	Crude Oil	USD/GBP	BTC
Minimum	1.83	1.91	1.79	1.76	1.98
5% Qnt.	1.99	2.17	1.92	2.28	2.04
10% Qnt.	2.10	2.31	2.22	2.54	2.05
Median	2.73	2.77	3.11	3.24	2.54
90% Qnt.	3.88	3.34	4.04	7.01	2.84
95% Qnt.	4.59	3.55	4.27	8.41	2.91
Maximum	7.61	4.63	5.33	11.48	3.38
Mean	2.96	2.79	3.09	3.78	2.45
Std. Dev	0.92	0.44	0.68	1.82	0.31
Excess Kurtosis	7.48	2.67	0.32	4.19	− 0.06
Skewness	2.26	1.15	0.39	2.17	0.35
N	205	204	152	106	53

Note: This table reports descriptive statistics for estimated power-law exponents derived from rolling window estimations based on annualized daily range-based variances for various assets, including the S&*P* 500, gold, crude oil, the USD/GBP exchange rate, and Bitcoin (BTC). The annualized daily range-based variances for each asset are calculated using the Garman-Klass methodology. The power-law exponents are estimated using the maximum likelihood estimation approach. Quintiles (Qnt) are reported from the 5th to the 95th percentile.

Table A5. Important subfamilies of the GB2 distribution.

Subfamily	Condition
Generalized Gamma	$q \to \infty$
Beta Prime	$a = 1$
F-distribution	$a = 1, b = 1, p = \frac{d_1}{2}, q = \frac{d_2}{2}$
Singh–Maddala	$p = 1$
Dagum distribution	$q = 1$
Pareto type IV	$p = 1, b = 1$

Note: The Generalized Beta of the Second Kind (GB2) is a four-parameter family of continuous probability distributions that encompasses a wide variety of shapes. The GB2 distribution is defined as:

$$f(x; a; b; p; q) = \frac{a(x/b)^{ap-1}}{bB(p,q)[1 + (x/b)^a]^{p+q}},$$

where $B(p, q)$ is the Beta function, $a > 0$ is the tail parameter controlling the tail heaviness, $b > 0$ is a scale parameter, $p > 0$, $q > 0$ are shape parameters that also govern skewness. This table reports important subfamilies of the GB2 distribution.

Table A6. Estimated parameters of the GB2 distribution.

Range-based variance	a	b	p	q
S&*P* 500	125.30	4.44	67.73	0.016
Gold	114.95	6.36	116.26	0.015
Crude oil	73.95	22.47	113.15	0.022
USD/GBP	108.31	1.05	105.28	0.016
BTC	223.86	243.08	0.0255	0.0094

Note: The Generalized Beta of the Second Kind (GB2) is a four-parameter family of continuous probability distributions that encompasses a wide variety of shapes. The GB2 distribution is defined as:

$$f(x; a; b; p; q) = \frac{a(x/b)^{ap-1}}{bB(p,q)[1 + (x/b)^a]^{p+q}},$$

where $B(p, q)$ is the Beta function, $a > 0$ is the tail parameter controlling the tail heaviness, $b > 0$ is a scale parameter, $p > 0$, $q > 0$ are shape parameters that influence the distribution's form, including its skewness. This table reports the parameter estimates for GB2 distributions fitted to the data ranges defined by the power-law cutoffs listed in table 3.

Momentum Bubbles

A New Perspective Derived from Log-Periodicity

This Draft: August 6, 2025

Masoumeh Fathi[1a] Klaus Grobys[2a] Janne Äijö[3a]

This study examines the predictability of momentum bubbles by calibrating the Log-Periodic Power-Law Singularity (LPPLS) model using data from November 3, 1926, to November 30, 2023. We analyze LPPLS signatures preceding major momentum downturns and assess how the model's predictive accuracy evolves as bubbles approach their peak. Our findings indicate that LPPLS signatures provide early warning signals before significant momentum bubble formations, supporting the model's effectiveness in identifying early signs of bubble formation within a zero-cost trading strategy. By evaluating the accuracy of predictions derived from LPPLS calibrations 3, 6, 9, and 12 months before the bubble peak, we provide empirical evidence that predictive accuracy does not improve linearly as the peak approaches but exhibits partial improvement.

Keywords: financial bubble, log-periodic power law, momentum bubbles, finite-time singularity

JEL Classification: G01, G11, G12, G14, G17, H12.

[1] M. Fathi (corresponding author)

Finance Research Group, School of Accounting and Finance, University of Vaasa, Wolffintie 34, 65200 Vaasa, Finland; E-mail: masoumeh.fathi@uwasa.fi.

[2] K. Grobys

Finance Research Group, School of Accounting and Finance, University of Vaasa, Wolffintie 34, 65200 Vaasa, Finland; Innovation and Entrepreneurship (InnoLab), University of Vaasa, Wolffintie 34, 65200 Vaasa, Finland;

E-mail: klaus.grobys@uwasa.fi.

[3]J. Äijö

Finance Research Group, School of Accounting and Finance, University of Vaasa, Wolffintie 34, 65200 Vaasa, Finland;

E-mail: janne.aijo@uwasa.fi.

1. Introduction

The momentum strategy, first documented in Jegadeesh and Titman's (1993) seminal paper, is a well-established trading approach recognized for its ability to capture market trends and generate risk-adjusted average returns. Although the momentum strategy remains replicable in extended samples (Hou, Xue, & Zhang, 2020), it is prone to recurring bubbles and crashes (Daniel & Moskowitz, 2016). For instance, Daniel and Moskowitz (2016) observe that between March and May 2009, the past-loser decile increased by 163%, while past winners gained only 8%, resulting in an implied loss of 155%. A key behavioral factor contributing to bubbles in momentum investing is herding behavior, wherein individuals imitate others rather than relying on their own information. This collective behavior drives prices beyond their intrinsic values, leading to bubble formation (Bikhchandani, Hirshleifer, & Welch, 1992).

The very nature of momentum investing, which relies on past performance, leads investors to make similar trading decisions and, in turn, can lead to overconfidence and extrapolation of trends, which are behavioral biases that can inflate asset prices beyond their fundamental values (Daniel, Hirshleifer, & Subrahmanyam, 1998). Hong and Stein (1999) show that momentum trading can contribute to market inefficiencies, as it causes price trends to persist longer than they should, creating conditions suitable for bubble formation. This process of self-reinforcing price increases due to positive feedback loops, described as "positive feedback trading" by De Long et al. (1990), explains why momentum bubbles can rapidly escalate and collapse. This unique characteristic suggests that herding behavior—where investors mimic others' trades rather than relying on their own information—plays a significant role in the formation and bursting of momentum bubbles.

Given the potential role of herding behavior in the formation and collapse of momentum bubbles, there is a need for empirical validation and quantification of this relationship. One promising tool for this purpose is the LPPLS model, introduced by Sornette et al. (1996), which has been widely used

to diagnose critical market phenomena, such as bubbles and crashes. Although the LPPLS model shows potential for forecasting market bubbles, its use in long-short investment vehicles, such as the momentum strategy, remains underexplored in financial research. The existing literature mainly examines the LPPLS model's ability to identify bubble formations and crashes building up in broad market indices.

This study aims to fulfill two main objectives: (a) identifying momentum bubble formations and (b) predicting the timing of regime switches—potential stochastic turning points—using the LPPLS model for the period from November 3, 1926, to November 30, 2023. We identify sixteen bubbles in our sample by adapting the epsilon draw-down Method, following Johansen & Sornette (1998, 2001, 2010), Filimonov & Sornette (2015), and Gerlach et al. (2019), as shown in Table 1. Next, to calibrate the LPPLS model, we focus on bubbles with at least a four-year gap between them due to the nonlinear nature of the LPPLS model. This approach yields five momentum bubbles with peaks on December 21, 1929; March 10, 2000; July 14, 2008; January 20, 2016; and April 16, 2020, as candidates for potential prediction. We then examine the LPPLS model's prediction accuracy by analyzing calibration windows ending 3, 6, 9, and 12 months before each bubble peak to determine whether accuracy improves as the peak approaches.

This study makes a novel contribution to the existing literature. It is the first study to explore the predictability of momentum bubbles using the LPPLS model, which is uniquely suited for identifying self-reinforcing price escalations and log-periodic oscillations that characterize bubble dynamics and signal potential regime changes in the market (Sornette & Johansen, 2001). Previous literature has applied the LPPLS model to predict stock market bubbles and crashes (e.g., Johansen et al., 1999; Vandewalle et al., 1999; Chang & Feigenbaum, 2008), cryptocurrency bubbles (e.g., Wheatley et al., 2018), or sovereign debt crises (Geraskin & Fantazzini, 2013). Moreover, this study is the first to explore whether the LPPLS model can predict bubbles within an investment vehicle derived from a long-short strategy. It is important to note that Daniel and Moskowitz (2016, p. 222) highlight that

the payoffs of the short leg of the momentum portfolio are exposed to dynamic betas, which "implies that momentum strategies in bear markets behave like written call options on the market; that is, when the market falls, they gain a little, but when the market rises, they lose much." These option-like features—especially those produced by the momentum strategy—are distinct from broad market indices, making the momentum strategy extremely risky during periods of stress or market reversals.

In addition, our study investigates the effect of the calibration window on the precision of the LPPLS model. Using equity market data for LPPLS-based stock market crash predictions, Sornette (2017, p. 330) argues that "a year before the crash, the data is not sufficient to give any conclusive results at all. This point corresponds to the end of the fourth oscillation. Approximately a year before the crash, the fit begins to lock in on the date of the crash with increasing precision. In fact, in four of the last five time intervals, a fit exists with a t_c, which differs from the true date of the crash by only a few weeks." Furthermore, the study by Brée et al. (2013) casts doubt on the reliability of the LPPLS model by documenting that the estimated critical time fluctuates significantly when the time-series input data change and refer to t_c as a "sloppy parameter." This raises the question: How reliable are LPPLS model predictions for the critical time? Adding to this literature, our study is the first to explore the reliability of LPPLS predictions**,** conditional on the calibration window, by adopting Grobys' (2023) recently proposed three-stage estimation approach.

Furthermore, another novel contribution of this study is that it explores the predictability of momentum bubbles and their potential risk management. While Barroso and Santa-Clara (2015) argue that the large returns from momentum strategies may not compensate for the risks of sudden crashes, we argue that momentum risk is manageable if bubbles and subsequent crashes are predictable. If investors expect a bubble to burst, they can close their long-short positions and reopen them after the bubble bursts and the market stabilizes. Therefore, this is the first study to examine whether momentum bubbles are predictable and, thus, manageable.

Our study provides a behavioral explanation for the occurrence of momentum bubbles. As will be discussed in Section 2.2, the LPPLS model is derived from the theoretical framework of imitation and herding behavior among traders. Consequently, we argue that evidence of significant LPPLS signatures in the log-compounded return of the momentum factor implies that momentum bubbles are not random events but rather artifacts of herding behavior. Essentially, individual traders are likely to follow the herd rather than make independent decisions. This behavior can create speculative bubbles, which, when they burst, can lead to crashes. We argue that the optionality effect discovered in Daniel and Moskowitz's (2016) study has a behavioral cause rooted in imitation among traders. Moreover, this study offers practical implications for portfolio management and risk assessment. Integrating the LPPLS model with momentum strategies could enhance decision-making, allowing for timely adjustments to portfolios in anticipation of potential reversals.

The results of our study confirm the presence of LPPLS signatures before each of the five bubbles in our sample—December 21, 1929; March 10, 2000; July 14, 2008; January 20, 2016; and April 16, 2020. Through the application of the LPPLS model and analysis of prediction accuracy at intervals of 3, 6, 9, and 12 months prior to each bubble peak, we demonstrate that these signatures can effectively serve as early indicators of future bubbles. Furthermore, our findings indicate that although the prediction accuracy of the LPPLS model does not consistently increase as the bubble approaches its peak, it shows partial improvement. These results support the previous literature suggesting that the prediction accuracy of the LPPLS model improves as we approach the critical time (Sornette, 2017). In addition, a cross-sectional comparison of all five bubble predictions shows that the LPPLS model tends to offer more accurate predictions when utilizing data from 9 months prior to the critical point compared to data from 3, 6, or 12 months beforehand.

The results of this study align with previous research in several ways. First, the LPPLS model's ability to detect early warnings of momentum bubbles is confirmed, which is consistent with findings for broad market indices, as documented in Sornette (2017). This study further supports the probabilistic

rather than deterministic nature of the LPPLS model, as emphasized by Sornette (2017), and acknowledges the model's sensitivity to parameter selection, a challenge noted by Zhou and Sornette (2006) and Brée et al. (2013). Second, this study confirms previous literature on the predictability of momentum bubbles and crashes (e.g., Daniel & Moskowitz, 2016; Barroso & Santa-Clara, 2015; Chabot et al., 2014), which found them to be predictable, especially when momentum performed well or interest rates were low. Adding to this literature, our study provides novel insights by demonstrating that the LPPLS model can reveal its signatures before the onset of momentum-driven bubble bursts, validating its ability to capture pre-peak dynamics. Lastly, our study's analysis of prediction accuracy using data up to 3, 6, 9, and 12 months prior to the bubbles peak provides a comprehensive view of the model's performance, aligning with prior literature acknowledging the inherent limitations and challenges of forecasting complex dynamic systems (Sornette & Johansen, 2001; Cont, 2001; Sornette, 2003).

The remainder of this paper is organized as follows: Section 2 provides the background of the study, Section 3 details the data and methodology, Section 4 presents the results and findings, Section 5 outlines the robustness tests, Section 6 analyzes these findings, and Section 7 concludes the paper.

2. Background Discussion

2.1. Understanding momentum bubbles through the framework provided by the LPPLS model

Unlike bubbles driven by external factors such as technological innovation or macroeconomic shifts, momentum bubbles emerge from the synchronized behavior of investors who follow similar strategies, often disregarding the intrinsic values of the assets. This behavior is particularly prevalent during periods of high uncertainty, as investors tend to follow the crowd to avoid the perceived risks of acting independently (Avery & Zemsky, 1998). Consequently, herding behavior has been identified as a key driver of these speculative bubbles, amplifying price movements and market volatility (Lux, 1995).

We argue that the LPPLS model is particularly well-suited for analyzing momentum-driven bubbles because it effectively captures the herding behavior that characterizes such speculative periods. This model, rooted in the physics of complex systems, hypothesizes that financial market bubbles and subsequent crashes are preceded by super-exponential growth in asset prices, terminating at a singularity point—referred to as a critical point—at which a regime switch may occur. As markets approach this critical point, price dynamics begin to exhibit power-law behavior, characterized by accelerating growth rates and increasing volatility (Sornette, 2003; 2017). In these regimes, small changes in market conditions can lead to disproportionately large effects on asset prices, making the system highly sensitive to minor disturbances (Sornette & Cauwels, 2014). Thus, once the endogenous system approaches the critical point, not only does the price process follow a power law, but price changes themselves are also governed by a power-law distribution—a finding first documented in Mandelbrot's (1963) seminal study, which identified that price changes of speculative financial assets exhibit power-law behavior. Recognizing the characteristic patterns of price movements close to critical points provides valuable insights into how herding behavior contributes to systemic risk and market vulnerability.

In addition, the LPPLS model benefits from a nonlinear mathematical framework that describes the behavior of financial markets during speculative bubbles and crashes (Sornette, 2003; 2017). Its nonlinearity arises from the combination of a power-law singularity with log-periodic oscillations, capturing the accelerating growth and increasing volatility observed in asset prices as they approach critical points, ultimately leading to regime changes (Sornette et al., 1996; Sornette, 2003). This nonlinear structure makes the LPPLS model particularly suited for modeling market returns during bubble phases. While traditional models often assume linear or constant growth rates that fail to capture the accelerating price increases seen in bubbles, the LPPLS model's power-law component effectively reflects this acceleration (Sornette, 2003). In this framework, log-periodic oscillations account for the increasing volatility and frequency of price fluctuations as the market approaches a

critical point (Sornette et al., 1996). Empirical studies have shown that the LPPLS model can successfully identify and predict the timing of market bubbles and crashes by fitting historical price data to its nonlinear structure (Zhou & Sornette, 2006).

However, despite its appealing features, the LPPLS model has some shortcomings. Specifically, it (a) cannot predict the exact timing of a market bubble but can only estimate its likelihood—including timing and market level—through a specific probability density function (Zhang et al., 2016), (b) is prone to false positives (Feigenbaum & Freund, 1996), and (c) faces estimation difficulties due to numerous local maxima in the likelihood function (Feigenbaum, 2001).[1] Chang and Feigenbaum (2008) examine the limitations of the LPPLS model and find that it failed to provide informative insights into the oscillatory nature of market dynamics leading up to the October 1987 crash.[2] Consequently, theoretical research on the LPPLS model has focused on refining its mathematical framework and improving its predictive accuracy.[3]

2.2. Modeling the interactions between traders

The LPPLS model is built on the dynamics of two distinct market agents: rational investors, who understand fundamental asset values, and noise traders, whose herding behavior can deviate prices from these values. Rational investors continue to invest despite inflated prices, expecting higher

[1] Feigenbaum (2001) asserts that this complexity results in the small-sample properties of this estimation method being poorly understood, making it difficult to effectively assess the statistical significance of a non-zero log-periodic component.

[2] For a more in-depth discussion regarding the capabilities and limitations of log-periodic precursors to financial crashes, refer to the foundational works of Graf and Meister (2003) and the comprehensive review of Geraskin and Fantazzini (2013), among others. These studies provide valuable insight into the statistical, Bayesian, and critical analyses of log-periodicity and its role in financial market predictions.

[3] A number of improvements have been made on the calibration of the LPPLS model, including meta-search heuristics (Zhou & Sornette, 2006) and the reformulation of equations aimed at reducing the number of nonlinear parameters (Filimonov & Sornette, 2013). Furthermore, Filimonov et al. (2016) present a detailed methodological approach to calibrating nonlinear financial models, including the LPPLS, and introduce the modified profile likelihood method, which significantly reduces the number of local extrema by constructing simpler and smoother log-likelihood landscapes. Their method provides a rigorous interval-based estimation for the critical time rather than a single-point estimate, accounting for the effect of other nonlinear parameters. Notably, in a recent study, Grobys (2023) proposes a novel three-stage estimation approach for implementing the LPPLS, which demonstrates superiority over some earlier methods.

returns unless the market bubble bursts. A bubble bursts when imitation among traders increases to a critical point, causing noise traders to shift collectively to selling, as outlined by Johansen et al. (2000). According to Johansen et al. (2000), the hazard of a stock market crash is derived using mean-field theory from statistical mechanics, which describes the imitative process among traders. This hazard rate is computed by the following differential equation:

$$\frac{dh}{dt} = Ch^{\delta}, \tag{1}$$

where C is a positive constant, and δ is the average number of interactions among traders minus 1. The term h^{δ} in the Eq. (1) signifies that the hazard rate fluctuates, either increasing or decreasing, based on the interactions among traders.

By taking the integral from Eq. (1), we derive

$$h(t) = \left(\frac{h_0}{t_c - t}\right)^{\alpha}, \tag{2}$$

where $\alpha \equiv \frac{1}{\delta - 1}$, t_c represents the critical time when the system is expected to change state (i.e., investors converge to the sell-side), and the exponent α, constrained to $0 < \alpha < 1$, is chosen to ensure that the price remains finite as the system approaches this critical time. This imposes the condition $2 < \delta < +\infty$.

In Johansen et al. (2000), the market is modeled at a micro level by linking traders in a network, inspired by the Ising model from statistical mechanics. In this framework, each trader is characterized by two states buying or selling (s_i = 1 or -1), influenced by adjacent traders, global news, and personal tendencies. The state transition of a trader (s_i) is formulated through a Markovian process as

$$s_t = sign\left\{K \sum_{k \in N(i)} s_k + \sigma\varepsilon_i + G\right\}, \tag{3}$$

where the function $sign(x)$ is defined as +1 for $x > 1$ and -1 for $x < 1$. Here, K is a positive constant, while ε_i is independently distributed by the standard normal distribution.

In this framework, the parameter K, referred to as the coupling strength, controls the tendency to imitate, whereas σ governs the tendency toward idiosyncratic behavior. Here, $N(i)$ denotes the set of traders directly linked to the trader i, and G signifies the global influences accessible by all market participants. Consequently, the relationship between K and σ determines the balance between order and disorder, which in turn influences the probability of a crash. Specifically, as K increases, the order in the network increases, and as σ increases, the order in the network decreases.

Suppose that K_c serves as a critical point that separates different market behaviors based on the level of imitation among traders. At the critical point K_c, if $K < K_c$, the level of imitation among traders is low, and the system exhibits low sensitivity to minor changes in global factors. Conversely, when $K > K_c$, nodes can be converted into the same state (selling), potentially causing a bubble to burst or leading to a crash. These dynamics, in which the level of imitation among traders leads to significant market movements, can be statistically represented by a power law. According to Sornette (2017), the emergence of power-law behavior indicates a form of fractality in the underlying return-generating process. In this context, fractality implies a degree of scaling in which events occurring on smaller scales exhibit a similar pattern of behavior to those occurring on larger scales.

This concept, in the context of the LPPLS model, refers to the idea that patterns of trader behavior and market dynamics are self-similar across different scales. Essentially, this means that what happens on a small scale (e.g., the behavior of a small group of traders) is transmitted to or indicative of what happens on a larger scale (e.g., the behavior of the market as a whole). This is a key characteristic of fractal patterns, as introduced by Mandelbrot (1963) in his seminal work on cotton price changes, which ascertains the self-similar nature of market dynamics, where the patterns observed on one scale are replicated on other scales.

The system's susceptibility (χ) measures its sensitivity to minor shifts in global dynamics. Specifically, it measures how the average state of traders adjusts in response to these small external changes. Johansen et al. (2000) express this concept of susceptibility simply as follows:

$$\chi = \frac{d(E[M])}{dG}\Big|_{G=0}, \tag{4}$$

where χ is the system's susceptibility, M is the market defined as the average state of traders ($\bar{s}$), and G is the global influence.

2.3. Derivation of the log-periodic power law

Johansen et al. (2000) introduce an exogenous crash probability using a jump process denoted by j, which switches to 0 before a crash and 1 afterward. They define the cumulative distribution function (CDF) of crash timing as $Q(t)$, with its probability density function (PDF) expressed as $q(t) = dQ/dt$, and the hazard rate as $h(t) = q(t)/[1 - Q(t)]$, indicating the immediate crash probability. Assuming a crash causes a fixed percentage drop in price $\kappa \in (0,1)$, the asset price dynamics pre-crash are modeled by

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

where $\mu(t)$ ensures that the price process satisfies the martingale condition as

$$E_t[dp] = \mu(t)p(t)dt - \kappa p(t)h(t)dt = 0. \tag{6}$$

Thus, ensuring the asset price remains stable on average requires $\mu(t) = \kappa h(t)$. Substituting Eq. (6) into Eq. (5) leads to an ordinary differential equation, whose solution is expressed as Eq. (7) before the crash:

$$log(\frac{p(t)}{p(t_0)}) = \kappa \int_{t_0}^{t} h(t')dt'. \tag{7}$$

Johansen et al. (2000) propose that, in line with rational expectations theory, traders incorporate the risk of a market crash into their price expectations. Consequently, market prices should increase to

offset the higher risk. The higher the crash probability, the faster the price must increase to reward investors for the increased crash risk in the market (Geraskin & Fantazzini, 2013).

Johansen et al. (2000) demonstrate that as K increases and approaches the critical point K_c, the system exhibits high sensitivity to even small global disturbances, and those with similar viewpoints tend to generate large clusters. Consequently, when imitation spreads across extensive distances, the susceptibility of the system (χ) approaches infinity according to a power law as follows:

$$\chi \approx A(K_c - K)^{-\gamma} , \tag{8}$$

where A is a positive constant, K_c is the critical imitation threshold, and $\gamma > 0$ is the critical exponent of susceptibility.

Johansen et al. (2000) solve this local imitation model by generating a network structure—so-called hierarchical diamond lattice—through an iterative process, starting with a simple connection between two traders. The model begins with two linked traders and evolves by replacing each link with a diamond structure and adding two new traders at each iteration. This process, repeated p times, creates a diamond lattice network. After p iterations, the network expands to $\frac{2}{3}(2 + 4^p)$ traders and 4^p links. The connectivity varies, with most traders having only two neighbors, a few traders (the original ones) having 2^p neighbors, and the others falling in between.

