ECONOMETRICS
EKONOMETRIA
Advances in Applied Data Analysis

Econometrics. Ekonometria.
Advances in Applied Data Analysis
Year 2026, Vol. 30, No. 1
ISSN 2449-9994
journals.ue.wroc.pl/eada

Is Persistent Entropy Simply a Volatility Proxy? Evidence from the WIG20 Index

Stanisław M. Halkiewicz

AGH University of Krakow, Poland

e-mail: smsh@student.agh.edu.pl

ORCID: 0009-0000-7344-7522

Quote as: Halkiewicz, M. S. (2026). Is Persistent Entropy Simply a Volatility Proxy? Evidence from the WIG20 Index. *Econometrics. Ekonometria. Advances in Applied Data Analysis*, *30*(1), 1-20.

DOI: 10.15611/eada.2026.1.01

JEL: C58, G17, C45

Abstract

Aim: This article investigates whether persistent homology and persistence entropy capture structural properties of financial time series beyond variance-based risk measures. Using data from the WIG20 index (2019–2024), the study examines whether topological descriptors reflect intrinsic geometric and temporal organization rather than merely volatility intensity.

Methodology: Logarithmic returns are embedded using sliding-window delay coordinates and analysed with Vietoris–Rips persistent homology. Betti numbers, persistence diagrams and rolling H_1 persistence entropy are computed. Relationships with classical risk diagnostics are evaluated using linear correlations, nonlinear dependence measures, regime comparisons and shuffle-based tests.

Results: Persistence entropy shows weak association with volatility and second-moment risk. Stronger relationships appear with higher-order distributional characteristics such as skewness and kurtosis. Volatility-based regimes do not significantly separate entropy, whereas a structural split around the 2022 geopolitical shock reveals a significant increase, indicating a shift in return geometry. Shuffle experiments confirm dependence on temporal ordering.

Implications: Persistence entropy captures structural and temporal organization of financial returns and may complement classical econometric risk measures.

Originality/value: The study shows that persistent homology reflects structural organization of return dynamics rather than acting as a volatility proxy.

Keywords: topological data analysis, persistence homology, WIG20, market indices, volatility

1. Introduction

Financial markets exhibit complex, nonlinear and dynamically evolving behaviour that often challenges classical econometric modelling. While volatility, autocorrelation and tail-risk measures provide valuable statistical summaries, they primarily capture distributional or moment-based properties of returns. Such measures do not directly describe the geometric or structural organization of return trajectories in state space.

Topological Data Analysis (TDA) offers a complementary framework for studying data through its geometric and multiscale structure. Rather than focusing solely on local statistical characteristics, TDA examines connectivity, loops and higher-dimensional features that emerge when observations are viewed as points in a high-dimensional space. Persistent homology, the central tool of TDA, tracks how such topological features appear and disappear across scales, allowing structural organization to be quantified in a robust and scale-sensitive manner.

In financial research, persistent homology has been used primarily in the context of crash detection, bubble identification and systemic risk analysis (Gidea, 2017; Yen & Cheong, 2021; Rai et al., 2024). However, most of these applications topological measures are interpreted as indicators of rising instability or volatility clustering, yet an important foundational question remains underexplored.

Do topological descriptors simply re-express volatility in geometric language, or do they capture structural properties of financial time series that extend beyond classical variance-based risk measures?

The study addressed this question using the WIG20 index (2019–2024) as an empirical case. Rather than focusing exclusively on crash detection, the author investigated whether persistence entropy reflects intrinsic geometric and temporal organization of embedded return trajectories.

Methodologically, logarithmic returns were embedded using sliding-window delay coordinates, constructing Vietoris–Rips filtrations, computing persistent homology, and analysing Betti numbers, persistence diagrams, and rolling H_1 persistence entropy. Next its relationship to classical risk diagnostics was evaluated using linear and nonlinear dependence measures, regime comparisons, and surrogate shuffling experiments.

The contribution of the paper is twofold. First, it provides a detailed structural characterization of the embedded WIG20 return geometry. Second, and more importantly, it demonstrates empirically that persistence entropy is not merely a volatility proxy but encodes structural and temporal organization of returns.