The hierarchical diamond lattice has a fractal framework and was initially solved by Derrida et al. (1983). Johansen et al. (2000) provide the following general solution for susceptibility to the diamond lattice model:

$$\chi \approx Re\left[A_0(K_c - K)^{-\gamma} + A_1(K_c - K)^{-\gamma+i\omega} + \cdots\right]$$

$$\approx A_0{}'(K_c - K)^{-\gamma} + A_1{}'(K_c - K)^{-\gamma} cos[\omega\, log(K_c - K) + \psi] + \cdots , \tag{9}$$

where $A_0{}'$, $A_1{}'$, ω, and ψ are real numbers, and Re represents the real part of a complex number.

The power law is adjusted by oscillations in which the frequency increases as the critical time approaches. The term "log-periodic" describes these accelerating oscillations, while the term "log-frequency" refers to the rate of these oscillations with $\omega/2\pi$ log-frequency.

Thus, the hazard rate of a crash is obtained as follows:

$$h(t) \approx B_0(t_c - t)^{-\alpha} + B_1(t_c - t)^{\beta} cos[\omega log(t_c - t) + \psi']. \quad (10)$$

By substituting Eq. (10) to Eq. (7), the dynamics of price before the critical date are computed by

$$log[p(t)] \approx \ log[p_c] - \frac{\kappa}{\beta}\left(B_0(t_c - t)^{\beta} \ + \ B_1(t_c - t)^{\beta} \cos[\omega \log(t_c - t) + \ \varphi]\right). \quad (11)$$

Finally, in accordance with Geraskin and Fantazzini (2013), Eq. (11) can be reformulated into a more appropriate form to analyze financial time series data as

$$log[E(p(t))] \approx \ A + B(t_c - t)^{\beta} - \ \ (1 + C \cos[\omega \log(t_c - t) + \ \varphi]). \quad (12)$$

Eq. (12), referred to as LPPLS, describes the dynamics of asset prices leading up to a critical point. In this equation, A represents the logarithm of the expected asset price at the critical time, B measures the rate of increase in the logarithm of the asset price as it approaches the critical time t_c , and β is a parameter that captures the rate of super-exponential growth. The coefficient C measures the relative scale of the oscillations throughout the growth path, ω denotes the angular frequency of these oscillations during the bubbles, and φ is the phase adjustment parameter.

3. Data and Methodology

3.1. Data

The data used in this study include the daily observations of 10 portfolios formed using the momentum strategy from November 03, 1926, to November 30, 2023, and obtained from the data library of Kenneth R. French. We compute the momentum factor known as winners-minus-losers (WML), which represents the difference between past winners and past losers for each day of the

time series data. Following mainstream literature, we focus on the cumulative past 12 months return to identify winners and losers, updating the strategy daily using prior return decile breakpoints.

We conduct our analysis on compounded returns using an equity momentum strategy that takes a long position in the winner decile and a short position in the loser decile. Compounded returns are particularly suited for this analysis because they capture the long-term growth dynamics that align with the LPPLS model's emphasis on detecting super-exponential patterns leading to bubbles. Further, by smoothing short-term volatility and reducing noise, compounded returns provide a clearer view of price trends, offering a more accurate reflection of investor experience and facilitating better historical comparisons, both of which are critical for bubble detection.

To compute the compounded return of the daily WML, we use the following formula:

$$COMP_t = \frac{100 + WML_t}{100} \cdot COMP_{t-1}, \tag{13}$$

where $COMP_t$ is the compounded return at time t, and WML_t is the winners-minus-losers portfolio at time t. We set $COMP_0$=100. Figure 1 shows the time series of the natural logarithm of the compounded return for the whole period spanning November 3, 1926, to November 30, 2023.

Next, we identify the peaks of potential financial bubbles. The peak detection technique is based on an adaptation of the epsilon draw-down method, proposed by Johansen and Sornette (1998, 2001, 2010), Filimonov and Sornette (2015), and Gerlach et al. (2019). Following this body of literature, the epsilon draw-down process is designed to systematically segment a return trajectory into consecutive periods of draw-up and draw-down phases, which are interpreted as bubble formations and crashes. A draw-up phase is defined as a period where the return generally trends upward while allowing for minor counter-movements, constrained by a pre-specified tolerance level ε. Specifically, a draw-up is considered a sequence of increasing returns that can only be interrupted by decreasing returns, which do not exceed the pre-specified tolerance level ε. Considering this definition, we define the end of the draw-up as the peak time of a bubble. Similarly, a draw-down is defined as a sequence

of decreasing returns that exceeds the pre-specified tolerance level ε. By definition, after the peak time of a bubble, a draw-down begins, leading to the crash or correction regime. Based on our data, we set $\varepsilon = -0.40$.

In this framework, a bubble can be interpreted as a period of sustained and accelerating growth in the time series return that significantly deviates from fundamental values (Kindleberger & Aliber, 2005), culminating in a peak at the time of a bubble, after which it enters a downtrend. This bubble phase is characterized by super-exponential growth and increasing volatility (Sornette, 2003). It often exhibits patterns consistent with the LPPLS model, such as power-law acceleration and log-periodic oscillations (Sornette, 2003).

Using this framework, we identified sixteen bubbles in our sample from November 3, 1926, to November 30, 2023, ranked in Table 1 by the economic magnitude of their fall in logarithmic compounded return, in descending order. For each bubble, the second and third columns list the corresponding peak time and crash end time, respectively. The change in compounded return (Δ_{CR}) reflects the decrease in logarithmic compounded return from the peak to the end of the crash, while Δ_{WML} represents the decrease in WML over the same period. The duration (days) column shows the time span between the peak and the end of the crash.

The bubble that burst on March 6, 2009, was the most severe, with a fall of -1.4 in compounded return over 45 days. The second most significant bubble burst occurred on June 2, 1932, with a substantial fall of -1.18 in compounded return over 77 days. The longest draw-down period belongs to the recent bubble that burst on November 6, 2020, with a -0.9 fall in compounded return over 174 days. Notable short but sharp bubble bursts occurred during the 2008 financial crisis, including November 20, 2008, with a -0.54 fall in compounded return over 31 days, and July 14, 2008, with a -0.44 drop in just 7 days, highlighting extreme short-term volatility. Another notable recent bubble burst on April 16, 2020, with a -0.85 fall in compounded return over just 37 days, highlighting the sharp decline within a short period. The bubble peak on December 21, 1929, coinciding with the aftermath of the Great

Depression, experienced a significant -0.68 fall in compounded return over 98 days. Figure 1 shows the peak dates of the sixteen bubbles listed in Table 1, represented by the red dashed lines.

Next, to calibrate the LPPLS model, we impose an additional criterion, selecting only bubbles with at least four years of data available between two consecutive bubbles. As a result, we identify five bubbles that meet this criterion: December 21, 1929; March 10, 2000; July 14, 2008; January 20, 2016; and April 16, 2020.

Classical theories, such as those proposed by Kindleberger and Aliber (2005) in their analysis of financial bubbles, suggest that bubbles follow a cyclical pattern, typically developing over several years before bursting. This perspective is supported by historical analyses of financial bubbles, from the Dutch Tulip Mania to the dot-com bubble, which generally show a formation period of seven to ten years (Reinhart & Rogoff, 2009). However, the globalization of financial markets and the rapid acceleration of information dissemination have led some scholars to reconsider the traditional timelines of bubble formation. For example, the mathematical modeling of bubble formation by Sornette and Cauwels (2014) and the analysis of market-specific dynamics by Hoberg and Phillips (2010) suggest that bubbles can form and burst over significantly shorter periods. Brunnermeier (2008) argues that the increased efficiency in financial markets, along with the rapid transmission of investor sentiment and market reactions, may shorten the lifecycle of bubbles. This variability is further underscored by research into the effects of high-frequency trading and the rapid rise and fall of assets, such as cryptocurrencies, illustrating that technological advancements and market innovations can drastically accelerate bubble dynamics (Brogaard et al., 2014; Gandal et al., 2018).

In addition, based on our analysis of the calibration windows—3, 6, 9, and 12 months prior to the actual event—the last year of data is allocated for calibration. This means that, for example, in 12 months prior calibration window, one year of data is excluded. For example, if only four years of data exist between two consecutive bubbles, excluding one year for calibration leaves three years of data for analysis. Therefore, our motivation for considering a minimum of four years for bubble formation

is supported by these insights. Given the highly nonlinear nature of the LPPLS model and its numerous parameters to estimate, having sufficient data is crucial for successful model calibration (Sornette, 2017).

3.2. Methodology

We calibrate the LPPLS model using a three-stage estimation approach, following Grobys (2023). This approach allows for a gradual refinement of model parameters, starting from a simple approximation and progressing to a more complex and accurate representation of the underlying dynamics captured by the LPPLS model. This methodology is designed to overcome the challenges of estimating highly nonlinear parameters by simplifying the LPPLS model into a form that can be calibrated using standard software.

The first stage of Grobys' (2023) estimation approach is called the worst-case scenario, which simplifies the LPPLS model using predefined values for some of its parameters. Thus, the first stage includes Sornette's (2017) simple power-law model for financial log-prices, which is defined as follows:

$$ln[p(t)] = A + B(t_c - t)^{\beta}, \tag{14}$$

where $ln[p(t)]$ is the natural logarithm of the price $p(t)$ at time t. Taking the logarithm can help linearize the exponential growth patterns, making them easier to analyze. A is a constant term representing the log-price level at the critical time t_c or an offset that adjusts the model vertically. The parameter B is a scaling factor that influences the magnitude of the price increase as the market nears the critical time, determining the rate at which the price accelerates as t approaches t_c. The parameter t_c indicates the critical time at which the bubble reaches its peak or a major correction is expected to occur. This is a key parameter of the model that indicates when the bubble will burst. β is a critical exponent, typically ranging between 0.1 and 0.9, that controls the rate of price increase as

the peak date approaches. The closer β is to 0, the more dramatic the increase in price as t_c is approached. Moreover, for this model specification following Grobys (2023), we include the additional conditions: $A > 0, B < 0$.

Substituting the price $P(t)$ with the compounded return while holding the exponent β "fixed" at selected values from the set $\bar{\beta} \in \{0.2, 0.4, 0.6, 0.8\}$ and fixing $B = -1$, Eq. (14) simplifies into a more straightforward form:

$$ln[comp(t)] = Ln[comp(T)] - (t_c - t)^{\bar{\beta}}. \quad (15)$$

Note that $A = Ln[com(T)]$ in association with $t_c = T + 1$ implies that the last compounded return notation of the momentum strategy represents the expected compounded return at the critical time t_c and that the critical event is expected to occur in the following time period.

In the second stage, the model is optimized by minimizing the sum of the squared residuals (*SSR*) for Eq. (15), allowing for the following constraints:

$$\begin{cases} t_c \geq T + 1 \\ 0.1 \leq \beta \leq 0.9 \end{cases}$$

The first constraint allows for the possibility that the finite-time singularity occurs at some unknown point in the future, whereas the second constraint ensures that the compounded return accelerates and remains finite.

In the third stage, using the optimized parameter values obtained from the second stage as the initial values $\Phi_1^* = (A^*, B^*, t_c^*, \beta^*)$, we calibrate a more complex version of the LPPLS model that includes log-periodic oscillations, as shown in Eq. (12). This stage improves the initial estimates by incorporating additional parameters (C, ω, ϕ) that account for the oscillatory behavior around the power-law growth, subject to the following constraints, with an additional constraint on ω added to those from Stage 2:

$$\begin{cases} t_c \geq T + 1 \\ 0.1 \leq \beta \leq 0.9 \\ 5 \leq \omega \leq 15 \end{cases}$$

In accordance with Sornette (2017), the parameter Φ is left unconstrained, as the phase parameter Φ should not be meaningfully restricted. Moreover, the last constraint on ω is based on Sornette's (2017) empirical observations, suggesting a typical range of $5 < \omega < 15$. Finally, the optimized parameter vector for Eq. (12) is reported as $\Phi_2^* = (A^{**}, B^{**}, t_c^{**}, \beta^{**}, C^{**}, \omega^{**}, \Phi^{**})$.

4. Results

In this section, we present the results of the three-stage LPPLS model estimation applied to the five bubbles included in our analysis: December 21, 1929; March 10, 2000; July 14, 2008; January 20, 2016; and April 16, 2020. We also evaluate the predictive power of the LPPLS model by assessing its performance at 3, 6, 9, and 12 months before the bubbles reach their peak. This two-fold analysis aims, first, to examine the forecasting ability of the LPPLS model and, second, to determine whether prediction accuracy improves as the bubble peak approaches.

4.1. Momentum bubble peak on December 21, 1929

For the first bubble in our study, the sample spans from March 11, 1926, to December 21, 1929, and includes 929 daily observations. We apply the LPPLS model at 3, 6, 9, and 12 months prior to the bubble peak on December 21, 1929. Therefore, for the first run of our analysis, 3 months before the peak, the sample spans from March 11, 1926, to September 21, 1929. Panel A of Table 2 presents the initial values from the first-stage estimation approach, based on the basic power-law model described in Eq. (14). In this setting, $A = Ln[comp(T)] = 5.85$ is the natural logarithm of the last compounded return in the sample on September 21, 1929. Given the daily frequency of the sample, where one unit of time equals 1/253 and the sample ends at $T = 3.403$, it follows that $T + 1 = 3.41$. As shown in panel B of Table 2, setting $\beta = 0.2$ leads to the $t_c^* = 3.52$ with $SSR = 4.20$. It should be noted that

these initial model specifications can be considered worst-case scenarios, as they imply an "immediate" occurrence of finite-time singularity, indicating a situation where no further action can be taken.

In the second stage, the model is optimized by minimizing the *SSR* of Eq. (14) by allowing for the following constraints:

$$\begin{cases} t_c \geq T + 1 \\ 0.1 \leq \beta \leq 0.9. \end{cases}$$

Panel B of Table 2 reports the optimal values for the parameter vector $\Phi^*_{\text{model 1}} = (A^*, B^*, t_c^*, \beta^*)$. From the results, the first specification of Model 1 yields the optimal parameter set $\Phi^*_{\text{model 1}} = (6.07, -0.76, 3.52, 0.55)$. In the third stage, using the optimized parameter values of Model 1 in Panel B of Table 2 as the initial values $\Phi^0_{\text{model 2}} = (A^*, B^*, t_c^*, \beta^*, C, \omega, \Phi)$ and setting $C = \phi = 0$ for varying values of $5 \leq \omega \leq 15$, we calibrate LPPLS Model 2, which includes log-periodic oscillations as defined in Eq. (12), by applying the following constraints:

$$\begin{cases} t_c \geq T + 1 \\ 0.1 \leq \beta \leq 0.9 \\ 5 \leq \omega \leq 15 \end{cases}$$

Panel C of Table 2 outlines the input data vectors $\Phi^0_{\text{model 2}} = (A^*, B^*, t_c^*, \beta^*, C, \omega, \Phi)$, while Panel D presents the optimized parameter vectors $\Phi^*_{\text{model 2}} = (A^{**}, B^{**}, t_c^{**}, \beta^{**}, C^{**}, \omega^{**}, \Phi^{**})$ for each iteration. In Panel D, the initial parameters $\Phi^0_{\text{model 2}} = (A^*, B^*, t_c^*, \beta^*, C, \omega, \Phi) = (6.07, -0.76, 3.52, 0.55, 0, 5, 0)$ lead to an optimized setup for Model 2, achieving the lowest $SSR = 1.58$. The resulting optimized parameter vector is $\Phi^*_{\text{model 2}} = (A^{**}, B^{**}, t_c^{**}, \beta^{**}, C, \omega^{**}, \Phi^{**}) = (8.48, -2.72, 4.34, 0.25, 0.02, 9.47, -2.05)$. The estimated parameter $t_c^{**} = 4.34$, corresponds to July 17, 1930, which is 8 months after the actual peak time on December 21, 1929.

In addition, Table 3 presents the estimation outcomes from 6 months prior to the bubble peak, covering the period from March 11, 1926, to June 21, 1929. As shown in Panel A of Table 3, $t_c =$

3.11, and in Panel D, the estimated $t_c^{**} = 4.03$, corresponding to April 11, 1930. Moreover, Table 4 reports the predictions made 9 months before the peak on December 21, 1929, with the dataset extending from March 11, 1926, to March 21, 1929. As shown in Panel A of Table 4, $t_c = 2.81$, and in Panel D, the estimated $t_c^{**} = 3.05$, with the lowest $SSR = 1.34$, corresponding to June 4, 1929. Finally, Table 5 shows the estimated data from 12 months before the peak on December 21, 1929, spanning from March 11, 1926, to December 21, 1928. As shown in Panel A of Table 5, $t_c = 2.53$, and in Panel D, the estimated $t_c^{**} = 5.15$ corresponds to March 24, 1931.

4.2. Momentum bubble peak on March 10, 2000

For the second momentum bubble, our dataset spans from November 20, 1939, to March 10, 2000, encompassing 15662 daily observations. Similarly, we apply the LPPLS model to periods 3, 6, 9, and 12 months prior to the peak on March 10, 2000. Thus, for the period 3 months before the peak, the dataset covers November 20, 1939, to December 10, 1999. Panel A of Table 6 reports the initial values of the first-stage estimation described by Eq. (14), where $A = Ln[comp(T)] = 13.28$ represents the natural logarithm of the last compounded return in the sample on December 10, 1999. Again, considering the daily frequency of the sample, where each unit of time equals 1/253 and ends at $T = 61.66$, it follows that $T + 1 = 61.67$. According to Panel A of Table 6, setting $\beta = 0.4$ results in the lowest $SSR = 74337.24$.

In the second stage, the model is optimized by minimizing the *SSR* of Eq. (14), subject to the following constraints:

$$\begin{cases} t_c \geq T + 1 \\ 0.1 \leq \beta \leq 0.9. \end{cases}$$

According to Panel B of Table 6, the third specification of Model 1 yields the optimal parameter set $\Phi^*_{\text{model 1}} = (A^*, B^*, t_c^*, \beta^*) = (14.82, -0.40, 66.82, 0.79)$, with the lowest $SSR = 1096.97$.

In the third stage, we use the optimized parameter values from Model 1 in Panel B of Table 6 as the initial values $\Phi^0_{\text{model 2}} = (A^*, B^*, t_c^*, \beta^*, C, \omega, \Phi)$, with $C = \phi = 0$ for the varying values of $5 \leq \omega \leq 15$, and calibrate LPPLS Model 2 as described in Eq. (12), subject to the following constraints:

$$\begin{cases} t_c \geq T + 1 \\ 0.1 \leq \beta \leq 0.9 \\ 5 \leq \omega \leq 15 \end{cases}$$

From Panel D of Table 6, the initial parameters $\Phi^0_{\text{model 2}} = (A^*, B^*, t_c^*, \beta^*, C, \omega, \Phi) = (14.82, -0.40, 66.82, 0.79, 0, 5, 0)$ yield the optimized parameter vector $\Phi^*_{\text{model 2}} = (A^{**}, B^{**}, t_c^{**}, \beta^{**}, C^{**}, \omega^{**}, \Phi^{**}) = (13.43, -0.25, 62.23, 0.88, 0.07, 5.04, 0.31)$ for Model 2, with the lowest $SSR = 302.26$, where the estimated parameter $t_c^{**} = 62.23$, corresponds to July 6, 2000.

Next, Table 7 shows the estimation results from 6 months prior to the peak, covering the period from November 20, 1939, to September 10, 1999. As shown in Panel A of Table 3, $t_c = 61.42$, and in Panel D, the estimated $t_c{}^{**} = 62.14$ corresponds to the lowers $SSR = 301.08$, indicating June 1, 2000, which is 3 months after the actual peak date.

Table 8 reports the predictions made 9 months prior to the peak date, using data spanning from November 20, 1939, to June 10, 1999. As shown in Panel A of Table 8, $t_c = 61.16$, and in Panel D, the estimated $t_c{}^{**} = 62.04$ with the lowest $SSR = 296.63$ corresponds to April 26, 2000. Finally, Table 9 presents the estimated results from 12 months before the peak date, spanning from November 20, 1939, to March 10, 1999. As shown in Panel A of Table 9, $t_c = 60.91$, and in Panel D, the estimated $t_c{}^{**} = 62.04$ yields the lowest $SSR = 293.92$, corresponding to May 1, 2000.

4.3. Momentum bubble peak on July 14, 2008

For the third momentum bubble on July 14, 2008, the dataset spans from September 30, 2002, to July 14, 2008, encompassing 1457 daily observations. Similarly, we apply the LPPLS model to periods 3,

6, 9, and 12 months prior to the bubble peak. For the period 3 months before the peak on July 14, 2008, the dataset covers September 30, 2002, to April 14, 2008. Panel A of Table 10 reports the initial value of the first-stage estimation described by Eq. (14), where $A = Ln[comp(T)] = 14.40$, which represents the natural logarithm of the last compounded return in the sample on April 14, 2008. Again, considering the daily frequency of the sample in which each unit of time equals 1/253 and ends at $T = 5.51$, it follows that $T + 1 = 5.52$. According to Table 10, setting $\beta = 0.2$ leads to the lowest $SSR = 776.12$.

In the second stage, the model is optimized by minimizing the *SSR* of Eq. (14) by allowing for the following constraints:

$$\begin{cases} t_c \geq T + 1 \\ 0.1 \leq \beta \leq 0.9. \end{cases}$$

According to Panel B of Table 10, the third specification of Model 1 yields the optimal parameter set $\Phi^*_{\text{model 1}} = (A^*, B^*, t_c^*, \beta^*) = (14.67, -0.67, 5, 52, 0.10)$, with the lowest $SSR = 21.06$. From the results in Panel B, regardless of the specific model specification, all optimized models consistently predict the finite-time singularity occurring at $t_c^* = T + 1 = 5.52$.

In the third stage, we use the optimized parameter values from Model 1 in Panel B of Table 10 as the initial values $\Phi^0_{\text{model 2}} = (A^*, B^*, t_c^*, \beta^*, C, \omega, \Phi)$, with $C = \phi = 0$ for varying values of $5 \leq \omega \leq 15$, and calibrate LPPLS Model 2 as defined in Eq. (12), subject to the following constraints:

$$\begin{cases} t_c \geq T + 1 \\ 0.1 \leq \beta \leq 0.9 \\ 5 \leq \omega \leq 15 \end{cases}$$

In Panel D of Table 10, the initial parameters $\Phi^0_{\text{model 2}} = (A^*, B^*, t_c^*, \beta^*, C, \omega, \Phi) = (14.67, -0.67, 5.52, 0.1, 0, 6, 0)$ yield the optimized parameter vector $\Phi^*_{\text{model 2}} = (A^{**}, B^{**}, t_c^{**}, \beta^{**}, C^{**}, \omega^{**}, \Phi^{**}) = (14.74, -0.73\ , 5.52, 0.10,\ \ 0.14, 6{,}38, -1{,}49)$ for Model 2, with the lowest $SSR = 12.38$. Since the estimated parameter $t_c^{**} = 5.52$, we conclude that the LPPLS

model cannot accurately forecast the exact timing of the crash in this case and only replicates the initial parameter t_c from the first stage. Table 11 presents the estimation results from 6 months prior to the bubble peak time, covering the period from September 30, 2002, to January 14, 2008. From Panel A of Table 11,$t_c = 5.27$, and Panel D, the estimated ${t_c}^{**} = 5.41$ corresponds to March 11, 2008. Thus, the estimated ${t_c}^{**}$ is 4 months earlier than the actual peak.

Additionally, Table 12 reports the predictions made 9 months prior to the peak, using data spanning from September 30, 2002, to October 12, 2007. From Panel A of Table 12, $t_c = 5.02$, and in Panel D, the estimated ${t_c}^{**} = 5.61$ corresponds to May 21, 2008. It should be noted that, as the initial estimated ${t_c}^{*}$ from Panel B of Table 8 exceeds $2T$ (not reported), we follow Sornette (2017) and include the additional condition $t_c \leq 2T$. Finally, Table 13 reports the estimated results from 12 months before the peak, spanning from September 30, 2002, to July 12, 2007. As shown in Panel A of Table 13, $t_c = 4.76$, and in Panel D, the estimated ${t_c}^{**} = 5.32$ corresponds to February 6, 2008.

4.4. Momentum bubble peak on January 20, 2016

For the fourth bubble, the dataset spans from March 6, 2009, to January 20, 2016, and includes 1731 daily observations. Applying the LPPLS model for 3, 6, 9, and 12 months before the bubble peak, the estimation results are presented in Tables 14, 15, 16, and 17, respectively. Panel A of Table 14 presents the initial values from the first-stage estimation approach described in Eq. (14), using data from 3 months before January 20, 2016. In this setting, $A = Ln[comp(T)] = 14.09$ corresponds to the last natural logarithm of compounded return in the sample on October 20, 2015. Again, as one unit of time equals 1/253 and ends at $T = 6.59$, it follows that $T + 1 = 6.60$. According to Panel A of Table 14, $\beta = 0.2$ results in the lowest $SSR = 1004.55$. In the second stage, the model is optimized by minimizing the SSR of Eq. (14) by allowing for the following constraints:

$$\begin{cases} t_c \geq T + 1 \\ 0.1 \leq \beta \leq 0.9. \end{cases}$$

Panel B of Table 14 reports the optimal values for the parameter vector $\Phi^*_{\text{model 1}} = (A^*, B^*, t_c^*, \beta^*)$. From the results, the third specification ($\beta = 0.6$) of Model 1 yields the optimal parameter set $\Phi^*_{\text{model 1}} = (14.40, -0.50, 6.89, 0.36)$, with the lowest $SSR = 35.67$. The critical time, t_c^*, is estimated to be $t_c^* = 6.89$. In the final stage, we use the optimized parameters from Model 1 and set $C = \phi = 0$ for varying ω in the range of $5 \leq \omega \leq 15$ to calibrate LPPLS Model 2, subject to the following constraints:

$$\begin{cases} t_c \geq T + 1 \\ 0.1 \leq \beta \leq 0.9 \\ 5 \leq \omega \leq 15. \end{cases}$$

Panel C of Table 14 outlines the input data vectors $\Phi^0_{\text{model 2}} = (A^*, B^*, t_c^*, \beta^*, C, \omega, \Phi)$, while Panel D presents the optimized parameter vectors $\Phi^*_{\text{model 2}} = (A^{**}, B^{**}, t_c^{**}, \beta^{**}, C^{**}, \omega^{**}, \Phi^{**})$ for each iteration. From Panel D, the initial parameters $\Phi^0_{\text{model 2}} = (A^*, B^*, t_c^*, \beta^*, C, \omega, \Phi) = (14.40, -0.50, 6.89, 0.36, 0, 12, 0)$ lead to an optimized setup for Model 2, achieving the lowest $SSR = 27.90$. The estimated $t_c^{**} = 7.58$ corresponds to October 19, 2016, which is 9 months after the actual peak.

Moreover, Table 15 outlines the predictions made 6 months before the peak, with the dataset spanning from March 6, 2009, to July 20, 2015. As shown in Panel A of Table 15, $t_c = 6.34$, and in panel D, the estimated $t_c{}^{**} = 8.18$, yields the lowest $SSR = 25.47$, corresponding to May 23, 2017, which is 16 months after the actual peak. Table 16 presents the estimation outcomes from 9 months prior to the peak, covering the period from March 6, 2009, to April 20, 2015. From Panel A of Table 16, the initial $t_c = 6.10$, and in Panel D, the estimated $t_c{}^{**} = 7.89$, yields the lowest $SSR = 25.06$, corresponding to February 7, 2017, which is 13 months after the actual event. Finally, Table 17 presents the prediction results from 12 months prior to the peak, spanning from March 6, 2009, to January 20, 2015. From Panel A of Table 17, $t_c = 5.85$, and in Panel D, the estimated $t_c{}^{**} = 8.41$,

yields the lowest *SSR* = 24.97, corresponding to August 16, 2017. Note that, in this case, as the initial estimated t_c^* exceeds $2T$, we include the additional condition $t_c \leq 2T$.

4.5. Momentum bubble peak on April 16, 2020

The data for the last bubble, spanning from January 20, 2016, to April 16, 2020, includes 1068 daily observations. The LPPLS model is similarly applied for 3, 6, 9, and 12 months prior to the bubble peak date, with the estimation results for these periods presented in Tables 18, 19, 20, and 21, respectively.

Based on the findings in Table 18, the LPPLS model predicts the peak occurrence within a narrow time frame, approximately 3 days after the initial $t_c = 3.98$. This critical time coincides with January 23, 2020, as indicated by $t_c^{**} = 3.99$, with a corresponding minimum $SSR = 2.90$. Table 19 reports the predictions made 6 months before the peak, with the sample spanning January 20, 2016 to October 16, 2019. From Panel A of Table 19, $t_c = 3.73$, and in Panel D, the estimated $t_c^{**} = 3.98$, with the lowest $SSR = 2.80$ corresponding to January 21, 2020, which is 3 months prior to the actual peak on April 16, 2020. Also, Table 20 outlines the estimation results 9 months prior to the peak, with the sample spanning from January 20, 2016 to July 16, 2019. From Panel A of Table 20, the initial $t_c = 3.47$, and in Panel D, the estimated $t_c^{**} = 4.01$, with the lowest $SSR = 2.33$ corresponding to January 29, 2020, which is 3 months prior to the actual peak. Finally, Table 21 presents the prediction results 12 months before the peak, spanning from January 20, 2016 to April 16, 2019. From Panel A of Table 21, $t_c = 3.23$, and in Panel D, the estimated $t_c^{**} = 4.31$, with the lowest $SSR = 1.94$ corresponding to May 15, 2020.

5. Robustness Check

To validate the reliability of our predictions and ensure that the model's forecasts are not merely the result of data selection biases or methodological idiosyncrasies, we perform robustness tests.

According to Lin et al. (2014), residuals from the LPPLS model should exhibit stationary, mean-reverting behavior, as the presence of non-stationary residuals could indicate a spurious regression, thereby confusing the clarity of statistical analyses. We examine these residuals using the Augmented Dickey–Fuller (ADF) test. According to Jiang et al. (2010), residuals that demonstrate stationarity at the 99% confidence level should pass the test, suggesting the statistical significance of the LPPLS model. Therefore, if the residual process is stationary at the 1% significance level, the LPPLS signature is statistically significant.

In this context, and following Grobys (2023), we perform the ADF test on the residuals of the LPPLS model using the following regression equation:

$$\Delta e_t = \delta_0 + \delta_1 t + \delta_2 u_{t-1} + \gamma_1 \Delta u_{t-1} + \cdots + \gamma_p \Delta u_{t-p} + \varepsilon_t, \tag{16}$$

where Δe_t represents the residuals of the LPPLS model, defined as the difference between $Ln[comp(t)]$ and the predicted value from the LPPLS model, Eq. (12), at time t. The parameters $\gamma_1, \dots, \gamma_p$ measure exposures to $\Delta u_{t-1}, \dots, \Delta u_{t-p}$ and ε_t is assumed to be a white noise process. Setting $\delta_0 = \delta_1 = 0$ corresponds to a model specification with no deterministic terms, while $\delta_1 = 0$ represents a model that includes only a constant term. In addition, the parameters $\gamma_1, \dots, \gamma_p$ correspond to the model's autoregressive components. The model with $\delta_0 = \delta_1 = 0$ represents a random walk under the null hypothesis, whereas setting $\delta_1 = 0$ models a random walk with drift under the null hypothesis. Consequently, the ADF test is structured in three main forms, with the null hypothesis stating that $\hat{\delta}_2 = 0$, against the alternative hypothesis that $\hat{\delta}_2 < 0$. The computed test statistic $\hat{\lambda}$ is the *t*-statistic that corresponds to the estimated coefficient $\hat{\delta}_2$. The estimated test statistics and the corresponding *p*-values for various ADF tests for the bubbles on December 21, 1929, March 10, 2000, July 14, 2008, January 20, 2016 and April 16, 2020, are reported in Panels A, B, C, D and E of Table 22, respectively. The critical values for the 5% and 1% significance levels are presented in columns

4 and 5, respectively. It should be noted that a lag order of $p = 2$ is used for all model specifications, as recommended by the Schwarz information criterion.

The results of the ADF test performed for the bubble peak on December 21, 1929, in Panel A of Table 22 show significant predictive capability across all three model specifications (i.e., no deterministic terms, constant, and constant with trend) for all time frames. The computed test statistics ($\hat{\lambda}$) for the data points 3, 6, 9, and 12 months prior to the bubble peak indicate statistically significant values, with *p*-values well below the thresholds of $p < 0.05$ and $p < 0.01$. This provides strong evidence against the null hypothesis of non-stationarity in the residuals, thereby confirming the reliability of the LPPLS model's predictions for the bubble peak on December 21, 1929. Additionally, Panel B of Table 22 presents the results for the bubble peak on March 10, 2000, with significant $\hat{\lambda}$ values and low *p*-values, confirming residual stationarity and validating the LPPLS model's predictions.

Further, Panel C of Table 22 reports the ADF test results for the residual series of the LPPLS calibration conducted 6, 9, and 12 months before the bubble peak on July 14, 2008 (the 3 months calibration is not reported, as it did not yield any predictions). The computed test statistics $\hat{\lambda}$ for the data 9 and 12 months prior to the peak on July 14, 2008, exhibit statistically significant values, supported by *p*-values well below the significance levels, suggesting strong evidence against the null hypothesis of non-stationarity in the residuals, thereby confirming the robustness of the LPPLS model's predictions. However, the residuals obtained from prediction using data 6 months prior to the peak do not exhibit statistically significant results, as indicated by the ADF test. Across all model specifications (i.e., no deterministic terms, constant, and constant with trend), the computed test statistics ($\hat{\lambda}$) fail to reach the levels of 1% or 5% significance, with *p*-values above the threshold used for statistical significance. This suggests that for the 6 months before the peak on July 14, 2008, there is insufficient evidence to reject the null hypothesis of non-stationarity in the residuals. These results indicate a potential limitation or lack of predictive power of the LPPLS model for this specific period

before the peak. This implies that the model may not effectively capture the underlying dynamics during this time frame, highlighting the importance of considering additional factors or refining the model's specifications to enhance its predictive performance.

In addition, as shown in Panels D and E of Table 22 for the bubbles peaks on January 20, 2016, and April 16, 2020, respectively, the ADF test results demonstrate similarly significant outcomes across different model specifications and time frames in both panels. The computed test statistics ($\hat{\lambda}$) consistently exhibit statistically significant values, accompanied by remarkably low *p*-values, providing convincing evidence against the presence of non-stationarity in the residual series of these predictions. Based on the results in Table 22, we conclude that the residuals of the LPPLS model adhere to a stationary mean-reverting process. This observation suggests that the LPPLS signature holds statistical significance, strengthening confidence in the model's predictive capabilities. In addition, Figures A.1, A.2, A.3, A.4, and A.5 plot the residuals covering the in-sample time window for each of the five bubbles analyzed, with respect to the 3, 6, 9, and 12 months prior to each peak date. Visual inspection of these figures indicates that the residuals exhibit stationarity, despite presumably high levels of autocorrelation.

6. Discussion

In this study, we investigate the capability of the LPPLS model in predicting five significant momentum bubbles using a sample from November 3, 1926, to November 30, 2023. Table 23 presents the comparative results of the LPPLS model's forecasting ability for the five momentum bubbles using data from 3, 6, 9, and 12 months prior to each bubble peak. The results reveal the LPPLS model's performance across different forecasting intervals.

The first row in Table 23 compares the prediction outcomes for the momentum bubble peak of December 21, 1929, using almost four years of data preceding the bubble peak. According to the results, when using data up to 3 months before the peak, the model predicts the bubble would burst

seven months later than the actual event. Using data up to 6 months prior to the peak date, the LPPLS model predicts the peak will occur approximately 4 months later than the actual peak. However, using data from 9 months prior, the LPPLS model predicts the peak will occur 6 months earlier than it actually did. Finally, using data up to 12 months before the peak, the predicted peak occurs 15 months after the actual peak. The 9 months window, in particular, stands out for predicting an earlier peak, highlighting the importance of selecting an appropriate timeframe for optimal accuracy.

The second row in Table 23 presents the prediction results for the bubble peak on March 10, 2000. The results show that when using data up to 3 and 6 months before the peak date, the model predicts the peak occurring 4 and 3 months earlier than the actual event, respectively. In contrast, when using data from 9 and 12 months prior to the peak, the model predicts the peak will occur approximately one month after the actual event.

The third row in Table 23 compares the prediction results for the momentum bubble peak on July 14, 2008, using approximately six years of data leading up to the bubble peak. Although the LPPLS model failed to forecast the peak date when using data up to 3 months beforehand, the predictions made using data from 6, 9, and 12 months prior provide more accurate approximations. For instance, using data from 6 months before the peak, the model predicts the peak occurring 4 months earlier than the actual event. Using data from 9 months prior, the prediction is approximately 2 months before the peak, and with data from 12 months before, the estimation is 5 months earlier than the actual peak. Based on these results, while the model may not always predict the exact timing of the peak, it offers investors a valuable window of opportunity to mitigate potential losses. Thus, despite the lack of a clear pattern when using data up to 3, 6, 9, and 12 months prior to the bubble peak on July 14, 2008, the LPPLS model's predictions indicate that the peak is likely to occur within a time frame of 2–5 months before the actual event. This time frame is crucial for investors because although the model may not predict the exact date of the peak, it provides a warning signal that a crash is imminent. This

warning allows investors to "close" or exit their positions before the bubble burst, thereby protecting their investments from its full impact.

The fourth row in Table 19 presents the estimation results for the third bubble peak on January 20, 2016, using approximately seven years of data. Using data up to 3 months before the bubble peak, the model predicts the peak occurring 9 months later than the actual event. However, predictions using data from 6, 9, and 12 months prior place the peak date more than a year later, in 2017. Finally, the fifth row of Table 19 shows the findings for the most recent bubble in our analysis, which occurred on April 16, 2020, using approximately four years of data. Based on these findings, using data from 3, 6, and 9 months before the peak, the model predicts the event occurring 3 months prior to the actual peak, whereas using data from 12 months before, the model predicts the peak to occur one month after the actual event.[4]

Although our predictions are not accurate in determining the exact timing of market bubbles, they provide a practically useful estimate with a maximum deviation of 19 months after the actual peak. In this regard, Sornette (2017) argues that achieving absolute prediction accuracy in financial systems is impossible due to the complex and unpredictable dynamics of financial markets. This underscores the fact that the LPPLS model is inherently probabilistic rather than deterministic (Sornette, 2017). Zhou and Sornette (2006) also acknowledge the model's sensitivity to parameter selection and the challenge of distinguishing between genuine bubble signals and the noise inherent in financial data. Bree and Joseph's (2013) critical analysis of the LPPLS model's application in historical bubbles supports the view that, although the model is capable of identifying periods of strong speculation that

[4] We acknowledge that the COVID-19 market crash in early 2020 represents a special case. Unlike typical bubbles that result from endogenous market dynamics—often driven by herding behavior, speculative momentum, and feedback loops—the COVID-19 crisis was triggered by an abrupt, exogenous health shock with profound global economic implications. This distinction is supported by macro-financial indicators such as the U.S. Treasury yield curve, which maintained a normal upward slope in the lead-up to the crash, contrasting with the inverted or flat yield curves that typically precede endogenous market downturns. Accordingly, while the LPPLS model may still capture certain speculative dynamics preceding the COVID-19 crash, its applicability for identifying the tipping point in this specific context should be interpreted with caution.

may lead to a crash, accurately pinpointing the exact timing is challenging due to the complex interaction of market factors and external influences.

In addition, the last row of Table 23 provides the average of predictions with respect to different time windows. Based on the results, prediction accuracy improves as we approach the bubble peak, with the deviation reducing from +6.2 months at 12 months before the event to +3.4 months at 3 months prior. These results confirm previous literature (Sornette, 2017), suggesting that the prediction accuracy of the LPPLS model improves as the critical time approaches. Notably, the 9 months window provides the most accurate average forecast, with a deviation of only +0.6 months, suggesting that predictions made within a relatively short period before the actual peak are more aligned with real-time market conditions.

Further, the predicted times from the LPPLS model, on average, tend to occur later than the actual events. However, this tendency aligns with Sornette's (2017) observation that the model generally "overshoots" the true crash date. According to Sornette (2017), the critical time t_c does not represent the exact crash date but rather the most probable timing for a crash, where the asymmetric distribution of potential bubble burst dates reaches its peak. Sornette (2017) describes the crash as a stochastic event with a probability that increases as t_c approaches. As a result, model fits typically yield t_c values that are close to, but usually later than the actual crash. This pattern arises because t_c is part of the log-periodic power-law structure, while the crash itself is randomly triggered with a probability that increases sharply as t_c approaches. This effect might be more pronounced as we approach the peak of the bubble and uncertainty in the endogenous system increases.

The results provide valuable insights into the two main objectives of this study. First, by fitting the LPPLS model in the five significant momentum bubbles identified in our data we successfully detect the LPPLS signatures prior to each bubble peak. This empirical evidence verifies the hypothesis that LPPLS signatures serve as early indicators of significant momentum bubble formations, thus

validating the model's capability to capture the dynamics leading up to a bubble formation. Second, our analysis of prediction accuracy using data from 3, 6, 9, and 12 months prior to the bubble peak reveals a comprehensive view of the model's performance. Based on the results, the prediction accuracy does not exhibit a completely linear improvement as the peak approaches, but a partial improvement is observed as the peak looms closer. The average prediction error for the 3 months window is +3.4, suggesting a relatively strong predictive capability when closer to the bubble peak. However, the 12 months window shows a significantly higher average error of +6.2, indicating that while some signals might be captured earlier, these longer-term forecasts could potentially be influenced by broader market trends or noise, which sometimes leads to overestimations. The average errors for the 6 and 9 months windows, recorded at +3.2 and +0.6 respectively, reflect a fluctuating accuracy that does not show a clear progression as the bubble approaches. This lack of a consistently improving trend suggests that the model's performance is not solely dependent on proximity to the peak but is also influenced by other factors, such as market volatility and the unique characteristics of each bubble. Overall, these average values highlight the need for a careful interpretation of the LPPLS model's predictions, recognizing both its strengths and limitations in different temporal contexts. In addition, the 9 months prediction values, in particular, indicate a period where the model's forecasting accuracy is strongest, as the deviation of +0.6 is smaller compared to predictions made 3, 6, and 12 months before the peak.

The observed nonlinear improvement in prediction accuracy as the critical time approaches may be attributed to various factors. First, the "herding behavior" often observed in financial markets can intensify as a bubble approaches its peak, making predictions more difficult (Tsuchiya, 2021; Shantha, 2019). Second, we use Grobys' (2023) three-stage estimation approach, which starts with a simplified version of the model and gradually adds complexity. While this approach allows for a more accurate representation of market dynamics, it may also make the model more sensitive to recent market behavior. Therefore, if market behavior changes dramatically as the critical time approaches,

the model might struggle to adjust its parameters quickly enough to predict the peak time accurately. Finally, zero-cost portfolios, such as those used in this study, may behave differently from other types of portfolios, making them more sensitive to market dynamics and potentially more challenging to predict.

7. Conclusion

In this study, we investigate the presence of LPPLS signatures before five significant momentum bubble bursts on December 21, 1929, March 10, 2000, July 14, 2008, January 20, 2016, and April 16, 2020. By applying the LPPLS model to predict the peak dates of these bubbles and analyzing prediction accuracies at intervals of 3, 6, 9, and 12 months prior to each bubble peak, this study confirms the presence of LPPLS signatures before significant momentum bubble formations, thus supporting the hypothesis that these signatures can act as early indicators of bubble formation in a zero-cost trading strategy. This not only validates the LPPLS model's theoretical foundation but also emphasizes its practical utility.

Although our results show that the prediction accuracy of the LPPLS model does not improve linearly as the bubble peaks approach, a partial improvement is still observed. The average prediction error for the 3 months window is +3.4, reflecting strong accuracy when approaching the bubble peak. In contrast, using data from 12 months before the peak results in a higher average error of +6.2, indicating that while longer-term forecasts may detect signals earlier, they are more susceptible to broader market trends and noise, which can lead to overestimations. In the cross-sectional comparison of all five bubbles, the LPPLS model provides more accurate predictions when using data from 9 months prior to the critical time, compared to data from 3, 6, or 12 months earlier. This highlights the inherent complexity of financial markets and the challenges in forecasting their behavior.

Overall, while the LPPLS model may not deliver absolute precision in predicting momentum bubbles, its value lies in its ability to offer a probabilistic assessment that can guide investors in managing

investment risk, particularly in terms of exiting trading positions. Further research is encouraged to explore the model's applicability across different investment strategies and market conditions, with the ultimate goal of enhancing the predictability and manageability of severe market downturns.

Our novel findings contribute significantly to the field of financial market analysis by providing empirical support for the predictive power of the LPPLS model in relation to significant momentum bubbles and drawdowns. However, our results should be viewed in light of the acknowledged limitations of the LPPLS model, as it serves more as a tool for risk assessment and the prediction of market corrections rather than a precise timing mechanism. It is also important to note that, while the LPPLS model has been successful in predicting certain market bubbles, it should be considered one of several tools within a comprehensive market analysis strategy. Comparing the model's predictions with other indicators and models can provide a more balanced and nuanced understanding of market risks.

This study offers a behavioral explanation for the documented optionality effect of momentum bubbles. Specifically, the dynamics of imitation follow a power law, where a larger group of traders tends to behave like a smaller group, amplifying collective market movements and contributing to bubble formations. This herd-like behavior can have significant implications for market dynamics, particularly during periods of market stress. A momentum bubble refers to a scenario in which a momentum portfolio contains low-beta stocks in the long leg and high-beta stocks in the short leg. Such a bubble burst when the high-beta stocks (from the short leg) experience an unexpected increase in value, outperforming the low-beta stocks (from the long leg) and disrupting the expected momentum pattern. A momentum bubble bursts when high-beta stocks outperform low-beta stocks as the market approaches a reversal. In this scenario, the larger group of traders follows the smaller group, which persists in remaining unchanged in its positions despite the market dynamics, indicating a need for adjustment. This pattern of behavior provides a new perspective on the dynamics of financial bubbles. This suggests that the actions of a small group of traders can have a

disproportionate impact on market outcomes, particularly during periods of high uncertainty. These results are in line with Yan et al. (2023), which demonstrate that slow transmission of correct information leads to the persistence of the momentum effect, while slow transmission of rumors contributes to reversal effects and bubble phenomena. Our identification of LPPLS signatures before major momentum bubbles formations supports this by highlighting how delays in information flow can sustain market momentum before eventual corrections occur. This finding highlights the importance of understanding behavioral factors in market analysis and the need for further research in this area.

References

Avery, C., & Zemsky, P. (1998). Multidimensional uncertainty and herd behavior in financial markets. *American Economic Review*, *88*(4), 724-748.

Barroso, P., & Santa-Clara, P. (2015). Momentum has its moments. *Journal of Financial Economics*, *116*, 111–120.

Bikhchandani, S., Welch, I., & Hirshleifer, D. (1992). A theory of fads, fashion, custom, and cultural change as informational cascades. *Journal of Political Economy*, *100*(5), 992-1026.

Bree, D. S., & Joseph, N. L. (2013). Fitting the log periodic power law to financial crashes: A critical analysis. *Physica A: Statistical Mechanics and its Applications*, *392*(16), 3444–3452.

Brée, D. S., Challet, D., P. P. Peirano (2013). Prediction accuracy and sloppiness of log-periodic functions. *Quantitative Finance 13*(2), 275—280.

Brogaard, J., Hendershott, T., & Riordan, R. (2014). High-frequency trading and price discovery. *The Review of Financial Studies*, *27*(8), 2267–2306.

Brunnermeier, M. K. (2008). Deciphering the liquidity and credit crunch 2007–2008. *Journal of Economic Perspectives*, *23*(1), 77–100.

Chabot, B., Ghysels, E., & Jagannathan, R. (2014). Momentum trading, return chasing, and predictable crashes (Working Paper No. 20660). *National Bureau of Economic Research*.

Chang, G., & Feigenbaum, J. (2008). Detecting log-periodicity in a regime-switching model of stock returns. *Quantitative Finance*, *8*(7), 723–738.

Cont, R. (2001). Empirical properties of asset returns: Stylized facts and statistical issues. *Quantitative Finance*, *1*(2), 223–236.

Daniel, K., & Moskowitz, T. J. (2016). Momentum crashes. *Journal of Financial Economics*, *122*(2), 221–247.

Daniel, K., Hirshleifer, D., & Subrahmanyam, A. (1998). Investor Psychology and Security Market under- and Overreactions. *The Journal of Finance*, *53*(6), 1839–1885.

De Long, J. B., Shleifer, A., Summers, L. H., & Waldmann, R. J. (1990). Positive feedback investment strategies and destabilizing rational speculation. *Journal of Finance*, *45*(2), 379-395.