The remainder of the paper is organized as follows. Section 2 presents the methodological framework. Section 3 reports empirical results. Section 4 discusses implications, limitations and directions for future research.

1.1. Applications of TDA

TDA has been successfully applied in image analysis, genomics, neuroscience, sensor networks, electronics, and dynamical systems (Carlsson, 2009; Ghrist, 2008; Zabaleta-Ortega et al., 2023; Zheng et al., 2024). More recently, researchers have explored its applications in time series analysis (Ravishanker & Chen, 2021; Ichinomiya, 2023; Chaudhari & Singh, 2024; de Jesus et al., 2025), particularly in economics and finance (e.g., Gidea, 2017; Gidea & Katz, 2018; Gidea et al., 2020; de Jesus et al., 2025), but also in health data (Zheng et al., 2024) and predictive maintenance (El Yaagoubi, Freyermuth & Ombao, 2025), among other domains.

1.1.1. TDA in Time Series Analysis

Time series data often exhibit complex dependencies and nonlinear behaviours, making them difficult to analyse using traditional models. TDA offers a novel approach by transforming time series into point

clouds via embedding techniques. One such approach is delay-coordinate embedding or sliding window embedding (Perea & Harer, 2015). This transformation allows the geometry of time series to be studied in high-dimensional spaces.

One of the earliest works applying persistent homology to time series is that of Perea and Harer (2015), in which the authors introduced sliding window embeddings and demonstrated how topological recurrence structures can be used to detect periodicity in dynamical systems. Other studies have shown that TDA can effectively characterise chaotic behaviour (Emrani et al., 2014) and detect regime changes in complex systems (Seversky et al., 2016).

Recently, another application of computational topology to time series has emerged, namely the use of topological invariants to detect regime change points (commonly referred to as changepoint analysis). Zheng et al. (2023) proposed PERCEPT, a novel online changepoint detection method that uses TDA as its technical backbone. Sugathadasa, Erfani and Leckie (2025) addressed the same task using topological smoothing, another method from computational topology, and developed the S-CPD framework. Yao et al. (2025) conducted a benchmark study of various changepoint detection methods in which they analysed 26 stocks over a 12-year period. Their results confirmed that TDA-based indicators consistently outperformed other approaches (in terms of the F1 score) in detecting "four major financial extreme events […]: the intensification of the European debt crisis in 2011, Brexit in 2016, the outbreak of the COVID-19 pandemic in 2020, and the energy crisis triggered by the Russia–Ukraine war in 2022" (Yao et al., 2025). El-Yaagoubi et al. (2025) introduced a robust method for changepoint detection in complex, non-stationary data in multi-trial experiments and validated its practical usefulness by correctly identifying bearing failures in NASA vibration signals.

An overview of TDA applications to time series data can be found in Chaudhari and Singh (2024), Leaverton (2020), and the thesis by Bois (2024). A gentle introduction to TDA concepts for statistical audiences is provided by El-Yaagoubi, Chung and Ombao (2023).

1.1.2. TDA for Market Data

The use of persistent homology for financial time series specifically is another growing area of research. One of the first studies to apply TDA in financial markets was by Gidea (2017) and Gidea & Katz (2018), who analysed stock indices using topological recurrence networks. Their results suggested that persistent homology could detect pre-crash signals in stock markets, therefore rendering a useful tool in risk avoidance and mitigation. Aguilar and Ensor (2020) applied the theory to US stock market indices and some ETF contracts' prices in order to detect changes in the market structure. In another study, Yen and Cheong (2021) used TDA to model volatility of stock markets in Singapore and Taiwan and found that a drop in persistency among homology groups was strongly related to market turbulence, as reflected in Betti numbers and persistent entropy. Similarly, Li et al. (2023) and Rai et al. (2024), applied Vietoris-Rips complexes to price series and demonstrated that market crashes coincide with an increase in topological complexity. One could say that detecting market crashes is currently the main area of use for TDA in financial and economic literature, as during financial crashes, the complexity of associate time-series rises, which corresponds with increase in β_0 and β_1 components (Yen & Cheong, 2021; Aguilar & Ensor, 2020). For example, Girish and Shahabudheen (2024) analysed the Indian stock market, aiming to detect crashes thereof. Rai et al. (2024) extended the usage of TDA to detecting widely understood extreme events on the market – which of course includes crashes.