Derrida, B., De Seze, L., & Itzykson, C. (1983). Fractal structure of zeros in hierarchical models. *Journal of Statistical Physics*, *33*(3), 559–569.

Feigenbaum, J. (2001). A statistical analysis of log-periodic precursors to financial crashes. *Quantitative, Finance, 1*, 346–360.

Feigenbaum, J. A., & Freund, P. G. (1996). Discrete scale invariance in stock markets before crashes. *International Journal of Modern Physics B, 10*, 3737–37451.

Filimonov, V., & Sornette, D. (2013). A stable and robust calibration scheme of the log-periodic power law model. *Physica A: Statistical Mechanics and its Applications*, *392*(17), 3698-3707.

Filimonov, V., & Sornette, D. (2015). Power law scaling and "Dragon-Kings" in distributions of intraday financial drawdowns. *Chaos, Solitons & Fractals, 74*, 27–45.

Filimonov, V., Demos, G., & Sornette, D. (2016). Modified profile likelihood inference and interval forecast of the burst of financial bubbles. *Swiss Finance Institute Research Paper No. 16–12*.

Gandal, N., Hamrick, J. T., Moore, T., & Oberman, T. (2018). Price manipulation in the Bitcoin ecosystem. *Journal of Monetary Economics*, *95*, 86–96.

Geraskin, P., & Fantazzini, D. (2013). Everything you always wanted to know about log-periodic power laws for bubble modeling but were afraid to ask. *The European Journal of Finance*, *19*(5), 366–391.

Gerlach, J. C., Demos, G., & Sornette, D. (2019). Dissection of Bitcoin's multiscale bubble history from January 2012 to February 2018. *Royal Society Open Science*, *6*(7), 180643.

Graf, V. B., & Meister, C. (2003). Predicting critical crashes? A new restriction for the free variables. *Physica A: Statistical Mechanics and its Applications*, 539–547.

Grobys, K. (2023). A finite-time singularity in the dynamics of the US equity market: Will the US equity market eventually collapse? *International Review of Financial Analysis*, *89*, 102787.

Hoberg, G., & Phillips, G. (2010). Real and financial industry booms and busts. *Journal of Finance*, *65*(1), 45–86.

Hong, H., & Stein, J. C. (1999). A unified theory of underreaction, momentum trading, and overreaction in asset markets. *The Journal of Finance*, *54*(6), 2143-2184.

Hou, K., Xue, C., & Zhang, L. (2020). Replicating anomalies. *The Review of Financial Studies*, *33*(5), 2019–2133.

Jegadeesh, N., & Titman, S. (1993). Returns to buying winners and selling losers: Implications for stock market efficiency. *Journal of Finance*, *48*(1), 65–91.

Jiang, Z. Q., Zhou, W. X., Sornette, D., Woodard, R., Bastiaensen, K., & Cauwels, P. (2010). Bubble diagnosis and prediction of the 2005–2007 and 2008–2009 Chinese stock market bubbles. *Journal of Economic Behavior & Organization*, *74*(3), 149–162.

Johansen, A., & Sornette, D. (1998). Stock market crashes are outliers. *The European Physical Journal B - Condensed Matter and Complex Systems*, *1*(2), 141–143.

Johansen, A., & Sornette, D. (2001). Large stock market price drawdowns are outliers. *Journal of Risk*, *4*(2), 69–110.

Johansen, A., & Sornette, D. (2010). Shocks, crashes and bubbles in financial markets. *Brussels Economic Review (Cahiers Économiques de Bruxelles)*, *53*(2), 201–253.

Johansen, A., Ledoit, O., & Sornette, D. (2000). Crashes as critical points. *International Journal of Theoretical and Applied Finance*, *3*(02), 219–255.

Johansen, D., Sornette, D., & Ledoit, O. (1999). Predicting financial crashes using discrete scale invariance. *Journal of Risk*, *1*(4), 5–32.

Kindleberger, C. P., & Aliber, R. Z. (2005). *Manias, panics, and crashes: A history of financial crises (5th ed.).* John Wiley & Sons.

Lin, L., Ren, R. E., & Sornette, D. (2014). The volatility-confined LPPL model: A consistent model of "explosive" financial bubbles with mean-reverting residuals. *International Review of Financial Analysis*, *33*, 210–225

Lux, T. (1995). Herd behavior, bubbles and crashes. *The Economic Journal*, *105*(431), 881-896.

Mandelbrot, B. (1963). The variation of certain speculative prices. *Journal of Business*, *36*, 394–419.

Reinhart, C. M., & Rogoff, K. S. (2009). *This time is different: Eight centuries of financial folly*. Princeton University Press.

Shantha, K. V. A. (2019). The evolution of herd behavior: Will herding disappear over time? *Studies in Economics and Finance*, *36*(3), 637–661.

Sornette, D. (2003). Critical market crashes. *Physics Reports*, *378*(1), 1–98.

Sornette, D. (2017). Why stock markets crash: Critical events in complex financial systems. *Princeton University Press.*

Sornette, D., & Cauwels, P. (2014). Financial bubbles: Mechanisms and diagnostics. *Swiss Finance Institute Research Paper No. 14–28*.

Sornette, D., & Johansen, A. (2001). Large financial crashes. *Physica A: Statistical Mechanics and its Applications*, *299*(1–2), 40–59.

Sornette, D., Johansen, A., & Bouchaud, J. P. (1996). Stock market crashes, precursors, and replicas. *Journal de Physique I*, *6*(1), 167–175.

Tsuchiya, Y. (2021). Crises, market shocks, and herding behavior in stock price forecasts. *Empir Econ, 61*, 919–945

Vandewalle, N., Ausloos, M., Boveroux, P., & Minguet, A. (1999). Visualizing the log-periodic pattern before crashes. *European Physics Journal B*, *9*, 355–359.

Wheatley, S., Sornette, D., Huber, T., Reppen, M., & Gantner, R. N. (2018). Are bitcoin bubbles predictable? Combining a generalized Metcalfe's law and the LPPLS model (March 15, 2018). *Swiss Finance Institute Research Paper No. 18–22*.

Yan, Y., Tong, Y., & Wang, Y. (2023, November 27). *Momentum, bubble and information transmission*. SSRN.

Zhang, Q., Zhang, Q., & Sornette, D. (2016). Early warning signals of financial crises with multi-scale quantile regressions of log-periodic power law singularities. *PloS ONE*, *11*(11), e0165819.

Zhou, W.-X., & Sornette, D. (2006). Is there a real estate bubble in the US? *Physica A: Statistical Mechanics and Its Applications*, *361*(1), 297–308.

Figures and tables

Figure 1. Time series graph

This graph shows the time series of the natural logarithm of compounded returns computed from the daily observation of 10 portfolios formed in a momentum strategy obtained from the data library of Kenneth R. French. The two plots illustrate the time series fluctuations of the natural logarithm of compounded returns, divided into two subsamples: November 3, 1926, to December 31, 1974, and January 1, 1975, to November 30, 2023. The peak values of the sixteen bubbles listed in Table 1 are marked with red lines.

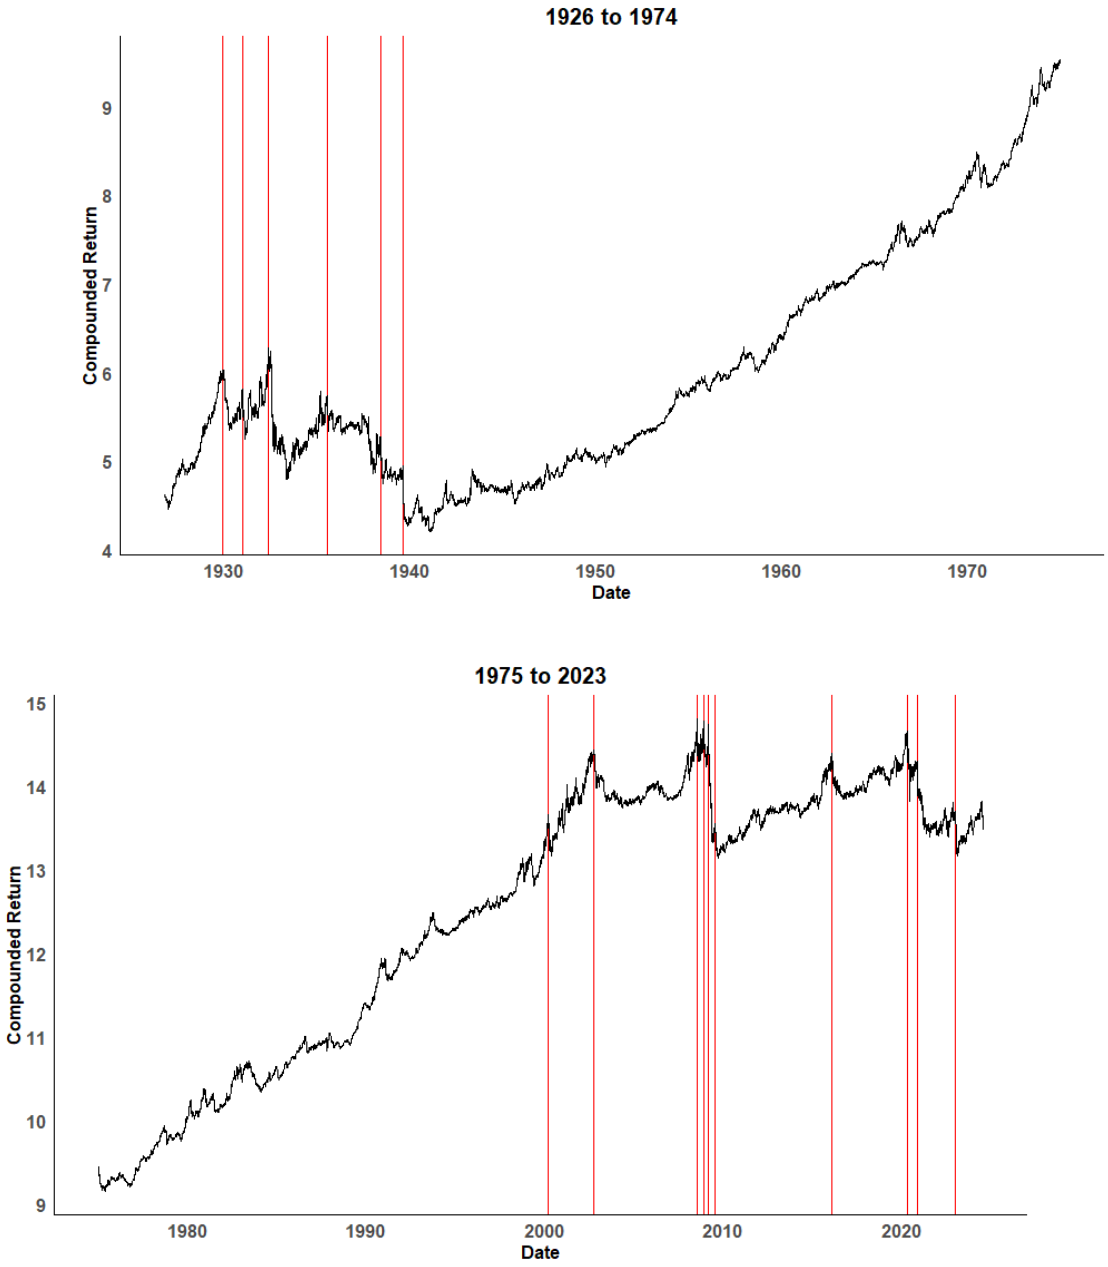

Table 1. Worst daily momentum returns

This table presents the 16 largest drawdowns of daily returns, sorted by the decline in logarithmic compounded return (CR) of daily WML, from November 3, 1926, to November 30, 2023. For each bubble, the second and third columns list the corresponding peak time and crash end time, respectively. The change in CR (Δ_{CR}) reflects the decrease in logarithmic compounded return from the peak to the end of the crash, while Δ_{WML} represents the decrease in WML over the same period. The Duration (days) column shows the time span between the peak and the end of the crash.

Rank	**Peak$_t$**	**Crash end$_t$**	**Δ_{CR}**	**Δ_{WML}**	**Duration(days)**
1	2009.03.06	2009.05.08	-1.40	-11.41	45
2	1932.06.02	1932.09.01	-1.18	-12.68	77
3	2020.11.06	2021.07.19	-0.9	-2.23	174
4	2020.04.16	2020.06.08	-0.85	-19.34	37
5	1929.12.21	1930.04.21	-0.68	-5.02	98
6	1939.08.23	1939.11.29	-0.67	-4.11	80
7	1930.12.29	1931.02.28	-0.57	-4.43	51
8	2008.11.20	2009.01.06	-0.54	-14.25	31
9	2022.12.27	2023.02.15	-0.54	-7.06	35
10	1938.06.17	1938.07.23	-0.53	-5.22	31
11	2000.03.10	2000.05.23	-0.49	-5.48	52
12	2016.01.20	2016.04.27	-0.49	-2.89	69
13	2002.09.30	2002.12.02	-0.47	-0.81	45
14	2008.07.14	2008.07.23	-0.44	-14.08	7
15	2009.07.08	2009.09.22	-0.41	-3.77	54
16	1935.07.20	1935.08.17	-0.40	-3.03	25

Table 2. Calibrating the power-law model using 3 months of data before the momentum bubble peak on December 21, 1929.

Panel A shows the initial value of the first-stage estimation approach $\Phi^0_{model\ 1} = (A, B, t_c, \beta)$, which is a basic power-law model, as $ln[comp(t)] = Ln[com(T)] - (t_c - t)^{\overline{\beta}}$. In this setting, $A = ln[p(T)]$ is the natural logarithm of the last compounded return in the sample. Given the daily frequency of the sample, one unit of time equals 1/253. In the second stage, the model is optimized by minimizing the *SSR* of $ln[comp(t)] = Ln[com(T)] - (t_c - t)^{\overline{\beta}}$ by allowing for the following constraints: $t_c \geq T + 1\ and\ 0.1 \leq \beta \leq 0.9$. Panel B reports the optimal values for the parameter vector $\Phi^*_{model\ 1} = (A^*, B^*, t_c^*, \beta^*)$. In the third stage, using the optimized parameter values of Model 1 in Panel B as the initial values $\Phi^0_{model\ 2} = (A^*, B^*, t_c^*, \beta^*, C, \omega, \Phi)$ and setting $C = \phi = 0$ for varying values of $5 \leq \omega \leq 15$, we calibrate the LPPLS model 2 that includes log-periodic oscillations as $log[E(p(t))] \approx A + B(t_c - t)^{\beta} - (1 + C\ cos[\omega\ log(t_c - t) + \varphi])$ by considering the following constraints: $t_c \geq T + 1,\ 0.1 \leq \beta \leq 0.9, and\ 5 \leq \omega \leq 15$. Panel C outlines the input data vectors $\Phi^0_{model\ 2} = (A^*, B^*, t_c^*, \beta^*, C, \omega, \Phi)$, while Panel D presents the optimized parameter vectors $\Phi^*_{model\ 2} = (A^{**}, B^{**}, t_c^{**}, \beta^{**}, C^{**}, \omega^{**}, \Phi^{**})$ for each iteration.

Panel A. Initial parameter values for Model 1					
Spec	**A**	**B**	**T_c=T+1**	**b**	***SSR***
1	5.85	-1.00	3.41	0.20	**93.20***
2	5.85	-1.00	3.41	0.40	124.37
3	5.85	-1.00	3.41	0.60	256.67
4	5.85	-1.00	3.41	0.80	543.29
Panel B. Optimized parameter values for Model 1					
Spec	**A***	**B***	**t_c^***	**b***	***SSR***
1	6.07	-0.76	3.52	0.55	**4.20***
2	6.12	-0.80	3.56	0.53	4.20
3	6.09	-0.78	3.54	0.54	4.20
4	6.09	-0.77	3.54	0.55	4.20

Panel C. Initial parameter values for Model 2											
Spec	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**	**11**
A*	6.07	6.07	6.07	6.07	6.07	6.07	6.07	6.07	6.07	6.07	6.07
B*	-0.76	-0.76	-0.76	-0.76	-0.76	-0.76	-0.76	-0.76	-0.76	-0.76	-0.76
t_c*	3.52	3.52	3.52	3.52	3.52	3.52	3.52	3.52	3.52	3.52	3.52
b*	0.55	0.55	0.55	0.55	0.55	0.55	0.55	0.55	0.55	0.55	0.55
C*	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
ω*	5.00	6.00	7.00	8.00	9.00	10.00	11.00	12.00	13.00	14.00	15.00
Φ*	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
SSR	4.20	4.20	4.20	4.20	4.20	4.20	4.20	4.20	4.20	4.20	4.20
Panel D. Optimized parameter values for Model 2											
Spec	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**	**11**
A*	5.87	7.59	7.98	7.64	7.08	6.88	8.48	7.85	5.92	5.92	6.14
B*	-0.59	-1.91	-2.26	-1.96	-1.47	-1.30	-2.72	-2.15	-0.67	-0.67	-0.84
t_c*	3.44	4.23	4.28	4.22	4.12	4.06	4.34	4.27	3.41	3.41	3.53
b*	0.71	0.33	0.29	0.32	0.39	0.43	0.25	0.30	0.57	0.57	0.50
C*	-0.11	-0.03	-0.03	-0.03	0.04	0.05	0.02	-0.03	-0.06	-0.06	0.05
ω*	5.62	8.97	9.21	8.95	8.47	8.19	9.47	9.14	13.77	13.85	14.54
φ*	1.04	-4.32	-4.74	-4.29	-0.33	0.13	-2.05	1.67	0.21	0.13	1.90
SSR	1.76	1.59	1.59	1.59	1.61	1.61	**1.58***	1.59	3.04	3.04	3.07

Table 3. Calibrating the power-law model using 6 months of data before the momentum bubble peak on December 21, 1929.

Panel A shows the initial value of the first-stage estimation approach $\Phi^0_{model\ 1} = (A, B, t_c, \beta)$, which is a basic power-law model, as $ln[comp(t)] = Ln[com(T)] - (t_c - t)^{\overline{\beta}}$. In this setting, $A = ln[p(T)]$ is the natural logarithm of the last compounded return in the sample. Given the daily frequency of the sample, one unit of time equals 1/253. In the second stage, the model is optimized by minimizing the *SSR* of $ln[comp(t)] = Ln[com(T)] - (t_c - t)^{\overline{\beta}}$ by allowing for the following constraints: $t_c \geq T + 1\ and\ 0.1 \leq \beta \leq 0.9$. Panel B reports the optimal values for the parameter vector $\Phi^*_{model\ 1} = (A^*, B^*, t_c^*, \beta^*)$. In the third stage, using the optimized parameter values of Model 1 in Panel B as the initial values $\Phi^0_{model\ 2} = (A^*, B^*, t_c^*, \beta^*, C, \omega, \Phi)$ and setting $C = \phi = 0$ for varying values of $5 \leq \omega \leq 15$, we calibrate the LPPLS model 2 that includes log-periodic oscillations as $log[E(p(t))] \approx A + B(t_c - t)^{\beta} - (1 + C\ cos[\omega\ log(t_c - t) + \varphi])$ by considering the following constraints: $t_c \geq T + 1,\ 0.1 \leq \beta \leq 0.9, and\ 5 \leq \omega \leq 15$. Panel C outlines the input data vectors $\Phi^0_{model\ 2} = (A^*, B^*, t_c^*, \beta^*, C, \omega, \Phi)$, while Panel D presents the optimized parameter vectors $\Phi^*_{model\ 2} = (A^{**}, B^{**}, t_c^{**}, \beta^{**}, C^{**}, \omega^{**}, \Phi^{**})$ for each iteration.

Panel A. Initial parameter values for Model 1					
Spec	**A**	**B**	**T_c=T+1**	**b**	***SSR***
1	5.59	-1.00	3.11	0.20	**182.44***
2	5.59	-1.00	3.11	0.40	233.17
3	5.59	-1.00	3.11	0.60	365.72
4	5.59	-1.00	3.11	0.80	611.96
Panel B. Optimized parameter values for Model 1					
Spec	**A***	**B***	**t_c^***	**b***	***SSR***
1	5.94	-0.70	3.35	0.56	**4.14***
2	5.89	-0.67	3.30	0.57	4.15
3	5.81	-0.62	3.24	0.60	4.16
4	5.85	-0.65	3.28	0.59	4.15

Panel C. Initial parameter values for Model 2											
Spec	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**	**11**
A*	5.94	5.94	5.94	5.94	5.94	5.94	5.94	5.94	5.94	5.94	5.94
B*	-0.70	-0.70	-0.70	-0.70	-0.70	-0.70	-0.70	-0.70	-0.70	-0.70	-0.70
t_c*	3.35	3.35	3.35	3.35	3.35	3.35	3.35	3.35	3.35	3.35	3.35
b*	0.56	0.56	0.56	0.56	0.56	0.56	0.56	0.56	0.56	0.56	0.56
C*	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
ω*	5.00	6.00	7.00	8.00	9.00	10.00	11.00	12.00	13.00	14.00	15.00
Φ*	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
SSR	4.14	4.14	4.14	4.14	4.14	4.14	4.14	4.14	4.14	4.14	4.14
Panel D. Optimized parameter values for Model 2											
Spec	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**	**11**
A*	9.02	7.42	7.58	8.44	7.60	8.75	5.84	5.84	5.84	5.84	5.84
B*	-3.43	-1.91	-2.06	-2.87	-2.08	-3.17	-0.67	-0.67	-0.67	-0.67	-0.67
t_c*	4.03	3.91	3.92	4.00	3.92	4.01	3.20	3.20	3.20	3.20	3.20
b*	0.19	0.30	0.28	0.22	0.28	0.20	0.54	0.54	0.54	0.54	0.54
C*	0.02	-0.04	-0.03	-0.03	-0.03	0.02	0.07	0.07	-0.07	-0.07	-0.07
ω*	8.69	8.08	8.10	8.55	8.12	8.61	11.55	11.55	11.55	11.55	11.55
φ*	-6.45	-2.37	-2.39	-3.10	-2.41	-0.03	0.23	0.23	3.37	3.37	3.37
SSR	**1.44***	1.47	1.46	1.45	1.46	1.45	2.76	2.76	2.76	2.76	2.76

Table 4. Calibrating the power-law model using 9 months of data before the momentum bubble peak on December 21, 1929.

Panel A shows the initial value of the first-stage estimation approach $\Phi^0_{model\ 1} = (A, B, t_c, \beta)$, which is a basic power-law model, as $ln[comp(t)] = Ln[com(T)] - (t_c - t)^{\overline{\beta}}$. In this setting, $A = ln[p(T)]$ is the natural logarithm of the last compounded return in the sample. Given the daily frequency of the sample, one unit of time equals 1/253. In the second stage, the model is optimized by minimizing the *SSR* of $ln[comp(t)] = Ln[com(T)] - (t_c - t)^{\overline{\beta}}$ by allowing for the following constraints: $t_c \geq T + 1\ and\ 0.1 \leq \beta \leq 0.9$. Panel B reports the optimal values for the parameter vector $\Phi^*_{model\ 1} = (A^*, B^*, t_c^*, \beta^*)$. In the third stage, using the optimized parameter values of Model 1 in Panel B as the initial values $\Phi^0_{model\ 2} = (A^*, B^*, t_c^*, \beta^*, C, \omega, \Phi)$ and setting $C = \phi = 0$ for varying values of $5 \leq \omega \leq 15$, we calibrate the LPPLS model 2 that includes log-periodic oscillations as $log[E(p(t))] \approx A + B(t_c - t)^{\beta} - (1 + C\ cos[\omega\ log(t_c - t) + \varphi])$ by considering the following constraints: $t_c \geq T + 1,\ 0.1 \leq \beta \leq 0.9, and\ 5 \leq \omega \leq 15$. Panel C outlines the input data vectors $\Phi^0_{model\ 2} = (A^*, B^*, t_c^*, \beta^*, C, \omega, \Phi)$, while Panel D presents the optimized parameter vectors $\Phi^*_{model\ 2} = (A^{**}, B^{**}, t_c^{**}, \beta^{**}, C^{**}, \omega^{**}, \Phi^{**})$ for each iteration.

Panel A. Initial parameter values for Model 1

Spec	**A**	**B**	**T_c=T+1**	**b**	***SSR***
1	5.50	-1.00	2.81	0.20	**165.33***
2	5.50	-1.00	2.81	0.40	201.65
3	5.50	-1.00	2.81	0.60	296.13
4	5.50	-1.00	2.81	0.80	464.57

Panel B. Optimized parameter values for Model 1

Spec	**A***	**B***	**t_c^***	**b***	***SSR***
1	5.58	-0.55	2.81	0.59	**3.79***
2	5.62	-0.58	2.84	0.57	3.80
3	5.58	-0.55	2.81	0.59	3.79
4	5.62	-0.58	2.83	0.56	3.80

Panel C. Initial parameter values for Model 2

Spec	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**	**11**
A*	5.58	5.58	5.58	5.58	5.58	5.58	5.58	5.58	5.58	5.58	5.58
B*	-0.55	-0.55	-0.55	-0.55	-0.55	-0.55	-0.55	-0.55	-0.55	-0.55	-0.55
t_c*	2.81	2.81	2.81	2.81	2.81	2.81	2.81	2.81	2.81	2.81	2.81
b*	0.59	0.59	0.59	0.59	0.59	0.59	0.59	0.59	0.59	0.59	0.59
C*	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
ω*	5.00	6.00	7.00	8.00	9.00	10.00	11.00	12.00	13.00	14.00	15.00
Φ*	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
SSR	3.79	3.79	3.79	3.79	3.79	3.79	3.79	3.79	3.79	3.79	3.79

Panel D. Optimized parameter values for Model 2

Spec	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**	**11**
A*	6.17	6.26	6.23	5.64	5.68	5.68	5.65	5.64	5.68	5.68	5.68
B*	-1.07	-1.16	-1.13	-0.62	-0.65	-0.65	-0.62	-0.61	-0.65	-0.65	-0.65
t_c*	3.04	3.05	3.05	2.81	2.83	2.83	2.81	2.81	2.83	2.83	2.83
b*	0.37	0.35	0.35	0.50	0.48	0.48	0.49	0.50	0.48	0.48	0.48
C*	0.08	0.07	0.07	0.10	-0.09	-0.09	-0.10	-0.10	0.09	0.09	0.09
ω*	5.29	5.37	5.34	7.80	8.36	8.36	8.31	8.34	8.36	8.36	8.36
φ*	-0.46	-0.53	-0.51	-1.25	1.41	1.41	1.50	1.50	4.56	4.56	4.56
SSR	1.35	**1.34***	1.35	2.13	2.11	2.11	2.11	2.11	2.11	2.11	2.11

Table 5. Calibrating the power-law model using 12 months of data before the momentum bubble peak on December 21, 1929.

Panel A shows the initial value of the first-stage estimation approach $\Phi^0_{model\ 1} = (A, B, t_c, \beta)$, which is a basic power-law model, as $ln[comp(t)] = Ln[com(T)] - (t_c - t)^{\overline{\beta}}$. In this setting, $A = ln[p(T)]$ is the natural logarithm of the last compounded return in the sample. Given the daily frequency of the sample, one unit of time equals 1/253. In the second stage, the model is optimized by minimizing the *SSR* of $ln[comp(t)] = Ln[com(T)] - (t_c - t)^{\overline{\beta}}$ by allowing for the following constraints: $t_c \geq T + 1\ and\ 0.1 \leq \beta \leq 0.9$. Panel B reports the optimal values for the parameter vector $\Phi^*_{model\ 1} = (A^*, B^*, t_c^*, \beta^*)$. In the third stage, using the optimized parameter values of Model 1 in Panel B as the initial values $\Phi^0_{model\ 2} = (A^*, B^*, t_c^*, \beta^*, C, \omega, \Phi)$ and setting $C = \phi = 0$ for varying values of $5 \leq \omega \leq 15$, we calibrate the LPPLS model 2 that includes log-periodic oscillations as $log[E(p(t))] \approx A + B(t_c - t)^{\beta} - (1 + C\ cos[\omega\ log(t_c - t) + \varphi])$ by considering the following constraints: $t_c \geq T + 1,\ 0.1 \leq \beta \leq 0.9, and\ 5 \leq \omega \leq 15$. Panel C outlines the input data vectors $\Phi^0_{model\ 2} = (A^*, B^*, t_c^*, \beta^*, C, \omega, \Phi)$, while Panel D presents the optimized parameter vectors $\Phi^*_{model\ 2} = (A^{**}, B^{**}, t_c^{**}, \beta^{**}, C^{**}, \omega^{**}, \Phi^{**})$ for each iteration.

Panel A. Initial parameter values for Model 1					
Spec	**A**	**B**	**T_c=T+1**	**b**	***SSR***
1	5.39	-1.00	2.53	0.20	**168.70***
2	5.39	-1.00	2.53	0.40	194.62
3	5.39	-1.00	2.53	0.60	261.71
4	5.39	-1.00	2.53	0.80	375.70
Panel B. Optimized parameter values for Model 1					
Spec	**A***	**B***	**t_c^***	**b***	***SSR***
1	5.57	-0.54	2.80	0.59	3.71
2	5.48	-0.48	2.71	0.64	3.68
3	5.35	-0.41	2.53	0.70	**3.59***
4	5.51	-0.53	2.69	0.58	3.68

Panel C. Initial parameter values for Model 2											
Spec	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**	**11**
A*	5.35	5.35	5.35	5.35	5.35	5.35	5.35	5.35	5.35	5.35	5.35
B*	-0.41	-0.41	-0.41	-0.41	-0.41	-0.41	-0.41	-0.41	-0.41	-0.41	-0.41
t_c*	2.53	2.53	2.53	2.53	2.53	2.53	2.53	2.53	2.53	2.53	2.53
b*	0.70	0.70	0.70	0.70	0.70	0.70	0.70	0.70	0.70	0.70	0.70
C*	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
ω*	5.00	6.00	7.00	8.00	9.00	10.00	11.00	12.00	13.00	14.00	15.00
Φ*	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
SSR	3.59	3.59	3.59	3.59	3.59	3.59	3.59	3.59	3.59	3.59	3.59
Panel D. Optimized parameter values for Model 2											
Spec	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**	**11**
A	8.47	5.45	5.45	5.45	5.45	5.45	5.45	5.45	5.45	5.36	5.36
B	-2.46	-0.52	-0.52	-0.52	-0.52	-0.52	-0.52	-0.52	-0.52	-0.42	-0.42
t_c*	5.15	2.53	2.53	2.53	2.53	2.53	2.53	2.53	2.53	2.53	2.53
b*	0.28	0.49	0.49	0.49	0.49	0.49	0.49	0.49	0.49	0.67	0.67
C*	0.03	0.14	0.14	0.14	0.14	0.14	-0.14	-0.14	-0.14	0.07	0.07
ω*	15.00	5.93	5.93	5.93	5.93	5.93	5.93	5.93	5.93	15.00	15.00
φ*	-18.05	1.02	1.02	1.02	1.02	1.02	4.16	4.16	4.16	0.48	0.48
SSR	**1.08***	1.61	1.61	1.61	1.61	1.61	1.61	1.61	1.61	3.27	3.27

Table 6. Calibrating the power-law model using 3 months of data before the momentum bubble peak on March 10, 2000.

Panel A shows the initial value of the first-stage estimation approach $\Phi^0_{model\ 1} = (A, B, t_c, \beta)$, which is a basic power-law model, as $ln[comp(t)] = Ln[com(T)] - (t_c - t)^{\overline{\beta}}$. In this setting, $A = ln[p(T)]$ is the natural logarithm of the last compounded return in the sample. Given the daily frequency of the sample, one unit of time equals 1/253. In the second stage, the model is optimized by minimizing the *SSR* of $ln[comp(t)] = Ln[com(T)] - (t_c - t)^{\overline{\beta}}$ by allowing for the following constraints: $t_c \geq T + 1\ and\ 0.1 \leq \beta \leq 0.9$. Panel B reports the optimal values for the parameter vector $\Phi^*_{model\ 1} = (A^*, B^*, t_c^*, \beta^*)$. In the third stage, using the optimized parameter values of Model 1 in Panel B as the initial values $\Phi^0_{model\ 2} = (A^*, B^*, t_c^*, \beta^*, C, \omega, \Phi)$ and setting $C = \phi = 0$ for varying values of $5 \leq \omega \leq 15$, we calibrate the LPPLS model 2 that includes log-periodic oscillations as $log[E(p(t))] \approx A + B(t_c - t)^{\beta} - (1 + C\ cos[\omega\ log(t_c - t) + \varphi])$ by considering the following constraints: $t_c \geq T + 1,\ 0.1 \leq \beta \leq 0.9, and\ 5 \leq \omega \leq 15$. Panel C outlines the input data vectors $\Phi^0_{model\ 2} = (A^*, B^*, t_c^*, \beta^*, C, \omega, \Phi)$, while Panel D presents the optimized parameter vectors $\Phi^*_{model\ 2} = (A^{**}, B^{**}, t_c^{**}, \beta^{**}, C^{**}, \omega^{**}, \Phi^{**})$ for each iteration.

Panel A. Initial parameter values for Model 1					
Spec	**A**	**B**	**T_c=T+1**	**b**	***SSR***
1	13.28	-1.00	61.67	0.20	253132.31
2	13.28	-1.00	61.67	0.40	**74337.24***
3	13.28	-1.00	61.67	0.60	86989.35
4	13.28	-1.00	61.67	0.80	1885434.69
Panel B. Optimized parameter values for Model 1					
Spec	**A***	**B***	**t_c^***	**b***	***SSR***
1	13.56	-0.32	61.76	0.83	1149.71
2	13.53	-0.31	61.70	0.84	1151.11
3	14.82	-0.40	66.82	0.79	**1096.97***
4	13.60	-0.34	61.67	0.82	1159.64

Panel C. Initial parameter values for Model 2											
Spec	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**	**11**
A*	14.82	14.82	14.82	14.82	14.82	14.82	14.82	14.82	14.82	14.82	14.82
B*	-0.40	-0.40	-0.40	-0.40	-0.40	-0.40	-0.40	-0.40	-0.40	-0.40	-0.40
t_c*	66.82	66.82	66.82	66.82	66.82	66.82	66.82	66.82	66.82	66.82	66.82
b*	0.79	0.79	0.79	0.79	0.79	0.79	0.79	0.79	0.79	0.79	0.79
C*	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
ω*	5.00	6.00	7.00	8.00	9.00	10.00	11.00	12.00	13.00	14.00	15.00
Φ*	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
SSR	1096.97	1096.97	1096.97	1096.97	1096.97	1096.97	1096.97	1096.97	1096.97	1096.97	1096.97
Panel D. Optimized parameter values for Model 2											
Spec	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**	**11**
A*	13.43	13.43	13.41	13.43	13.42	13.42	13.46	13.41	13.42	14.34	14.34
B*	-0.25	-0.25	-0.25	-0.25	-0.25	-0.25	-0.25	-0.25	-0.25	-0.34	-0.34
t_c*	62.23	62.23	62.16	62.23	62.22	62.19	62.37	62.16	62.23	65.32	65.32
b*	0.88	0.88	0.88	0.88	0.88	0.89	0.88	0.88	0.89	0.83	0.83
C*	0.07	-0.07	0.07	0.07	-0.07	0.07	0.07	-0.07	0.07	0.03	-0.03
ω*	5.04	5.04	5.00	5.04	5.03	5.02	5.00	5.00	5.04	15.00	15.00
φ*	0.31	3.46	6.76	12.89	16.05	19.24	25.58	28.75	31.75	-2.32	0.82
SSR	**302.26***	302.26	302.33	302.26	302.26	302.28	303.05	302.33	302.26	822.40	822.40

Table 7. Calibrating the power-law model using 6 months of data before the momentum bubble peak on March 10, 2000.

Panel A shows the initial value of the first-stage estimation approach $\Phi^0_{model\ 1} = (A, B, t_c, \beta)$, which is a basic power-law model, as $ln[comp(t)] = Ln[com(T)] - (t_c - t)^{\overline{\beta}}$. In this setting, $A = ln[p(T)]$ is the natural logarithm of the last compounded return in the sample. Given the daily frequency of the sample, one unit of time equals 1/253. In the second stage, the model is optimized by minimizing the *SSR* of $ln[comp(t)] = Ln[com(T)] - (t_c - t)^{\overline{\beta}}$ by allowing for the following constraints: $t_c \geq T + 1\ and\ 0.1 \leq \beta \leq 0.9$. Panel B reports the optimal values for the parameter vector $\Phi^*_{model\ 1} = (A^*, B^*, t_c^*, \beta^*)$. In the third stage, using the optimized parameter values of Model 1 in Panel B as the initial values $\Phi^0_{model\ 2} = (A^*, B^*, t_c^*, \beta^*, C, \omega, \Phi)$ and setting $C = \phi = 0$ for varying values of $5 \leq \omega \leq 15$, we calibrate the LPPLS model 2 that includes log-periodic oscillations as $log[E(p(t))] \approx A + B(t_c - t)^{\beta} - (1 + C\ cos[\omega\ log(t_c - t) + \varphi])$ by considering the following constraints: $t_c \geq T + 1,\ 0.1 \leq \beta \leq 0.9, and\ 5 \leq \omega \leq 15$. Panel C outlines the input data vectors $\Phi^0_{model\ 2} = (A^*, B^*, t_c^*, \beta^*, C, \omega, \Phi)$, while Panel D presents the optimized parameter vectors $\Phi^*_{model\ 2} = (A^{**}, B^{**}, t_c^{**}, \beta^{**}, C^{**}, \omega^{**}, \Phi^{**})$ for each iteration.

Panel A. Initial parameter values for Model 1

Spec	A	B	T_c=T+1	b	*SSR*
1	13.03	-1.00	61.42	0.20	229605.99
2	13.03	-1.00	61.42	0.40	**64747.43***
3	13.03	-1.00	61.42	0.60	102313.65
4	13.03	-1.00	61.42	0.80	1931376.86

Panel B. Optimized parameter values for Model 1

Spec	A*	B*	t_c^*	b*	*SSR*
1	13.50	-0.32	61.42	0.83	1149.82
2	13.48	-0.31	61.45	0.84	**1148.81***
3	13.48	-0.31	61.42	0.84	1149.20
4	13.62	-0.35	61.42	0.81	1168.59

Panel C. Initial parameter values for Model 2

Spec	1	2	3	4	5	6	7	8	9	10	11
A*	13.48	13.48	13.48	13.48	13.48	13.48	13.48	13.48	13.48	13.48	13.48
B*	-0.31	-0.31	-0.31	-0.31	-0.31	-0.31	-0.31	-0.31	-0.31	-0.31	-0.31
t_c*	61.45	61.45	61.45	61.45	61.45	61.45	61.45	61.45	61.45	61.45	61.45
b*	0.84	0.84	0.84	0.84	0.84	0.84	0.84	0.84	0.84	0.84	0.84
C*	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
ω*	5.00	6.00	7.00	8.00	9.00	10.00	11.00	12.00	13.00	14.00	15.00
Φ*	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
SSR	1148.81	1148.81	1148.81	1148.81	1148.81	1148.81	1148.81	1148.81	1148.81	1148.81	1148.81

Panel D. Optimized parameter values for Model 2

Spec	1	2	3	4	5	6	7	8	9	10	11
A*	13.42	13.42	13.41	15.00	13.41	13.42	13.42	13.42	13.73	13.99	13.37
B*	-0.25	-0.25	-0.25	-0.31	-0.25	-0.25	-0.25	-0.25	-0.30	-0.31	-0.27
t_c**	62.15	62.14	62.11	69.41	62.16	62.14	62.12	62.12	62.85	63.86	61.43
b*	0.88	0.88	0.88	0.85	0.88	0.88	0.88	0.88	0.85	0.84	0.87
C*	0.07	-0.07	0.07	-0.05	-0.07	0.07	-0.07	-0.07	0.03	-0.03	0.03
ω**	5.01	5.01	5.00	7.15	5.03	5.01	5.00	5.00	14.52	15.00	13.70
φ**	0.42	3.58	6.77	0.75	16.07	19.29	22.47	28.75	-5.99	-5.05	3.95
SSR	301.09	**301.08***	301.13	361.29	301.14	301.09	301.10	301.10	821.95	819.30	829.68

Table 8. Calibrating the power-law model using 9 months of data before the momentum bubble peak on March 10, 2000.

Panel A shows the initial value of the first-stage estimation approach $\Phi^0_{model\ 1} = (A, B, t_c, \beta)$, which is a basic power-law model, as $ln[comp(t)] = Ln[com(T)] - (t_c - t)^{\overline{\beta}}$. In this setting, $A = ln[p(T)]$ is the natural logarithm of the last compounded return in the sample. Given the daily frequency of the sample, one unit of time equals 1/253. In the second stage, the model is optimized by minimizing the *SSR* of $ln[comp(t)] = Ln[com(T)] - (t_c - t)^{\overline{\beta}}$ by allowing for the following constraints: $t_c \geq T + 1\ and\ 0.1 \leq \beta \leq 0.9$. Panel B reports the optimal values for the parameter vector $\Phi^*_{model\ 1} = (A^*, B^*, t_c^*, \beta^*)$. In the third stage, using the optimized parameter values of Model 1 in Panel B as the initial values $\Phi^0_{model\ 2} = (A^*, B^*, t_c^*, \beta^*, C, \omega, \Phi)$ and setting $C = \phi = 0$ for varying values of $5 \leq \omega \leq 15$, we calibrate the LPPLS model 2 that includes log-periodic oscillations as $log[E(p(t))] \approx A + B(t_c - t)^{\beta} - (1 + C\ cos[\omega\ log(t_c - t) + \varphi])$ by considering the following constraints: $t_c \geq T + 1,\ 0.1 \leq \beta \leq 0.9, and\ 5 \leq \omega \leq 15$. Panel C outlines the input data vectors $\Phi^0_{model\ 2} = (A^*, B^*, t_c^*, \beta^*, C, \omega, \Phi)$, while Panel D presents the optimized parameter vectors $\Phi^*_{model\ 2} = (A^{**}, B^{**}, t_c^{**}, \beta^{**}, C^{**}, \omega^{**}, \Phi^{**})$ for each iteration.

Panel A. Initial parameter values for Model 1

Spec	A	B	T_c=T+1	b	*SSR*
1	12.85	-1.00	61.16	0.20	213569.77
2	12.85	-1.00	61.16	0.40	**58757.78***
3	12.85	-1.00	61.16	0.60	113170.65
4	12.85	-1.00	61.16	0.80	1955527.20

Panel B. Optimized parameter values for Model 1

Spec	A*	B*	t_c^*	b*	*SSR*
1	13.49	-0.33	61.16	0.82	1140.31
2	13.46	-0.32	61.22	0.83	**1137.28***
3	13.45	-0.32	61.16	0.83	1138.78
4	13.53	-0.34	61.16	0.82	1145.25

Panel C. Initial parameter values for Model 2

Spec	1	2	3	4	5	6	7	8	9	10	11
A*	13.46	13.46	13.46	13.46	13.46	13.46	13.46	13.46	13.46	13.46	13.46
B*	-0.32	-0.32	-0.32	-0.32	-0.32	-0.32	-0.32	-0.32	-0.32	-0.32	-0.32
t_c*	61.22	61.22	61.22	61.22	61.22	61.22	61.22	61.22	61.22	61.22	61.22
b*	0.83	0.83	0.83	0.83	0.83	0.83	0.83	0.83	0.83	0.83	0.83
C*	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
ω*	5.00	6.00	7.00	8.00	9.00	10.00	11.00	12.00	13.00	14.00	15.00
Φ*	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
SSR	1137.28	1137.28	1137.28	1137.28	1137.28	1137.28	1137.28	1137.28	1137.28	1137.28	1137.28

Panel D. Optimized parameter values for Model 2

Spec	1	2	3	4	5	6	7	8	9	10	11
A*	13.42	13.43	13.42	13.42	13.20	13.20	13.42	13.97	13.97	13.97	13.97
B*	-0.26	-0.26	-0.26	-0.26	-0.24	-0.24	-0.26	-0.32	-0.32	-0.32	-0.32
t_c**	62.04	62.05	62.04	62.03	61.17	61.16	62.04	63.63	63.63	63.63	63.63
b*	0.88	0.88	0.88	0.88	0.90	0.90	0.88	0.84	0.84	0.84	0.84
C*	0.07	-0.07	0.07	-0.07	-0.07	0.07	-0.07	-0.03	0.03	-0.03	0.03
ω**	5.00	5.00	5.00	5.00	5.00	5.13	5.00	14.88	14.88	14.88	14.88
φ **	0.48	3.62	6.77	9.91	16.35	19.01	22.47	-10.78	-7.63	-4.49	-1.35
SSR	296.64	296.64	**296.63***	296.67	299.87	306.01	296.64	811.36	811.36	811.36	811.36

Table 9. Calibrating the power-law model using 12 months of data before the momentum bubble peak on March 10, 2000.

Panel A shows the initial value of the first-stage estimation approach $\Phi^0_{model\ 1} = (A, B, t_c, \beta)$, which is a basic power-law model, as $ln[comp(t)] = Ln[com(T)] - (t_c - t)^{\overline{\beta}}$. In this setting, $A = ln[p(T)]$ is the natural logarithm of the last compounded return in the sample. Given the daily frequency of the sample, one unit of time equals 1/253. In the second stage, the model is optimized by minimizing the *SSR* of $ln[comp(t)] = Ln[com(T)] - (t_c - t)^{\overline{\beta}}$ by allowing for the following constraints: $t_c \geq T + 1\ and\ 0.1 \leq \beta \leq 0.9$. Panel B reports the optimal values for the parameter vector $\Phi^*_{model\ 1} = (A^*, B^*, t_c^*, \beta^*)$. In the third stage, using the optimized parameter values of Model 1 in Panel B as the initial values $\Phi^0_{model\ 2} = (A^*, B^*, t_c^*, \beta^*, C, \omega, \Phi)$ and setting $C = \phi = 0$ for varying values of $5 \leq \omega \leq 15$, we calibrate the LPPLS model 2 that includes log-periodic oscillations as $log[E(p(t))] \approx A + B(t_c - t)^{\beta} - (1 + C\ cos[\omega\ log(t_c - t) + \varphi])$ by considering the following constraints: $t_c \geq T + 1,\ 0.1 \leq \beta \leq 0.9, and\ 5 \leq \omega \leq 15$. Panel C outlines the input data vectors $\Phi^0_{model\ 2} = (A^*, B^*, t_c^*, \beta^*, C, \omega, \Phi)$, while Panel D presents the optimized parameter vectors $\Phi^*_{model\ 2} = (A^{**}, B^{**}, t_c^{**}, \beta^{**}, C^{**}, \omega^{**}, \Phi^{**})$ for each iteration.

Panel A. Initial parameter values for Model 1

Spec	A	B	T_c=T+1	b	*SSR*
1	13.13	-1.00	60.91	0.20	239350.27
2	13.13	-1.00	60.91	0.40	**68944.56***
3	13.13	-1.00	60.91	0.60	88199.61
4	13.13	-1.00	60.91	0.80	1837956.02

Panel B. Optimized parameter values for Model 1

Spec	A*	B*	t_c^*	b*	*SSR*
1	13.52	-0.35	60.91	0.81	1145.21
2	13.43	-0.32	60.98	0.83	**1130.59***
3	13.41	-0.32	60.91	0.83	1132.29
4	13.52	-0.36	60.91	0.80	1150.02

Panel C. Initial parameter values for Model 2

Spec	1	2	3	4	5	6	7	8	9	10	11
A*	13.43	13.43	13.43	13.43	13.43	13.43	13.43	13.43	13.43	13.43	13.43
B*	-0.32	-0.32	-0.32	-0.32	-0.32	-0.32	-0.32	-0.32	-0.32	-0.32	-0.32
t_c*	60.98	60.98	60.98	60.98	60.98	60.98	60.98	60.98	60.98	60.98	60.98
b*	0.83	0.83	0.83	0.83	0.83	0.83	0.83	0.83	0.83	0.83	0.83
C*	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
ω*	5.00	6.00	7.00	8.00	9.00	10.00	11.00	12.00	13.00	14.00	15.00
Φ*	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
SSR	1130.59	1130.59	1130.59	1130.59	1130.59	1130.59	1130.59	1130.59	1130.59	1130.59	1130.59

Panel D. Optimized parameter values for Model 2

Spec	1	2	3	4	5	6	7	8	9	10	11
A*	13.45	13.43	13.44	13.43	13.48	13.44	13.46	13.95	13.95	13.95	13.95
B*	-0.26	-0.26	-0.26	-0.26	-0.27	-0.26	-0.26	-0.32	-0.32	-0.32	-0.32
t_c**	62.07	62.01	62.04	62.01	62.21	62.04	62.12	63.47	63.47	63.47	63.45
b*	0.87	0.87	0.87	0.87	0.87	0.87	0.87	0.84	0.84	0.84	0.84
C*	0.07	-0.07	0.07	-0.07	-0.07	0.07	-0.07	-0.03	0.03	-0.03	0.03
ω**	5.00	5.00	5.00	5.00	5.04	5.00	5.02	14.80	14.80	14.80	14.79
φ **	0.47	3.62	6.76	9.91	16.02	19.33	22.40	-10.42	-7.28	-4.14	-0.96
SSR	293.93	293.94	**293.92***	293.93	294.14	293.92	294.01	806.34	806.34	806.34	806.34

Table 10. Calibrating the power-law model using 3 months of data before the momentum bubble peak on July 14, 2008.