While topological methods in financial analysis were on the rise, so were cryptocurrencies. This may be why numerous studies utilising TDA in finance focused specifically on crypto markets. First, Gidea et al. (2020) used topological tools developed in their previous works and applied them to time series generated by cryptocurrency markets. Islambekov et al. (2024) used topological graph theory for capturing anomalies in crypto transactions. Song and Li (2025) discussed topological transmissions as early warnings for increased volatility in crypto prices. A sound overview of topological approach to cryptocurrencies was provided by Rudkin et al. (2023).

Another interesting use of TDA in market data was proposed by Akingbade et al. (2024), who analysed the capability of using TDA for detecting financial bubbles. Manziy et al. (2024) discussed how irrationality of human behaviour fits into this framework and, as a result, incorporated a behavioural model of information.

1.2. Aim of the Research

The primary objective of this research was to examine whether topological descriptors derived from persistent homology capture structural properties of financial return dynamics that extend beyond traditional variance-based risk measures.

Rather than focusing solely on crash detection or volatility amplification, the author investigated whether persistence entropy reflects intrinsic geometric and temporal organization of the embedded return trajectory.

Specifically, the research aimed to:

- extract multiscale topological features from delay-embedded WIG20 logarithmic returns using Vietoris–Rips persistent homology,
- characterise the static geometric structure of the embedded return cloud via Betti numbers, persistence diagrams, and barcodes,
- analyse rolling H_1 persistence entropy as a time-varying measure of structural complexity,
- evaluate linear and nonlinear dependence between persistence entropy and classical risk diagnostics,
- assess whether entropy differentiates volatility-defined regimes,
- examine whether entropy responds to macro-structural shifts (pre and post-February 2022),
- test whether persistence entropy depends on temporal ordering through surrogate shuffling experiments.

To address these objectives, the following research questions were formulated:

Q1. Does the delay-embedded WIG20 return trajectory exhibit persistent multiscale cyclic structure, as reflected in H_1 features?

Q2. Is persistence entropy linearly or nonlinearly dependent on classical risk measures such as volatility, tail risk, or autocorrelation?

Q3. Does persistence entropy differentiate regimes defined purely by volatility magnitude?

Q4. Is persistence entropy sensitive to broader structural changes in market dynamics, such as the macroeconomic shift associated with February 2022?

Q5. Does persistence entropy depend materially on temporal ordering of returns, or is it primarily driven by static distributional characteristics?

This paper systematically addressed these questions and by quantitative, topological analyses performed on the WIG20 data, evaluated whether persistent homology provides structural information about financial markets that is not reducible to conventional econometric diagnostics.

2. Methods

2.1. Data Collection and Preprocessing

Let P_t denote the adjusted closing price of the WIG20 index at time t. Logarithmic returns are defined as:

$$r_t = \log P_t - \log P_{t-1}, \tag{1}$$

where $\log$ denotes the natural logarithm.

To reconstruct the underlying dynamics, the author applied delay-coordinate embedding. For embedding dimension m, each embedded vector was defined as

$$\mathbf{x}_t = (r_t, r_{t+1}, \dots, r_{t+m-1}) \in \mathbb{R}^m. \tag{2}$$

This produced a point cloud

$$X = \{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_N\}, \tag{3}$$

where N denotes the number of embedded vectors and indices $i, j \in \{1, \dots, N\}$ refer to elements of this point cloud. m is the embedding dimension, chosen based on Taken's (1981) theorem, which ensures that the original dynamics are preserved in a sufficiently high-dimensional space (Noakes, 1991; Robinson, 2010; Barański et al., 2024; Koltai & Kunde, 2024). Following previous works on financial time series, $m = 30$ was selected (Gidea, 2017; Seversky et al., 2016).