Panel A shows the initial value of the first-stage estimation approach $\Phi^0_{model\ 1} = (A, B, t_c, \beta)$, which is a basic power-law model, as $ln[comp(t)] = Ln[com(T)] - (t_c - t)^{\overline{\beta}}$. In this setting, $A = ln[p(T)]$ is the natural logarithm of the last compounded return in the sample. Given the daily frequency of the sample, one unit of time equals 1/253. In the second stage, the model is optimized by minimizing the *SSR* of $ln[comp(t)] = Ln[com(T)] - (t_c - t)^{\overline{\beta}}$ by allowing for the following constraints: $t_c \geq T + 1\ and\ 0.1 \leq \beta \leq 0.9$. Panel B reports the optimal values for the parameter vector $\Phi^*_{model\ 1} = (A^*, B^*, t_c^*, \beta^*)$. In the third stage, using the optimized parameter values of Model 1 in Panel B as the initial values $\Phi^0_{model\ 2} = (A^*, B^*, t_c^*, \beta^*, C, \omega, \Phi)$ and setting $C = \phi = 0$ for varying values of $5 \leq \omega \leq 15$, we calibrate the LPPLS model 2 that includes log-periodic oscillations as $log[E(p(t))] \approx A + B(t_c - t)^{\beta} - (1 + C\ cos[\omega\ log(t_c - t) + \varphi])$ by considering the following constraints: $t_c \geq T + 1,\ 0.1 \leq \beta \leq 0.9, and\ 5 \leq \omega \leq 15$. Panel C outlines the input data vectors $\Phi^0_{model\ 2} = (A^*, B^*, t_c^*, \beta^*, C, \omega, \Phi)$, while Panel D presents the optimized parameter vectors $\Phi^*_{model\ 2} = (A^{**}, B^{**}, t_c^{**}, \beta^{**}, C^{**}, \omega^{**}, \Phi^{**})$ for each iteration.

Panel A. Initial parameter values for Model 1

Spec	A	B	T_c=T+1	b	*SSR*
1	14.40	-1.00	5.52	0.20	**776.12***
2	14.40	-1.00	5.52	0.40	1510.28
3	14.40	-1.00	5.52	0.60	2959.08
4	14.40	-1.00	5.52	0.80	5713.93

Panel B. Optimized parameter values for Model 1

Spec	A*	B*	t_c^*	b*	*SSR*
1	14.67	-0.67	5.52	0.10	**21.06***
2	14.67	-0.67	5.52	0.10	21.06
3	14.67	-0.67	5.52	0.10	21.06
4	14.57	-0.58	5.52	0.10	21.16

Panel C. Initial parameter values for Model 2

Spec	1	2	3	4	5	6	7	8	9	10	11
A*	14.67	14.67	14.67	14.67	14.67	14.67	14.67	14.67	14.67	14.67	14.67
B*	-0.67	-0.67	-0.67	-0.67	-0.67	-0.67	-0.67	-0.67	-0.67	-0.67	-0.67
t_c*	5.52	5.52	5.52	5.52	5.52	5.52	5.52	5.52	5.52	5.52	5.52
b*	0.10	0.10	0.10	0.10	0.10	0.10	0.10	0.10	0.10	0.10	0.10
C*	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
ω*	5.00	6.00	7.00	8.00	9.00	10.00	11.00	12.00	13.00	14.00	15.00
Φ*	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
SSR	14.65	14.65	14.65	14.65	14.65	14.65	14.65	14.65	14.65	14.65	14.65

Panel D. Optimized parameter values for Model 2

Spec	1	2	3	4	5	6	7	8	9	10	11
A*	14.73	14.74	14.74	14.72	14.72	14.58	14.71	14.71	14.71	14.70	14.70
B*	-0.71	-0.73	-0.74	-0.72	-0.72	-0.57	-0.71	-0.71	-0.71	-0.70	-0.70
t_c**	5.71	5.52	5.52	5.52	5.52	5.52	5.52	5.52	5.52	5.52	5.52
b*	0.10	0.10	0.10	0.10	0.10	0.13	0.10	0.10	0.10	0.10	0.10
C*	0.15	0.14	-0.14	-0.14	-0.14	0.18	-0.07	-0.07	0.07	0.07	-0.07
ω**	5.81	6.38	6.38	6.38	6.38	6.37	11.36	11.35	12.72	13.24	14.98
φ **	-1.05	-1.49	1.65	1.65	1.65	4.79	0.00	0.01	0.90	0.02	0.31
SSR	13.65	**12.38***	12.38	12.38	12.38	12.38	19.26	19.26	18.85	19.01	19.31

Table 11. Calibrating the power-law model using 6 months of data before the momentum bubble peak on July 14, 2008.

Panel A shows the initial value of the first-stage estimation approach $\Phi^0_{model\ 1} = (A, B, t_c, \beta)$, which is a basic power-law model, as $ln[comp(t)] = Ln[com(T)] - (t_c - t)^{\overline{\beta}}$. In this setting, $A = ln[p(T)]$ is the natural logarithm of the last compounded return in the sample. Given the daily frequency of the sample, one unit of time equals 1/253. In the second stage, the model is optimized by minimizing the *SSR* of $ln[comp(t)] = Ln[com(T)] - (t_c - t)^{\overline{\beta}}$ by allowing for the following constraints: $t_c \geq T + 1\ and\ 0.1 \leq \beta \leq 0.9$. Panel B reports the optimal values for the parameter vector $\Phi^*_{model\ 1} = (A^*, B^*, t_c^*, \beta^*)$. In the third stage, using the optimized parameter values of Model 1 in Panel B as the initial values $\Phi^0_{model\ 2} = (A^*, B^*, t_c^*, \beta^*, C, \omega, \Phi)$ and setting $C = \phi = 0$ for varying values of $5 \leq \omega \leq 15$, we calibrate the LPPLS model 2 that includes log-periodic oscillations as $log[E(p(t))] \approx A + B(t_c - t)^{\beta} - (1 + C\ cos[\omega\ log(t_c - t) + \varphi])$ by considering the following constraints: $t_c \geq T + 1,\ 0.1 \leq \beta \leq 0.9, and\ 5 \leq \omega \leq 15$. Panel C outlines the input data vectors $\Phi^0_{model\ 2} = (A^*, B^*, t_c^*, \beta^*, C, \omega, \Phi)$, while Panel D presents the optimized parameter vectors $\Phi^*_{model\ 2} = (A^{**}, B^{**}, t_c^{**}, \beta^{**}, C^{**}, \omega^{**}, \Phi^{**})$ for each iteration.

Panel A. Initial parameter values for Model 1

Spec	A	B	T_c=T+1	b	*SSR*
1	14.35	-1.00	5.27	0.20	**788.32***
2	14.35	-1.00	5.27	0.40	1471.12
3	14.35	-1.00	5.27	0.60	2775.61
4	14.35	-1.00	5.27	0.80	5190.73

Panel B. Optimized parameter values for Model 1

Spec	A*	B*	t_c^*	b*	*SSR*
1	14.36	-0.40	5.27	0.10	**18.47***
2	14.36	-0.40	5.27	0.10	18.47
3	14.36	-0.39	5.27	0.10	18.47
4	14.36	-0.40	5.27	0.10	18.47

Panel C. Initial parameter values for Model 2

Spec	1	2	3	4	5	6	7	8	9	10	11
A*	14.36	14.36	14.36	14.36	14.36	14.36	14.36	14.36	14.36	14.36	14.36
B*	-0.40	-0.40	-0.40	-0.40	-0.40	-0.40	-0.40	-0.40	-0.40	-0.40	-0.40
t_c*	5.27	5.27	5.27	5.27	5.27	5.27	5.27	5.27	5.27	5.27	5.27
b*	0.10	0.10	0.10	0.10	0.10	0.10	0.10	0.10	0.10	0.10	0.10
C*	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
ω*	5.00	6.00	7.00	8.00	9.00	10.00	11.00	12.00	13.00	14.00	15.00
Φ*	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
SSR	18.47	18.47	18.47	18.47	18.47	18.47	18.47	18.47	18.47	18.47	18.47

Panel D. Optimized parameter values for Model 2

Spec	1	2	3	4	5	6	7	8	9	10	11
A*	14.40	13.98	13.98	14.40	14.40	14.40	14.41	14.40	14.08	14.39	14.39
B*	-0.43	-0.02	-0.02	-0.44	-0.44	-0.44	-0.43	-0.43	-0.12	-0.43	-0.43
t_c**	5.27	5.41	5.41	5.27	5.27	5.27	5.55	5.27	5.27	5.27	5.27
b*	0.10	0.90	0.90	0.10	0.10	0.10	0.10	0.10	0.31	0.10	0.10
C*	0.25	3.01	-3.01	-0.24	0.11	0.11	-0.11	-0.12	0.80	0.11	-0.10
ω**	5.59	5.58	5.58	5.63	10.49	10.17	11.31	11.77	5.54	13.66	14.98
φ **	-0.09	-0.45	2.69	3.00	-1.19	-0.65	0.00	-0.09	12.51	0.01	1.02
SSR	10.20	**9.31***	9.31	10.20	16.50	16.51	17.45	16.44	9.79	16.61	16.93

Table 12. Calibrating the power-law model using 9 months of data before the momentum bubble peak on July 14, 2008.

Panel A shows the initial value of the first-stage estimation approach $\Phi^0_{model\ 1} = (A, B, t_c, \beta)$, which is a basic power-law model, as $ln[comp(t)] = Ln[com(T)] - (t_c - t)^{\overline{\beta}}$. In this setting, $A = ln[p(T)]$ is the natural logarithm of the last compounded return in the sample. Given the daily frequency of the sample, one unit of time equals 1/253. In the second stage, the model is optimized by minimizing the *SSR* of $ln[comp(t)] = Ln[com(T)] - (t_c - t)^{\overline{\beta}}$ by allowing for the following constraints: $t_c \geq T + 1\ and\ 0.1 \leq \beta \leq 0.9$. Panel B reports the optimal values for the parameter vector $\Phi^*_{model\ 1} = (A^*, B^*, t_c^*, \beta^*)$. In the third stage, using the optimized parameter values of Model 1 in Panel B as the initial values $\Phi^0_{model\ 2} = (A^*, B^*, t_c^*, \beta^*, C, \omega, \Phi)$ and setting $C = \phi = 0$ for varying values of $5 \leq \omega \leq 15$, we calibrate the LPPLS model 2 that includes log-periodic oscillations as $log[E(p(t))] \approx A + B(t_c - t)^{\beta} - (1 + C\ cos[\omega\ log(t_c - t) + \varphi])$ by considering the following constraints: $t_c \geq T + 1,\ 0.1 \leq \beta \leq 0.9, and\ 5 \leq \omega \leq 15$. Panel C outlines the input data vectors $\Phi^0_{model\ 2} = (A^*, B^*, t_c^*, \beta^*, C, \omega, \Phi)$, while Panel D presents the optimized parameter vectors $\Phi^*_{model\ 2} = (A^{**}, B^{**}, t_c^{**}, \beta^{**}, C^{**}, \omega^{**}, \Phi^{**})$ for each iteration.

Panel A. Initial parameter values for Model 1

Spec	A	B	T_c=T+1	b	*SSR*
1	14.08	-1.00	5.02	0.20	**1303.28***
2	14.08	-1.00	5.02	0.40	2056.47
3	14.08	-1.00	5.02	0.60	3385.12
4	14.08	-1.00	5.02	0.80	5699.62

Panel B. Optimized parameter values for Model 1

Spec	A*	B*	t_c^*	b*	*SSR*
1	13.43	0.37	10.03	0.14	**14.18***
2	13.78	0.09	7.39	0.30	14.21
3	13.82	0.07	6.34	0.27	14.25
4	13.60	0.28	6.90	0.10	14.26

Panel C. Initial parameter values for Model 2

Spec	1	2	3	4	5	6	7	8	9	10	11
A*	13.43	13.43	13.43	13.43	13.43	13.43	13.43	13.43	13.43	13.43	13.43
B*	0.37	0.37	0.37	0.37	0.37	0.37	0.37	0.37	0.37	0.37	0.37
t_c*	10.03	10.03	10.03	10.03	10.03	10.03	10.03	10.03	10.03	10.03	10.03
b*	0.14	0.14	0.14	0.14	0.14	0.14	0.14	0.14	0.14	0.14	0.14
C*	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
ω*	5.00	6.00	7.00	8.00	9.00	10.00	11.00	12.00	13.00	14.00	15.00
Φ*	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
SSR	14.18	14.18	14.18	14.18	14.18	14.18	14.18	14.18	14.18	14.18	14.18

Panel D. Optimized parameter values for Model 2

Spec	1	2	3	4	5	6	7	8	9	10	11
A*	10.07	11.78	11.35	13.89	13.89	13.89	13.89	13.89	13.89	13.89	13.91
B*	1.36	0.39	0.57	0.02	0.02	0.02	0.02	0.02	0.02	0.02	0.01
t_c**	10.03	10.03	10.03	5.61	5.61	5.61	5.61	5.61	5.61	5.61	9.28
b*	0.54	0.90	0.79	0.90	0.90	0.90	0.90	0.90	0.90	0.90	0.10
C*	0.15	-0.24	0.21	2.80	-2.80	2.80	2.80	-2.80	2.80	2.80	-14.08
ω**	5.00	5.00	5.00	5.65	5.65	5.65	5.65	5.65	5.65	5.65	14.98
φ **	-1.72	1.48	4.59	8.37	11.51	14.65	14.65	17.79	20.93	20.93	1.39
SSR	6.65	6.27	6.39	3.67	3.67	**3.66***	3.67	3.67	3.67	3.66	6.46

Table 13. Calibrating the power-law model using 12 months of data before the momentum bubble peak on July 14, 2008.

Panel A shows the initial value of the first-stage estimation approach $\Phi^0_{model\ 1} = (A, B, t_c, \beta)$, which is a basic power-law model, as $ln[comp(t)] = Ln[com(T)] - (t_c - t)^{\overline{\beta}}$. In this setting, $A = ln[p(T)]$ is the natural logarithm of the last compounded return in the sample. Given the daily frequency of the sample, one unit of time equals 1/253. In the second stage, the model is optimized by minimizing the *SSR* of $ln[comp(t)] = Ln[com(T)] - (t_c - t)^{\overline{\beta}}$ by allowing for the following constraints: $t_c \geq T + 1\ and\ 0.1 \leq \beta \leq 0.9$. Panel B reports the optimal values for the parameter vector $\Phi^*_{model\ 1} = (A^*, B^*, t_c^*, \beta^*)$. In the third stage, using the optimized parameter values of Model 1 in Panel B as the initial values $\Phi^0_{model\ 2} = (A^*, B^*, t_c^*, \beta^*, C, \omega, \Phi)$ and setting $C = \phi = 0$ for varying values of $5 \leq \omega \leq 15$, we calibrate the LPPLS model 2 that includes log-periodic oscillations as $log[E(p(t))] \approx A + B(t_c - t)^{\beta} - (1 + C\ cos[\omega\ log(t_c - t) + \varphi])$ by considering the following constraints: $t_c \geq T + 1,\ 0.1 \leq \beta \leq 0.9, and\ 5 \leq \omega \leq 15$. Panel C outlines the input data vectors $\Phi^0_{model\ 2} = (A^*, B^*, t_c^*, \beta^*, C, \omega, \Phi)$, while Panel D presents the optimized parameter vectors $\Phi^*_{model\ 2} = (A^{**}, B^{**}, t_c^{**}, \beta^{**}, C^{**}, \omega^{**}, \Phi^{**})$ for each iteration.

Panel A. Initial parameter values for Model 1

Spec	A	B	T_c=T+1	b	*SSR*
1	13.95	-1.00	4.76	0.20	**1539.16***
2	13.95	-1.00	4.76	0.40	2259.73
3	13.95	-1.00	4.76	0.60	3483.32
4	13.95	-1.00	4.76	0.80	5541.54

Panel B. Optimized parameter values for Model 1

Spec	A*	B*	t_c^*	b*	*SSR*
1	13.70	0.15	6.46	**0.27**	13.48
2	13.64	0.17	7.60	0.29	**13.43***
3	13.70	0.15	6.46	0.27	13.48
4	13.16	0.65	7.37	0.10	13.49

Panel C. Initial parameter values for Model 2

Spec	1	2	3	4	5	6	7	8	9	10	11
A*	13.64	13.64	13.64	13.64	13.64	13.64	13.64	13.64	13.64	13.64	13.64
B*	0.17	0.17	0.17	0.17	0.17	0.17	0.17	0.17	0.17	0.17	0.17
t_c*	7.60	7.60	7.60	7.60	7.60	7.60	7.60	7.60	7.60	7.60	7.60
b*	0.29	0.29	0.29	0.29	0.29	0.29	0.29	0.29	0.29	0.29	0.29
C*	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
ω*	5.00	6.00	7.00	8.00	9.00	10.00	11.00	12.00	13.00	14.00	15.00
Φ*	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
SSR	13.43	13.43	13.43	13.43	13.43	13.43	13.43	13.43	13.43	13.43	13.43

Panel D. Optimized parameter values for Model 2

Spec	1	2	3	4	5	6	7	8	9	10	11
A*	10.17	13.86	13.88	13.86	9.82	9.69	9.94	10.38	4.42	10.00	10.47
B*	0.48	0.03	0.02	0.03	0.57	0.54	0.52	0.46	1.00	0.45	0.39
t_c**	12.84	5.32	5.29	5.32	12.00	12.74	12.49	12.48	15.52	13.82	14.09
b*	0.90	0.90	0.90	0.90	0.90	0.90	0.90	0.90	0.89	0.90	0.90
C*	0.16	1.82	-2.42	-1.81	0.18	-0.16	-0.17	0.16	0.15	-0.14	-0.14
ω**	7.79	5.00	5.00	5.00	6.61	7.30	7.21	7.66	7.19	8.77	9.61
φ**	-10.01	3.33	6.55	6.47	-0.39	0.63	1.03	3.13	2.23	2.60	0.34
SSR	3.77	**3.13***	3.17	3.13	3.80	3.77	3.78	3.79	3.70	3.75	3.75

Table 14. Calibrating the power-law model using 3 months of data before the momentum bubble peak on January 20, 2016.

Panel A shows the initial value of the first-stage estimation approach $\Phi^0_{model\ 1} = (A, B, t_c, \beta)$, which is a basic power-law model, as $ln[comp(t)] = Ln[com(T)] - (t_c - t)^{\overline{\beta}}$. In this setting, $A = ln[p(T)]$ is the natural logarithm of the last compounded return in the sample. Given the daily frequency of the sample, one unit of time equals 1/253. In the second stage, the model is optimized by minimizing the *SSR* of $ln[comp(t)] = Ln[com(T)] - (t_c - t)^{\overline{\beta}}$ by allowing for the following constraints: $t_c \geq T + 1\ and\ 0.1 \leq \beta \leq 0.9$. Panel B reports the optimal values for the parameter vector $\Phi^*_{model\ 1} = (A^*, B^*, t_c^*, \beta^*)$. In the third stage, using the optimized parameter values of Model 1 in Panel B as the initial values $\Phi^0_{model\ 2} = (A^*, B^*, t_c^*, \beta^*, C, \omega, \Phi)$ and setting $C = \phi = 0$ for varying values of $5 \leq \omega \leq 15$, we calibrate the LPPLS model 2 that includes log-periodic oscillations as $log[E(p(t))] \approx A + B(t_c - t)^{\beta} - (1 + C\ cos[\omega\ log(t_c - t) + \varphi])$ by considering the following constraints: $t_c \geq T + 1,\ 0.1 \leq \beta \leq 0.9, and\ 5 \leq \omega \leq 15$. Panel C outlines the input data vectors $\Phi^0_{model\ 2} = (A^*, B^*, t_c^*, \beta^*, C, \omega, \Phi)$, while Panel D presents the optimized parameter vectors $\Phi^*_{model\ 2} = (A^{**}, B^{**}, t_c^{**}, \beta^{**}, C^{**}, \omega^{**}, \Phi^{**})$ for each iteration.

Panel A. Initial parameter values for Model 1

Spec	A	B	T_c=T+1	b	*SSR*
1	14.09	-1.00	6.60	0.20	**1004.55***
2	14.09	-1.00	6.60	0.40	2060.17
3	14.09	-1.00	6.60	0.60	4340.70
4	14.09	-1.00	6.60	0.80	9049.86

Panel B. Optimized parameter values for Model 1

Spec	A*	B*	t_c^*	b*	*SSR*
1	15.26	-1.21	7.66	0.21	35.80
2	14.40	-0.51	6.91	0.36	35.68
3	14.40	-0.50	6.89	0.36	**35.67***
4	14.47	-0.58	6.91	0.32	35.69

Panel C. Initial parameter values for Model 2

Spec	1	2	3	4	5	6	7	8	9	10	11
A*	14.40	14.40	14.40	14.40	14.40	14.40	14.40	14.40	14.40	14.40	14.40
B*	0.13	-0.50	-0.50	-0.50	-0.50	-0.50	-0.50	-0.50	-0.50	-0.50	-0.50
t_c*	6.89	6.89	6.89	6.89	6.89	6.89	6.89	6.89	6.89	6.89	6.89
b*	0.36	0.36	0.36	0.36	0.36	0.36	0.36	0.36	0.36	0.36	0.36
C*	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
ω*	5.00	6.00	7.00	8.00	9.00	10.00	11.00	12.00	13.00	14.00	15.00
Φ*	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
SSR	35.67	35.67	35.67	35.67	35.67	35.67	35.67	35.67	35.67	35.67	35.67

Panel D. Optimized parameter values for Model 2

Spec	1	2	3	4	5	6	7	8	9	10	11
A*	16.69	14.25	14.24	14.24	13.97	13.97	14.07	14.04	13.94	13.97	13.96
B*	-2.61	-0.43	-0.43	-0.43	-0.10	-0.13	-0.13	-0.11	-0.10	-0.13	-0.12
t_c**	8.07	6.60	6.60	6.60	6.91	6.60	7.60	7.58	6.60	6.60	6.60
b*	0.10	0.30	0.30	0.30	0.90	0.78	0.82	0.89	0.90	0.77	0.83
C*	-0.03	-0.18	0.18	0.18	-0.30	-0.29	0.23	-0.25	-0.33	0.29	0.30
ω**	6.63	5.00	5.00	5.00	10.40	9.80	12.88	12.82	9.87	9.80	10.03
φ**	-3.06	1.35	4.49	4.49	-2.01	-0.43	-4.73	-1.45	5.76	8.99	8.60
SSR	28.46	28.06	28.06	28.06	28.52	28.24	27.91	**27.90***	28.29	28.24	28.26

Table 15. Calibrating the power-law model using 6 months of data before the momentum bubble peak on January 20, 2016.

Panel A shows the initial value of the first-stage estimation approach $\Phi^{0}_{model\ 1} = (A, B, t_c, \beta)$, which is a basic power-law model, as $ln[comp(t)] = Ln[com(T)] - (t_c - t)^{\overline{\beta}}$. In this setting, $A = ln[p(T)]$ is the natural logarithm of the last compounded return in the sample. Given the daily frequency of the sample, one unit of time equals 1/253. In the second stage, the model is optimized by minimizing the *SSR* of $ln[comp(t)] = Ln[com(T)] - (t_c - t)^{\overline{\beta}}$ by allowing for the following constraints: $t_c \geq T + 1\ and\ 0.1 \leq \beta \leq 0.9$. Panel B reports the optimal values for the parameter vector $\Phi^{*}_{model\ 1} = (A^*, B^*, t_c^*, \beta^*)$. In the third stage, using the optimized parameter values of Model 1 in Panel B as the initial values $\Phi^{0}_{model\ 2} = (A^*, B^*, t_c^*, \beta^*, C, \omega, \Phi)$ and setting $C = \phi = 0$ for varying values of $5 \leq \omega \leq 15$, we calibrate the LPPLS model 2 that includes log-periodic oscillations as $log[E(p(t))] \approx A + B(t_c - t)^{\beta} - (1 + C\ cos[\omega\ log(t_c - t) + \varphi])$ by considering the following constraints: $t_c \geq T + 1,\ 0.1 \leq \beta \leq 0.9, and\ 5 \leq \omega \leq 15$. Panel C outlines the input data vectors $\Phi^{0}_{model\ 2} = (A^*, B^*, t_c^*, \beta^*, C, \omega, \Phi)$, while Panel D presents the optimized parameter vectors $\Phi^{*}_{model\ 2} = (A^{**}, B^{**}, t_c^{**}, \beta^{**}, C^{**}, \omega^{**}, \Phi^{**})$ for each iteration.

Panel A. Initial parameter values for Model 1

Spec	A	B	T_c=T+1	b	*SSR*
1	14.06	-1.00	6.34	0.20	**971.17***
2	14.06	-1.00	6.34	0.40	1950.51
3	14.06	-1.00	6.34	0.60	4013.78
4	14.06	-1.00	6.34	0.80	8191.13

Panel B. Optimized parameter values for Model 1

Spec	A*	B*	t_c^*	b*	*SSR*
1	16.64	-1.68	13.15	0.26	33.48
2	15.06	-0.59	11.24	0.43	33.49
3	14.70	-0.32	11.21	0.59	33.46
4	15.13	-0.53	12.25	0.47	**33.46***

Panel C. Initial parameter values for Model 2

Spec	1	2	3	4	5	6	7	8	9	10	11
A*	15.13	15.13	15.13	15.13	15.13	15.13	15.13	15.13	15.13	15.13	15.13
B*	-0.53	-0.53	-0.53	-0.53	-0.53	-0.53	-0.53	-0.53	-0.53	-0.53	-0.53
t_c*	12.25	12.25	12.25	12.25	12.25	12.25	12.25	12.25	12.25	12.25	12.25
b*	0.59	0.59	0.59	0.59	0.59	0.59	0.59	0.59	0.59	0.59	0.59
C*	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
ω*	5.00	6.00	7.00	8.00	9.00	10.00	11.00	12.00	13.00	14.00	15.00
Φ*	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
SSR	33.46	33.46	33.46	33.46	33.46	33.46	33.46	33.46	33.46	33.46	33.46

Panel D. Optimized parameter values for Model 2

Spec	1	2	3	4	5	6	7	8	9	10	11
A*	13.65	14.39	14.07	13.97	14.09	14.18	14.18	14.23	14.29	14.34	14.33
B*	0.00	-0.45	-0.09	-0.08	-0.09	-0.09	-0.09	-0.09	-0.10	-0.10	-0.10
t_c**	24.00	13.03	9.32	8.18	9.44	10.39	10.46	10.99	11.58	12.17	12.07
b*	0.90	0.23	0.90	0.90	0.90	0.90	0.90	0.90	0.90	0.90	0.90
C*	12.57	-0.20	0.26	0.35	-0.25	0.21	0.20	-0.18	0.17	-0.15	-0.16
ω**	15.00	6.86	10.06	8.10	10.27	11.85	11.96	12.85	13.85	14.83	14.65
φ**	-34.43	-1.45	-1.72	3.61	0.84	-0.66	-0.99	-0.54	-0.49	-0.45	0.10
SSR	29.66	31.03	25.50	**25.47***	25.51	25.56	25.56	25.58	25.61	25.63	25.63

Table 16. Calibrating the power-law model using 9 months of data before the momentum bubble peak on January 20, 2016.

Panel A shows the initial value of the first-stage estimation approach $\Phi^0_{model\ 1} = (A, B, t_c, \beta)$, which is a basic power-law model, as $ln[comp(t)] = Ln[com(T)] - (t_c - t)^{\overline{\beta}}$. In this setting, $A = ln[p(T)]$ is the natural logarithm of the last compounded return in the sample. Given the daily frequency of the sample, one unit of time equals 1/253. In the second stage, the model is optimized by minimizing the *SSR* of $ln[comp(t)] = Ln[com(T)] - (t_c - t)^{\overline{\beta}}$ by allowing for the following constraints: $t_c \geq T + 1\ and\ 0.1 \leq \beta \leq 0.9$. Panel B reports the optimal values for the parameter vector $\Phi^*_{model\ 1} = (A^*, B^*, t_c^*, \beta^*)$. In the third stage, using the optimized parameter values of Model 1 in Panel B as the initial values $\Phi^0_{model\ 2} = (A^*, B^*, t_c^*, \beta^*, C, \omega, \Phi)$ and setting $C = \phi = 0$ for varying values of $5 \leq \omega \leq 15$, we calibrate the LPPLS model 2 that includes log-periodic oscillations as $log[E(p(t))] \approx A + B(t_c - t)^{\beta} - (1 + C\ cos[\omega\ log(t_c - t) + \varphi])$ by considering the following constraints: $t_c \geq T + 1,\ 0.1 \leq \beta \leq 0.9, and\ 5 \leq \omega \leq 15$. Panel C outlines the input data vectors $\Phi^0_{model\ 2} = (A^*, B^*, t_c^*, \beta^*, C, \omega, \Phi)$, while Panel D presents the optimized parameter vectors $\Phi^*_{model\ 2} = (A^{**}, B^{**}, t_c^{**}, \beta^{**}, C^{**}, \omega^{**}, \Phi^{**})$ for each iteration.

Panel A. Initial parameter values for Model 1

Spec	A	B	T_c=T+1	b	*SSR*
1	13.79	-1.00	6.10	0.20	**1619.91***
2	13.79	-1.00	6.10	0.40	2726.65
3	13.79	-1.00	6.10	0.60	4867.65
4	13.79	-1.00	6.10	0.80	8945.32

Panel B. Optimized parameter values for Model 1

Spec	A*	B*	t_c^*	b*	*SSR*
1	16.44	-1.63	12.18	0.25	33.18
2	15.12	-0.61	11.42	0.43	**33.17***
3	14.53	-0.39	8.94	0.49	33.26
4	14.45	-0.33	8.91	0.54	33.23

Panel C. Initial parameter values for Model 2

Spec	1	2	3	4	5	6	7	8	9	10	11
A*	15.12	15.12	15.12	15.12	15.12	15.12	15.12	15.12	15.12	15.12	15.12
B*	-0.61	-0.61	-0.61	-0.61	-0.61	-0.61	-0.61	-0.61	-0.61	-0.61	-0.61
t_c*	11.42	11.42	11.42	11.42	11.42	11.42	11.42	11.42	11.42	11.42	11.42
b*	0.43	0.43	0.43	0.43	0.43	0.43	0.43	0.43	0.43	0.43	0.43
C*	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
ω*	5.00	6.00	7.00	8.00	9.00	10.00	11.00	12.00	13.00	14.00	15.00
Φ*	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
SSR	33.17	33.17	33.17	33.17	33.17	33.17	33.17	33.17	33.17	33.17	33.17

Panel D. Optimized parameter values for Model 2

Spec	1	2	3	4	5	6	7	8	9	10	11
A*	13.95	13.95	14.00	14.06	14.13	14.14	14.19	14.24	14.24	14.29	14.34
B*	-0.08	-0.08	-0.08	-0.09	-0.09	-0.09	-0.09	-0.10	-0.10	-0.10	-0.10
t_c**	7.89	7.90	8.47	9.03	9.76	9.84	10.35	10.90	10.92	11.45	11.95
b*	0.90	0.90	0.90	0.90	0.90	0.90	0.90	0.90	0.90	0.90	0.90
C*	0.39	-0.38	-0.32	0.27	-0.23	-0.23	0.20	-0.18	-0.18	0.17	-0.16
ω**	7.66	7.68	8.70	9.68	10.95	11.08	11.97	12.93	12.96	13.88	14.74
φ **	-1.50	1.60	-1.04	-0.56	-1.01	-1.40	-0.84	-0.59	-0.70	-0.38	0.05
SSR	**25.06***	25.06	25.08	25.10	25.13	25.14	25.16	25.19	25.19	25.21	25.24

Table 17. Calibrating the power-law model using 12 months of data before the momentum bubble peak on January 20, 2016.

Panel A shows the initial value of the first-stage estimation approach $\Phi^0_{model\ 1} = (A, B, t_c, \beta)$, which is a basic power-law model, as $ln[comp(t)] = Ln[com(T)] - (t_c - t)^{\overline{\beta}}$. In this setting, $A = ln[p(T)]$ is the natural logarithm of the last compounded return in the sample. Given the daily frequency of the sample, one unit of time equals 1/253. In the second stage, the model is optimized by minimizing the *SSR* of $ln[comp(t)] = Ln[com(T)] - (t_c - t)^{\overline{\beta}}$ by allowing for the following constraints: $t_c \geq T + 1\ and\ 0.1 \leq \beta \leq 0.9$. Panel B reports the optimal values for the parameter vector $\Phi^*_{model\ 1} = (A^*, B^*, t_c^*, \beta^*)$. In the third stage, using the optimized parameter values of Model 1 in Panel B as the initial values $\Phi^0_{model\ 2} = (A^*, B^*, t_c^*, \beta^*, C, \omega, \Phi)$ and setting $C = \phi = 0$ for varying values of $5 \leq \omega \leq 15$, we calibrate the LPPLS model 2 that includes log-periodic oscillations as $log[E(p(t))] \approx A + B(t_c - t)^{\beta} - (1 + C\ cos[\omega\ log(t_c - t) + \varphi])$ by considering the following constraints: $t_c \geq T + 1,\ 0.1 \leq \beta \leq 0.9, and\ 5 \leq \omega \leq 15$. Panel C outlines the input data vectors $\Phi^0_{model\ 2} = (A^*, B^*, t_c^*, \beta^*, C, \omega, \Phi)$, while Panel D presents the optimized parameter vectors $\Phi^*_{model\ 2} = (A^{**}, B^{**}, t_c^{**}, \beta^{**}, C^{**}, \omega^{**}, \Phi^{**})$ for each iteration.

Panel A. Initial parameter values for Model 1

Spec	A	B	T_c=T+1	b	*SSR*
1	13.90	-1.00	5.85	0.20	**1190.31***
2	13.90	-1.00	5.85	0.40	2096.20
3	13.90	-1.00	5.85	0.60	3875.41
4	13.90	-1.00	5.85	0.80	7281.09

Panel B. Optimized parameter values for Model 1

Spec	A*	B*	t_c^*	b*	*SSR*
1	16.31	-1.58	11.70	0.25	33.11
2	15.11	-0.57	11.70	0.45	**33.08***
3	14.27	-0.19	8.99	0.71	33.10
4	14.21	-0.24	7.72	0.61	33.21

Panel C. Initial parameter values for Model 2

Spec	1	2	3	4	5	6	7	8	9	10	11
A*	15.11	15.11	15.11	15.11	15.11	15.11	15.11	15.11	15.11	15.11	15.11
B*	-0.57	-0.57	-0.57	-0.57	-0.57	-0.57	-0.57	-0.57	-0.57	-0.57	-0.57
t_c*	11.70	11.70	11.70	11.70	11.70	11.70	11.70	11.70	11.70	11.70	11.70
b*	0.45	0.45	0.45	0.45	0.45	0.45	0.45	0.45	0.45	0.45	0.45
C*	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
ω*	5.00	6.00	7.00	8.00	9.00	10.00	11.00	12.00	13.00	14.00	15.00
Φ*	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
SSR	33.08	33.08	33.08	33.08	33.08	33.08	33.08	33.08	33.08	33.08	33.08

Panel D. Optimized parameter values for Model 2

Spec	1	2	3	4	5	6	7	8	9	10	11
A*	14.01	14.00	14.24	14.05	14.10	14.25	14.17	14.99	14.29	14.30	14.32
B*	-0.08	-0.08	-0.10	-0.09	-0.09	-0.10	-0.09	-0.58	-0.10	-0.10	-0.10
t_c**	8.48	8.41	10.82	8.88	9.44	10.95	10.10	10.87	11.40	11.43	11.65
b*	0.90	0.90	0.90	0.90	0.90	0.90	0.90	0.42	0.90	0.90	0.90
C*	0.31	-0.32	-0.19	0.28	-0.24	-0.18	0.21	-0.08	0.17	0.17	-0.16
ω**	8.72	8.60	12.88	9.44	10.45	13.10	11.62	12.70	13.88	13.93	14.32
φ**	-4.24	-0.76	-12.94	0.13	0.45	-7.32	0.24	-0.07	-0.32	-0.46	1.46
SSR	24.97	**24.97***	25.06	24.98	25.00	25.07	25.03	25.53	25.09	25.09	25.10

Table 18. Calibrating the power-law model using 3 months of data before the momentum bubble peak on April 16, 2020.

Panel A shows the initial value of the first-stage estimation approach $\Phi^{0}_{model\ 1} = (A, B, t_c, \beta)$, which is a basic power-law model, as $ln[comp(t)] = Ln[com(T)] - (t_c - t)^{\overline{\beta}}$. In this setting, $A = ln[p(T)]$ is the natural logarithm of the last compounded return in the sample. Given the daily frequency of the sample, one unit of time equals 1/253. In the second stage, the model is optimized by minimizing the *SSR* of $ln[comp(t)] = Ln[com(T)] - (t_c - t)^{\overline{\beta}}$ by allowing for the following constraints: $t_c \geq T + 1\ and\ 0.1 \leq \beta \leq 0.9$. Panel B reports the optimal values for the parameter vector $\Phi^{*}_{model\ 1} = (A^{*}, B^{*}, t_c^{*}, \beta^{*})$. In the third stage, using the optimized parameter values of Model 1 in Panel B as the initial values $\Phi^{0}_{model\ 2} = (A^{*}, B^{*}, t_c^{*}, \beta^{*}, C, \omega, \Phi)$ and setting $C = \phi = 0$ for varying values of $5 \leq \omega \leq 15$, we calibrate the LPPLS model 2 that includes log-periodic oscillations as $log[E(p(t))] \approx A + B(t_c - t)^{\beta} - (1 + C\,cos[\omega\, log(t_c - t) + \varphi])$ by considering the following constraints: $t_c \geq T + 1,\ 0.1 \leq \beta \leq 0.9, and\ 5 \leq \omega \leq 15$. Panel C outlines the input data vectors $\Phi^{0}_{model\ 2} = (A^{*}, B^{*}, t_c^{*}, \beta^{*}, C, \omega, \Phi)$, while Panel D presents the optimized parameter vectors $\Phi^{*}_{model\ 2} = (A^{**}, B^{**}, t_c^{**}, \beta^{**}, C^{**}, \omega^{**}, \Phi^{**})$ for each iteration.

Panel A. Initial parameter values for Model 1

Spec	**A**	**B**	**T_c=T+1**	**b**	***SSR***
1	14.25	-1.00	3.98	0.20	**866.29***
2	14.25	-1.00	3.98	0.40	1218.05
3	14.25	-1.00	3.98	0.60	1820.84
4	14.25	-1.00	3.98	0.80	2813.23

Panel B. Optimized parameter values for Model 1

Spec	**A***	**B***	**t_c^{*}**	**b***	***SSR***
1	15.06	-0.83	4.72	0.20	**8.24***
2	14.59	-0.40	4.48	0.35	8.29
3	14.36	-0.21	4.11	0.51	8.34
4	14.64	-0.44	4.42	0.30	8.26

Panel C. Initial parameter values for Model 2

Spec	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**	**11**
A*	15.06	15.06	15.06	15.06	15.06	15.06	15.06	15.06	15.06	15.06	15.06
B*	-0.83	-0.83	-0.83	-0.83	-0.83	-0.83	-0.83	-0.83	-0.83	-0.83	-0.83
t_c*	4.72	4.72	4.72	4.72	4.72	4.72	4.72	4.72	4.72	4.72	4.72
b*	0.20	0.20	0.20	0.20	0.20	0.20	0.20	0.20	0.20	0.20	0.20
C*	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
ω*	5.00	6.00	7.00	8.00	9.00	10.00	11.00	12.00	13.00	14.00	15.00
Φ*	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
SSR	8.24	8.24	8.24	8.24	8.24	8.24	8.24	8.24	8.24	8.24	8.24

Panel D. Optimized parameter values for Model 2

Spec	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**	**11**
A	14.26	14.26	14.26	14.26	14.26	14.25	14.27	14.25	14.25	14.34	14.32
B	-0.10	-0.11	-0.11	-0.11	-0.10	-0.10	-0.12	-0.10	-0.10	-0.11	-0.10
t_c*	4.09	3.98	3.99	3.98	3.98	3.98	3.98	3.98	3.98	4.87	4.75
b*	0.75	0.68	0.68	0.68	0.70	0.72	0.61	0.90	0.75	0.89	0.90
C*	0.68	0.72	-0.72	-0.72	-0.73	0.74	0.66	-0.33	-0.78	0.22	0.23
ω*	5.41	5.01	5.05	5.01	5.03	5.03	5.00	8.91	5.00	15.00	14.32
***φ **	0.86	1.48	4.57	4.62	4.61	7.75	7.77	6.55	10.93	-1.54	-0.12
SSR	2.94	2.90	**2.90***	2.90	2.90	2.90	2.91	6.31	2.91	6.76	6.76

Table 19. Calibrating the power-law model using 6 months of data before the momentum bubble peak on April 16, 2020.

Panel A shows the initial value of the first-stage estimation approach $\Phi^0_{model\ 1} = (A, B, t_c, \beta)$, which is a basic power-law model, as $ln[comp(t)] = Ln[com(T)] - (t_c - t)^{\overline{\beta}}$. In this setting, $A = ln[p(T)]$ is the natural logarithm of the last compounded return in the sample. Given the daily frequency of the sample, one unit of time equals 1/253. In the second stage, the model is optimized by minimizing the *SSR* of $ln[comp(t)] = Ln[com(T)] - (t_c - t)^{\overline{\beta}}$ by allowing for the following constraints: $t_c \geq T + 1\ and\ 0.1 \leq \beta \leq 0.9$. Panel B reports the optimal values for the parameter vector $\Phi^*_{model\ 1} = (A^*, B^*, t_c^*, \beta^*)$. In the third stage, using the optimized parameter values of Model 1 in Panel B as the initial values $\Phi^0_{model\ 2} = (A^*, B^*, t_c^*, \beta^*, C, \omega, \Phi)$ and setting $C = \phi = 0$ for varying values of $5 \leq \omega \leq 15$, we calibrate the LPPLS model 2 that includes log-periodic oscillations as $log[E(p(t))] \approx A + B(t_c - t)^{\beta} - (1 + C\ cos[\omega\ log(t_c - t) + \varphi])$ by considering the following constraints: $t_c \geq T + 1,\ 0.1 \leq \beta \leq 0.9, and\ 5 \leq \omega \leq 15$. Panel C outlines the input data vectors $\Phi^0_{model\ 2} = (A^*, B^*, t_c^*, \beta^*, C, \omega, \Phi)$, while Panel D presents the optimized parameter vectors $\Phi^*_{model\ 2} = (A^{**}, B^{**}, t_c^{**}, \beta^{**}, C^{**}, \omega^{**}, \Phi^{**})$ for each iteration.

Panel A. Initial parameter values for Model 1

Spec	A	B	T_c=T+1	b	*SSR*
1	14.27	-1.00	3.73	0.20	**729.39***
2	14.27	-1.00	3.73	0.40	1010.87
3	14.27	-1.00	3.73	0.60	1490.22
4	14.27	-1.00	3.73	0.80	2267.56

Panel B. Optimized parameter values for Model 1

Spec	A*	B*	t_c^*	b*	*SSR*
1	14.85	-0.70	4.08	0.19	**8.06***
2	14.48	-0.35	3.95	0.34	8.13
3	14.36	-0.25	3.73	0.39	8.11
4	14.53	-0.40	3.91	0.28	8.07

Panel C. Initial parameter values for Model 2

Spec	1	2	3	4	5	6	7	8	9	10	11
A*	14.85	14.85	14.85	14.85	14.85	14.85	14.85	14.85	14.85	14.85	14.85
B*	-0.70	-0.70	-0.70	-0.70	-0.70	-0.70	-0.70	-0.70	-0.70	-0.70	-0.70
t_c*	4.08	4.08	4.08	4.08	4.08	4.08	4.08	4.08	4.08	4.08	4.08
b*	0.19	0.19	0.19	0.19	0.19	0.19	0.19	0.19	0.19	0.19	0.19
C*	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
ω*	5.00	6.00	7.00	8.00	9.00	10.00	11.00	12.00	13.00	14.00	15.00
Φ*	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
SSR	8.06	8.06	8.06	8.06	8.06	8.06	8.06	8.06	8.06	8.06	8.06

Panel D. Optimized parameter values for Model 2

Spec	1	2	3	4	5	6	7	8	9	10	11
A*	14.26	14.26	14.26	14.27	14.22	14.22	14.22	14.22	14.28	14.26	14.25
B*	-0.11	-0.11	-0.11	-0.12	-0.07	-0.07	-0.07	-0.07	-0.12	-0.12	-0.10
t_c**	3.98	3.99	3.98	3.98	3.99	4.01	3.99	3.99	4.05	3.94	3.94
b*	0.64	0.64	0.64	0.61	0.90	0.90	0.90	0.90	0.81	0.82	0.90
C*	0.71	0.70	0.71	-0.69	-0.96	-0.94	0.96	0.96	-0.28	-0.30	0.33
ω**	5.00	5.04	5.00	5.00	5.00	5.07	5.00	5.00	14.98	13.52	13.47
φ **	1.47	1.42	1.47	4.62	4.62	4.51	7.76	7.76	-1.68	0.64	3.86
SSR	2.80	2.80	**2.80***	2.80	2.83	2.83	2.83	2.83	6.36	6.28	6.29

Table 20. Calibrating the power-law model using 9 months of data before the momentum bubble peak on April 16, 2020.

Panel A shows the initial value of the first-stage estimation approach $\Phi^0_{model\ 1} = (A, B, t_c, \beta)$, which is a basic power-law model, as $ln[comp(t)] = Ln[com(T)] - (t_c - t)^{\overline{\beta}}$. In this setting, $A = ln[p(T)]$ is the natural logarithm of the last compounded return in the sample. Given the daily frequency of the sample, one unit of time equals 1/253. In the second stage, the model is optimized by minimizing the *SSR* of $ln[comp(t)] = Ln[com(T)] - (t_c - t)^{\overline{\beta}}$ by allowing for the following constraints: $t_c \geq T + 1\ and\ 0.1 \leq \beta \leq 0.9$. Panel B reports the optimal values for the parameter vector $\Phi^*_{model\ 1} = (A^*, B^*, t_c^*, \beta^*)$. In the third stage, using the optimized parameter values of Model 1 in Panel B as the initial values $\Phi^0_{model\ 2} = (A^*, B^*, t_c^*, \beta^*, C, \omega, \Phi)$ and setting $C = \phi = 0$ for varying values of $5 \leq \omega \leq 15$, we calibrate the LPPLS model 2 that includes log-periodic oscillations as $log[E(p(t))] \approx A + B(t_c - t)^{\beta} - (1 + C\ cos[\omega\ log(t_c - t) + \varphi])$ by considering the following constraints: $t_c \geq T + 1,\ 0.1 \leq \beta \leq 0.9, and\ 5 \leq \omega \leq 15$. Panel C outlines the input data vectors $\Phi^0_{model\ 2} = (A^*, B^*, t_c^*, \beta^*, C, \omega, \Phi)$, while Panel D presents the optimized parameter vectors $\Phi^*_{model\ 2} = (A^{**}, B^{**}, t_c^{**}, \beta^{**}, C^{**}, \omega^{**}, \Phi^{**})$ for each iteration.