2.2. Vietoris-Rips Filtration

Let $\delta(\mathbf{x}_i, \mathbf{x}_j)$ denote the Euclidean distance between points

$$\delta(\mathbf{x}_i, \mathbf{x}_j) = \| \mathbf{x}_i - \mathbf{x}_j \|_2. \tag{4}$$

For filtration parameter $\varepsilon > 0$, the Vietoris–Rips complex was defined as

$$VR_\varepsilon(X) = \{\sigma \subset X : \delta(\mathbf{x}_i, \mathbf{x}_j) \leq \varepsilon \text{ for all } \mathbf{x}_i, \mathbf{x}_j \in \sigma\}. \tag{5}$$

This produced a nested sequence of complexes as ε increases.

2.3. Persistent Homology

Persistent homology tracks the birth and death of topological features across filtration scales.

Each feature is represented by a pair (b_ℓ, d_ℓ), where

- b_ℓ = birth scale
- d_ℓ = death scale

Lifetime of feature

$$\ell_\ell = d_\ell - b_\ell, \tag{6}$$

where index $\ell = 1, \dots, L$ runs over persistent features.

Persistence entropy is computed from lifetimes:

$$p_\ell = \frac{\ell_\ell}{\sum_{k=1}^{L} \ell_k} \tag{7}$$

$$H = -\sum_{\ell=1}^{L} p_\ell \log p_\ell. \tag{8}$$

2.4. Betti Numbers and Betti Curves

For each filtration scale $\varepsilon \geq 0$, the Vietoris–Rips complex $VR_\varepsilon(X)$ induces homology groups $H_k(VR_\varepsilon(X))$. The Betti number $\beta_k(\varepsilon)$ is defined as the rank (dimension) of the k-th homology group

$$\beta_k(\varepsilon) := \operatorname{rank} H_k(VR_\varepsilon(X)). \tag{9}$$

In practice, persistent homology is computed once for the entire filtration, producing a persistence diagram (or barcode) for each homological dimension k. For a fixed k, the diagram is a multiset of birth–death pairs

$$D_k = \{(b_\ell^{(k)}, d_\ell^{(k)})\}_{\ell=1}^{L_k}, \tag{10}$$

where $b_\ell^{(k)}$ and $d_\ell^{(k)}$ denote the birth and death filtration scales of the ℓ-th topological feature in dimension k, and L_k is the number of detected features.

The Betti number $\beta_k(\varepsilon)$ at scale ε can be recovered directly from the diagram as the number of intervals alive at ε, i.e. the number of features whose birth occurs at or before ε and whose death occurs after ε:

$$\beta_k(\varepsilon) = \#\left\{\ell \in \{1, \dots, L_k\} : b_\ell^{(k)} \leq \varepsilon < d_\ell^{(k)}\right\}. \tag{11}$$

Thus, $\varepsilon \mapsto \beta_k(\varepsilon)$ defines the Betti curve (sometimes called the Betti function) for dimension k. This study focused on:

- $\beta_0(\varepsilon)$, the number of connected components of $VR_\varepsilon(X)$,
- $\beta_1(\varepsilon)$, the number of independent one-dimensional loops.

For reporting purposes, Betti curves were evaluated on a discrete grid of filtration values $\{\varepsilon_1, \dots, \varepsilon_M\}$. The resulting table $\{(\varepsilon_m, \beta_0(\varepsilon_m), \beta_1(\varepsilon_m))\}_{m=1}^{M}$ provides a compact summary of how connectivity and loop structure evolve as the scale parameter increases.

In the context of delay-embedded financial returns, connected components (β_0) represent clusters of return states that remain disconnected at small filtration scales. As the scale increases, components merge, reflecting increasing structural coherence in state space.

Loops (β_1) correspond to cyclic geometric structures in the embedded trajectory. In a financial interpretation, such loops may reflect recurrent or oscillatory patterns in return dynamics, however these structures are purely geometric and should not be interpreted as deterministic cycles in price levels.

2.5. Computational Implementation

In this study, the giotto-tda Python library, developed by Tauzin et al. (2021), was employed to perform the required analyses, recommend the use of Python 3.12.12 and packages versions as follows:

- giotto-tda 0.6.2,
- numpy 1.24+,
- scipy 1.10+,
- scikit-learn 1.3+,
- matplotlib 3.7+,
- pandas 1.5+.