Panel A. Initial parameter values for Model 1

Spec	**A**	**B**	**T_c=T+1**	**b**	***SSR***
1	14.22	-1.00	3.47	0.20	**718.70***
2	14.22	-1.00	3.47	0.40	953.30
3	14.22	-1.00	3.47	0.60	1343.43
4	14.22	-1.00	3.47	0.80	1958.45

Panel B. Optimized parameter values for Model 1

Spec	**A***	**B***	**t_c^***	**b***	***SSR***
1	14.98	-0.77	4.61	0.20	**7.75***
2	14.49	-0.33	4.25	0.35	7.78
3	14.27	-0.16	3.73	0.55	7.83
4	14.54	-0.40	4.08	0.29	7.76

Panel C. Initial parameter values for Model 2

Spec	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**	**11**
A*	14.98	14.98	14.98	14.98	14.98	14.98	14.98	14.98	14.98	14.98	14.98
B*	-0.77	-0.77	-0.77	-0.77	-0.77	-0.77	-0.77	-0.77	-0.77	-0.77	-0.77
t_c*	4.61	4.61	4.61	4.61	4.61	4.61	4.61	4.61	4.61	4.61	4.61
b*	0.20	0.20	0.20	0.20	0.20	0.20	0.20	0.20	0.20	0.20	0.20
C*	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
ω*	5.00	6.00	7.00	8.00	9.00	10.00	11.00	12.00	13.00	14.00	15.00
Φ*	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
SSR	7.76	7.76	7.76	7.76	7.76	7.76	7.76	7.76	7.76	7.76	7.76

Panel D. Optimized parameter values for Model 2

Spec	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**	**11**
A*	14.19	14.19	14.20	14.19	14.19	14.19	14.19	14.19	14.19	14.29	14.19
B*	-0.05	-0.05	-0.06	-0.06	-0.05	-0.05	-0.05	-0.05	-0.05	-0.09	-0.05
t_c*	4.01	4.01	4.01	4.01	4.01	4.01	4.01	4.01	4.01	4.78	4.01
b*	0.90	0.90	0.90	0.90	0.90	0.90	0.90	0.90	0.90	0.90	0.90
C*	1.25	1.25	-1.15	-1.18	-1.25	1.25	1.25	-1.25	-1.25	0.25	1.25
ω*	5.00	5.00	5.00	5.00	5.00	5.00	5.00	5.00	5.00	14.98	5.00
φ*	1.39	1.39	4.54	4.53	4.53	7.67	7.67	10.81	10.81	-1.25	13.95
SSR	**2.33***	2.33	2.34	2.33	2.33	2.33	2.33	2.33	2.33	6.15	2.33

Table 21. Calibrating the power-law model using 12 months of data before the momentum bubble peak on April 16, 2020.

Panel A shows the initial value of the first-stage estimation approach $\Phi^0_{model\ 1} = (A, B, t_c, \beta)$, which is a basic power-law model, as $ln[comp(t)] = Ln[com(T)] - (t_c - t)^{\overline{\beta}}$. In this setting, $A = ln[p(T)]$ is the natural logarithm of the last compounded return in the sample. Given the daily frequency of the sample, one unit of time equals 1/253. In the second stage, the model is optimized by minimizing the *SSR* of $ln[comp(t)] = Ln[com(T)] - (t_c - t)^{\overline{\beta}}$ by allowing for the following constraints: $t_c \geq T + 1\ and\ 0.1 \leq \beta \leq 0.9$. Panel B reports the optimal values for the parameter vector $\Phi^*_{model\ 1} = (A^*, B^*, t_c^*, \beta^*)$. In the third stage, using the optimized parameter values of Model 1 in Panel B as the initial values $\Phi^0_{model\ 2} = (A^*, B^*, t_c^*, \beta^*, C, \omega, \Phi)$ and setting $C = \phi = 0$ for varying values of $5 \leq \omega \leq 15$, we calibrate the LPPLS model 2 that includes log-periodic oscillations as $log[E(p(t))] \approx A + B(t_c - t)^{\beta} - (1 + C\ cos[\omega\ log(t_c - t) + \varphi])$ by considering the following constraints: $t_c \geq T + 1,\ 0.1 \leq \beta \leq 0.9, and\ 5 \leq \omega \leq 15$. Panel C outlines the input data vectors $\Phi^0_{model\ 2} = (A^*, B^*, t_c^*, \beta^*, C, \omega, \Phi)$, while Panel D presents the optimized parameter vectors $\Phi^*_{model\ 2} = (A^{**}, B^{**}, t_c^{**}, \beta^{**}, C^{**}, \omega^{**}, \Phi^{**})$ for each iteration.

Panel A. Initial parameter values for Model 1

Spec	A	B	T_c=T+1	b	*SSR*
1	14.03	-1.00	3.23	0.20	**941.95***
2	14.03	-1.00	3.23	0.40	1151.64
3	14.03	-1.00	3.23	0.60	1491.06
4	14.03	-1.00	3.23	0.80	2007.13

Panel B. Optimized parameter values for Model 1

Spec	A*	B*	t_c^*	b*	*SSR*
1	14.88	-0.72	4.09	0.19	**7.43***
2	14.46	-0.34	3.81	0.34	7.48
3	14.36	-0.25	3.61	0.40	7.49
4	14.47	-0.36	3.63	0.29	7.45

Panel C. Initial parameter values for Model 2

Spec	1	2	3	4	5	6	7	8	9	10	11
A*	14.88	14.88	14.88	14.88	14.88	14.88	14.88	14.88	14.88	14.88	14.88
B*	-0.72	-0.72	-0.72	-0.72	-0.72	-0.72	-0.72	-0.72	-0.72	-0.72	-0.72
t_c*	4.09	4.09	4.09	4.09	4.09	4.09	4.09	4.09	4.09	4.09	4.09
b*	0.19	0.19	0.19	0.19	0.19	0.19	0.19	0.19	0.19	0.19	0.19
C*	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
ω*	5.00	6.00	7.00	8.00	9.00	10.00	11.00	12.00	13.00	14.00	15.00
Φ*	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
SSR	7.43	7.43	7.43	7.43	7.43	7.43	7.43	7.43	7.43	7.43	7.43

Panel D. Optimized parameter values for Model 2

Spec	1	2	3	4	5	6	7	8	9	10	11
A	14.08	14.08	14.08	14.07	14.07	14.11	14.12	14.12	14.17	14.17	14.17
B	-0.03	-0.03	-0.03	-0.02	-0.02	-0.01	-0.01	-0.01	-0.09	-0.09	-0.09
t_c**	9.10	9.10	9.09	9.10	9.10	4.31	4.29	4.29	3.23	3.23	3.23
b*	0.10	0.10	0.10	0.10	0.10	0.90	0.90	0.90	0.90	0.90	0.90
C*	3.98	4.78	4.40	-5.70	-5.69	-8.28	5.70	5.87	-0.49	-0.49	0.49
ω**	14.99	14.98	14.99	15.00	14.99	5.00	5.00	5.00	10.21	10.21	10.21
φ **	-24.84	-24.81	-24.81	-21.71	-21.71	3.86	7.05	7.04	7.14	7.14	10.28
SSR	2.04	2.04	2.04	2.04	2.04	**1.94***	1.96	1.96	5.18	5.18	5.18

Table 22. Testing the residuals of the LPPLS model.

This table presents the estimated test statistics and the corresponding *p*-values for various Augmented Dickey–Fuller (ADF) tests conducted on the residual series of optimized Model 2 predictions for each bubble event occurring on 1939.08.23, 2008.07.14, 2016.01.20, and 2020.04.16, the results of which are presented in Panels A, B, C, and D of Table 14, respectively. We perform the ADF test using the regression equation $\Delta e_t = \delta_0 + \delta_1 t + \delta_2 u_{t-1} + \gamma_1 \Delta u_{t-1} + \cdots + \gamma_p \Delta u_{t-p} + \varepsilon_t$, where Δe_t represents the residuals of the LPPLS model, defined as the difference between $Ln[comp(t)]$ and the predicted value from the LPPLS model, Eq. (12), at time t. The parameters $\gamma_1, \ldots, \gamma_p$ measure exposures to $\Delta u_{t-1}, \ldots, \Delta u_{t-p}$ and ε_t is assumed to be a white noise process. We set $\delta_0 = \delta_1 = 0$ when assuming no deterministic terms and $\delta_1 = 0$ when including only a constant in the model. Thus, the ADF test is structured in three primary forms with the null hypothesis that $\hat{\delta}_2 = 0$, which is contrary to the alternative hypothesis $\hat{\delta}_2 < 0$. The computed test statistic $\hat{\lambda}$ is the *t*-statistic that corresponds to the estimated coefficient $\hat{\delta}_2$. The critical values for the 5% and 1% significance levels are presented in columns 4 and 5, respectively. It is important to note that a lag order of $p = 2$ is utilized for all model specifications, as recommended by the Schwarz information criterion. *p*-values are reported in parentheses. The significance levels at 1% and 5% are indicated by *** and **, respectively.

Panel A. Bubble on December 21, 1929				
	Model specification	$\hat{\lambda}$	**Critical Values 5%**	**Critical Values 1%**
3 months (1926.11.03 to 1929.09.21)	**No deterministic terms**	-4.91*** (1.161E-06)	-1.95	-2.58
	Constant	-4.91*** (3.161E-05)	-2.86	-3.43
	Constant with trend	-4.90*** (2.83E-04)	-3.41	-3.96
6 months (1926.11.03 to 1929.06.21)	**No deterministic terms**	-4.53*** (6.778E-06)	-1.95	-2.58
	Constant	-4.53*** (1.741E-04)	-2.86	-3.43
	Constant with trend	-4.51*** (1.397E-03)	-3.41	-3.96
9 months (1926.11.03 to 1929.03.21)	**No deterministic terms**	-4.25*** (2.267E-05)	-1.95	-2.58
	Constant	-4.25*** (5.354E-04)	-2.86	-3.43
	Constant with trend	-4.25*** (3.665E-03)	-3.41	-3.96
12months (1926.11.03 to 1928.12.21)	**No deterministic terms**	-4.12*** (3.949E-005)	-1.95	-2.58
	Constant	-4.12*** (9.009E-04)	-2.86	-3.43
	Constant with trend	-4.12*** (5.896E-03)	-3.41	-3.96
Panel B. Bubble on March 10, 2000				
	Model specification	$\hat{\lambda}$	**Critical values 5%**	**Critical values 1%**
3 months (1939.11.29 to 1999.12.10)	**No deterministic terms**	-4.86*** (1.477E-06)	-1.95	-2.58
	Constant	-4.86*** (3.979e-05)	-2.86	-3.43
	Constant with trend	-4.86*** (3.40E-04)	-3.41	-3.96
6 months (1939.11.29 to 1999.09.10)	**No deterministic terms**	-4.74*** (2.549E-06)	-1.95	-2.58
	Constant	-4.74*** (6.755E-05)	-2.86	-3.43
	Constant with trend	-4.74*** (5.571E-04)	-3.41	-3.96
9 months	**No deterministic terms**	-4.56*** (5.756e-006)	-1.95	-2.58
	Constant	-4.56*** (1.477E-04)	-2.86	-3.43

(1939.11.29 to 1999.06.07)	**Constant with trend**	-4.56*** (1.147E-03)	-3.41	-3.96
12months (1939.11.29 to 1990.03.10)	**No deterministic terms**	-4.78*** (2.143e-006)	-1.95	-2.58
	Constant	-4.78*** (5.71e-005)	-2.86	-3.43
	Constant with trend	-4.78*** (4.767E-04)	-3.41	-3.96
Panel C. Bubble on July 14, 2008				
	Model specification	$\hat{\lambda}$	**Critical Values 5%**	**Critical Values 1%**
6 months (2002.09.30 to 2008.01.11)	**No deterministic terms**	-2.04** (0.0401)	-1.95	-2.58
	Constant	-2.03 (0.2723)	-2.86	-3.43
	Constant with trend	-1.98 (0.6112)	-3.41	-3.96
9 months (2002.09.30 to 2007.10.12)	**No deterministic terms**	-5.41*** (1.005E-07)	-1.95	-2.58
	Constant	-5.41*** (2.759E-06)	-2.86	-3.43
	Constant with trend	-5.41*** (2.62E-05)	-3.41	-3.96
12 months (2002.09.30 to 2007.07.14)	**No deterministic terms**	-6.01*** (4.426E-09)	-1.95	-2.58
	Constant	-6.01*** (1.125E-07)	-2.86	-3.43
	Constant with trend	-6.02*** (1.018E-06)	-3.41	-3.96
Panel D. Bubble on January 20, 2016				
	Model specification	$\hat{\lambda}$	**Critical values 5%**	**Critical values 1%**
3 months (2009.03.6 to 2015.10.20)	**No deterministic terms**	-8.07*** (3.157E-14)	-1.95	-2.58
	Constant	-8.09*** (2.714 E-13)	-2.86	-3.43
	Constant with trend	-8.20*** (4.134E-13)	-3.41	-3.96
6 months (2009.03.6 to 2015.07.20)	**No deterministic terms**	-32.30*** (9.67 E-30)	-1.95	-2.58
	Constant	-32.67*** (4.161E-43)	-2.86	-3.43
	Constant with trend	-32.76*** (7.649E-126)	-3.41	-3.96
9 months (2009.03.6 to 2015.4.20)	**No deterministic terms**	-12.24*** (1.879E-25)	-1.95	-2.58
	Constant	-12.25*** (1.634E-26)	-2.86	-3.43
	Constant with trend	-12.28*** (1.079E-29)	-3.41	-3.96
12 months (2009.03.6 to 2015.01.20)	**No deterministic terms**	-12.20*** (2.377E-25)	-1.95	-2.58
	Constant	-12.21*** (2.203E-26)	-2.86	-3.43
	Constant with trend	-12.24*** (1.539E-29)	-3.41	-3.96
Panel E. Bubble on April 16, 2020				
	Model specification	$\hat{\lambda}$	**Critical values 5%**	**Critical values 1%**
	No deterministic terms	-5.61*** (3.576E-08)	-1.95	-2.58

3 months (2016.01.20 to 2020.01.16)	**Constant**	-5.61*** (9.769E-07)	-2.86	-3.43
	Constant with trend	-6.28*** (2.341E-07)	-3.41	-3.96
6 months (2016.01.20 to 2019.10.16)	**No deterministic terms**	-6.01*** (4.491E-09)	-1.95	-2.58
	Constant	-6.00*** (1.169E-07)	-2.86	-3.43
	Constant with trend	-5.97*** (1.389E-06)	-3.41	-3.96
9 months (2016.01.20 to 2019.07.16)	**No deterministic terms**	-7.71*** (2.767E-13)	-1.95	-2.58
	Constant	-8.11*** (2.369E-13)	-2.86	-3.43
	Constant with trend	-8.00*** (2.011E-12)	-3.41	-3.96
12 months (2016.01.20 to 2019.04.16)	**No deterministic terms**	-6.12*** (2.478E-09)	-1.95	-2.58
	Constant	-6.11*** (6.336E-08)	-2.86	-3.43
	Constant with trend	-6.06*** (8.111E-07)	-3.41	-3.96

Table 23. Performance comparison.

This table compares the prediction results of the LPPLS model's forecasting capability for five separate momentum bubbles on December 21, 1929, March 10, 2000, July 14, 2008, January 20, 2016 and April 16, 2020, using data 3, 6, 9, and 12 months prior to the critical time.

Bubble peak date	**Prediction**			
	3 months	**6 months**	**9 months**	**12 months**
1929.12.21	+7M	+4M	-6M	+15M
2000.03.10	+4M	+3M	+1M	+1M
2008.07.14	---	-4M	-2M	-5M
2016.01.20	+9M	+16M	+13M	+19M
2020.04.16	-3M	-3M	-3M	+1M
Average	+3.4M	+3.2M	+0.6M	+6.2M

Appendix

Figure A. 1. Residuals of the calibrated LPPLS Model 2 for the momentum bubble peak on December 21, 1929.

This figure plots the residuals of Model 2, covering the in-sample time window for 3, 6, 9, and 12 months of data before the bubble on December 21, 1929.

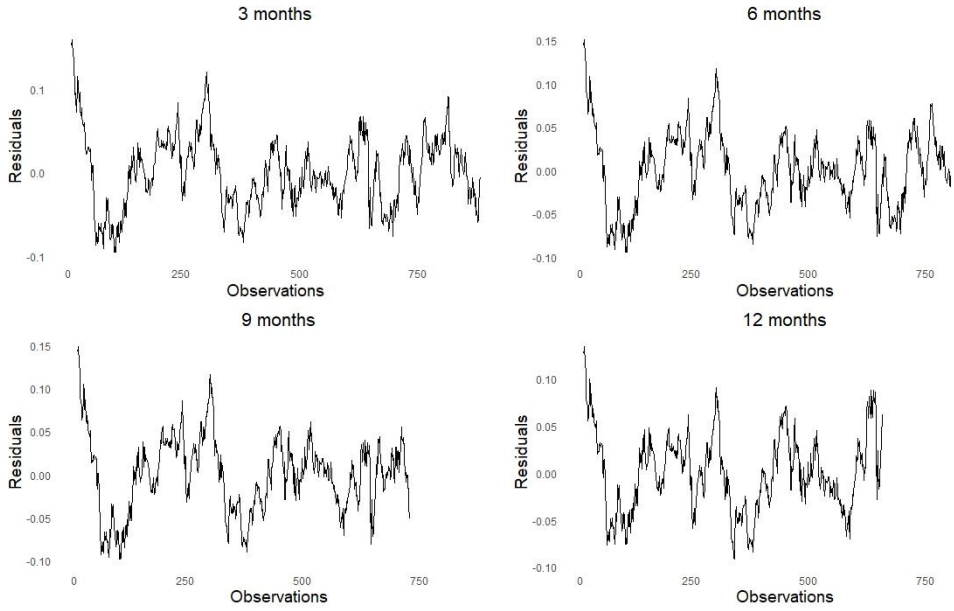

Figure A. 2. Residuals of the calibrated LPPLS Model 2 for the momentum bubble peak on March 10, 2000.

This figure plots the residuals of Model 2, covering the in-sample time window for 3, 6, 9, and 12 months of data before the bubble on March 10, 2000.

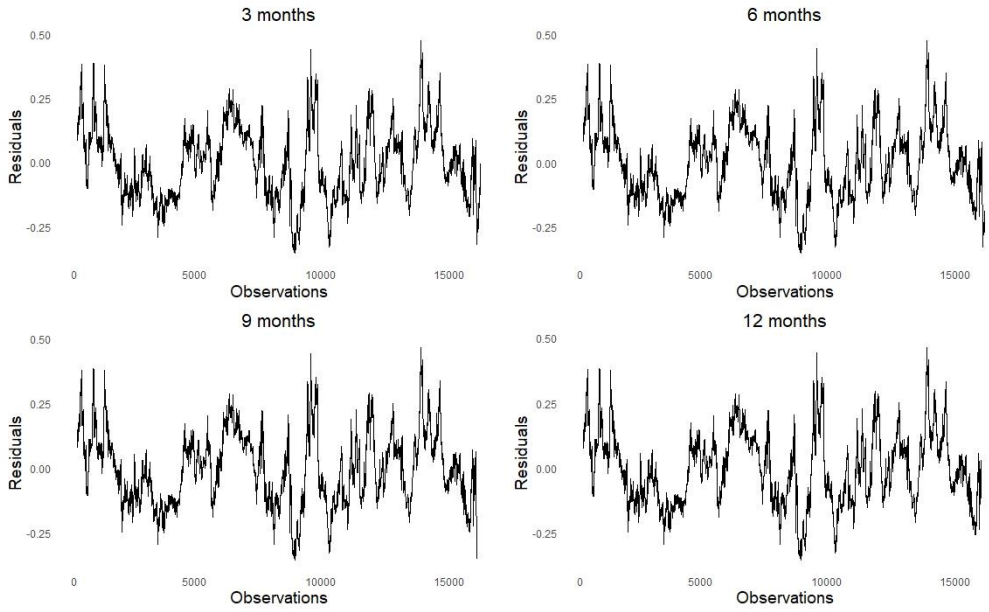

Figure A. 3. Residuals of the calibrated LPPLS Model 2 for the momentum bubble peak on July 14, 2008.

This figure plots the residuals of model 2, covering the in-sample time window for 6, 9, and 12 months of data before the bubble on July 14, 2008. (the 3 months calibration is excluded as it did not yield a prediction).

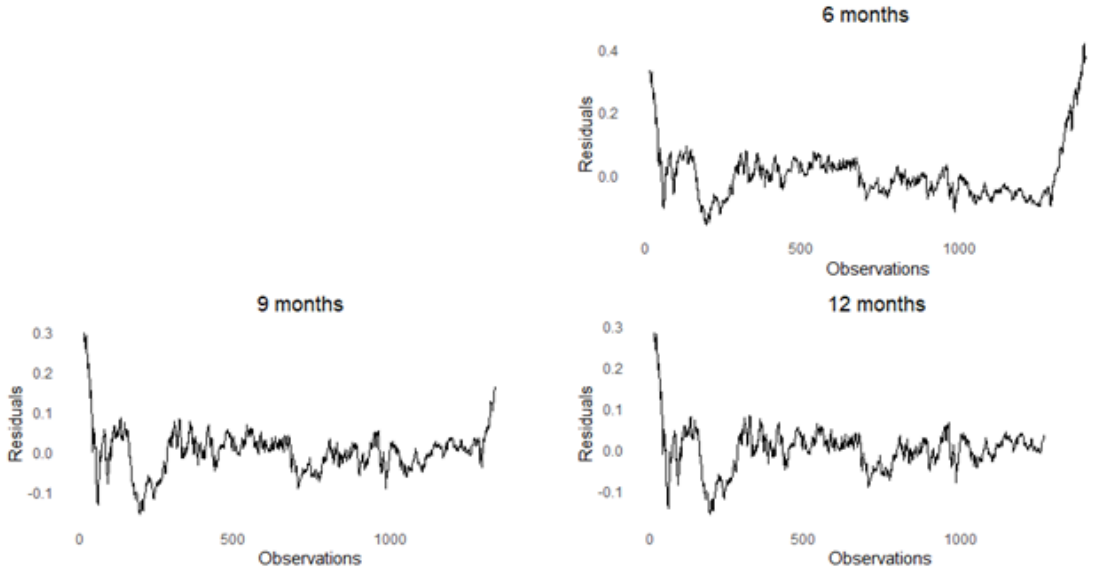

Figure A. 4. Residuals of the calibrated LPPLS Model 2 for the momentum bubble peak on January 20, 2016.

This figure plots the residuals of model 2, covering the in-sample time window for 3, 6, 9, and 12 months of data before the bubble on January 20, 2016.

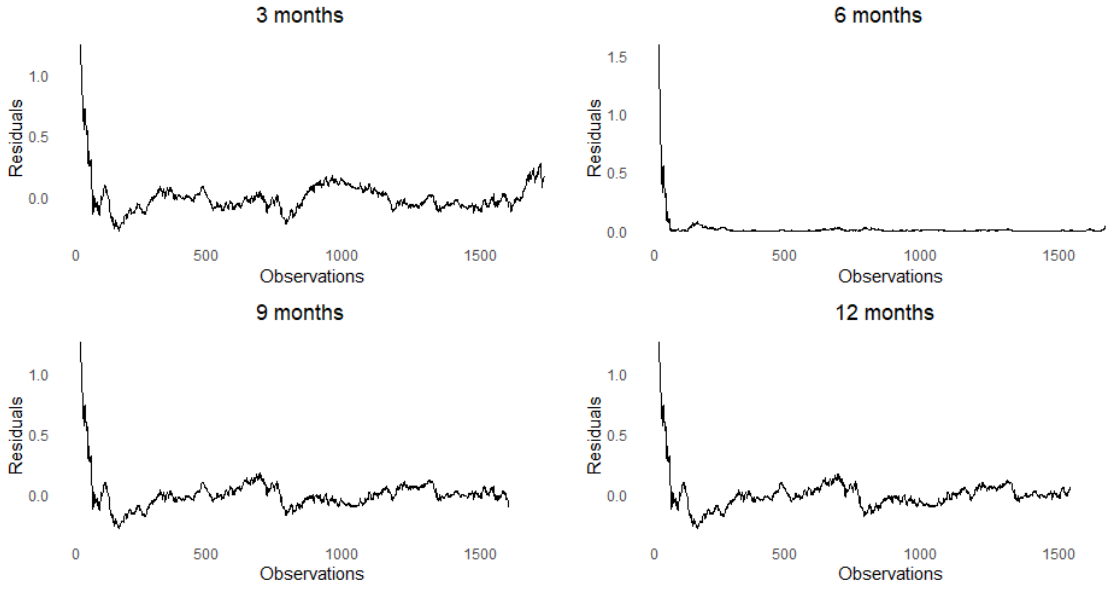

Figure A. 5. Residuals of the calibrated LPPLS Model 2 for the momentum bubble peak on April 16, 2020.

This figure plots the residuals of Model 2, covering the in-sample time window for 3, 6, 9, and 12 months of data before the bubble on April 16, 2020.