The daily adjusted closing prices (Adj Close) of the WIG20 index were obtained from Stooq.pl. The sample covered the period from January 2019 to December 2024. The dataset contains observations for trading days only. Weekends and bank holidays were naturally excluded, and no artificial calendar alignment was performed. No missing observations were present within the trading-day series. Logarithmic returns were computed using consecutive trading days without interpolation or forward filling.

3. Results

3.1. Static Topological Structure of Embedded Returns

The study began by analysing the persistent homology of the delay-embedded WIG20 log-return trajectory. Returns were embedded in dimension $m = 30$, and the Vietoris–Rips filtration was constructed using the Euclidean metric. Within each rolling window, returns were standardised (z-score normalisation), ensuring comparability of filtration scales across time.

3.1.1. Betti Numbers and Betti Curves

Table 1 reports representative Betti numbers for selected filtration values computed from the full-sample embedding, calculated as in Eq. 9.

Table 1. Betti numbers across filtration scales (full-sample embedding)

Filtration scale $\boldsymbol{\varepsilon}$	$\boldsymbol{\beta_0}$ *(components)*	$\boldsymbol{\beta_1}$ *(loops)*
4.426	572	62
4.5482	404	30
4.8537	297	30
5.1593	218	13
5.4649	172	9
…	…	0
13.1043	16	0
13.4099	11	0
13.7154	5	0
14.0210	1	0

Source: own calculations in Python, using giotto-tda library (Tauzin et al., 2021); data from Stooq.pl.

The β_0 profile shows a gradual decrease in the number of connected components as the filtration scale increases, consistent with progressive merging of components in Vietoris–Rips complexes. In contrast, the $\beta_1$profile is concentrated at the smallest filtration scales: loops are present only for ε near the lower end of the grid ($\beta_1 = 62$ at $\varepsilon \approx 4.24$, declining to $\beta_1 = 9$at $\varepsilon \approx 5.46$), after which $\beta_1$becomes identically zero over the remainder of the filtration range.

This pattern indicates that cyclic structure in the embedded return cloud is strictly local in scale: H_1features exist only at very small radii and disappear rapidly as ε increases. Consequently, there was no evidence of persistent loop structure at moderate or large filtration scales.

Figure 1 shows the corresponding Betti curves. The β_0 curve decreases smoothly, while the β_1 curve is sharply concentrated at the smallest filtration values and collapses to zero thereafter. Importantly, this behaviour reflects that loop-like geometry, if present, is confined to a narrow band of small scales.

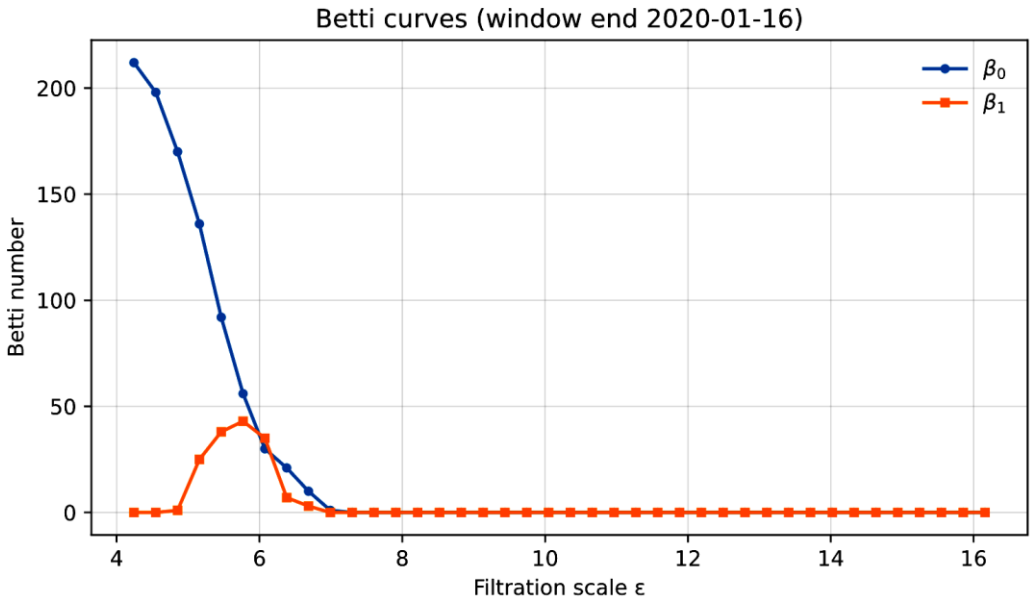

Fig. 1. Betti curves from full-sample embedding

Source: own calculations in Python, using giotto-tda library (Tauzin et al., 2021); data from Stooq.pl.

3.1.2. Persistence Diagram

The persistence diagram presented in Figure 2 supports the Betti-number evidence. H_0 features merge as ε grows, while H_1 features are predominantly short-lived and concentrated near the diagonal, consistent with loop structure that vanishes quickly when moving away from the smallest scales.

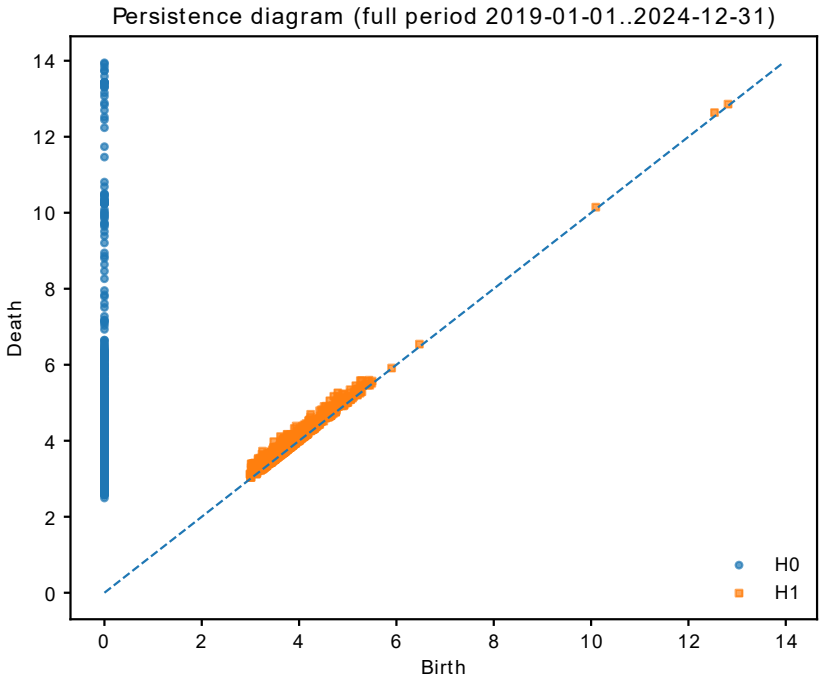

Fig. 2. Persistence diagram calculated from the full-sample embedding

Source: own calculations in Python, using giotto-tda library (Tauzin et al., 2021); data from Stooq.pl.

3.1.3. Persistence Barcode

The persistence diagrams represent each topological feature as a point (b, d) in the birth–death plane. The persistence barcode, on the other hand, provides an equivalent but often more interpretable representation. Formally, for each homological feature (either connected component or loop) with birth time $b$and death time d, the barcode represents the feature as a horizontal interval $[b, d)$ on the filtration axis. The length of the bar equals its lifetime $d - b$, and longer bars correspond to more persistent (and therefore potentially more structurally relevant) topological features.

The barcode presented in Figure 3 is illustrated with separate panels for H_0 and H_1:

- the upper panel corresponds to H_0 (connected components),
- the lower panel corresponds to H_1 (loops).

This separation is necessary because H_0 typically contains a large number of intervals (one per initial component at small scales), which would otherwise visually dominate the barcode and obscure the behaviour of H_1 features.

The H_0 panel shows many intervals born at small filtration values that terminate as components merge. As the filtration scale increases, these intervals progressively collapse until a single connected component remains.

The H_1 panel reveals that loop intervals occur only within a narrow range of small filtration values. In line with the Betti-number evidence reported in Table 1, H_1 bars are short-lived and disappear rapidly as ε increases and no long-persistent H_1 intervals are observed at moderate or large scales.

By analysing Figure 3, one can reinforce the conclusion that cyclic structures in the embedded WIG20 return trajectory are localised and scale-limited rather than globally persistent.

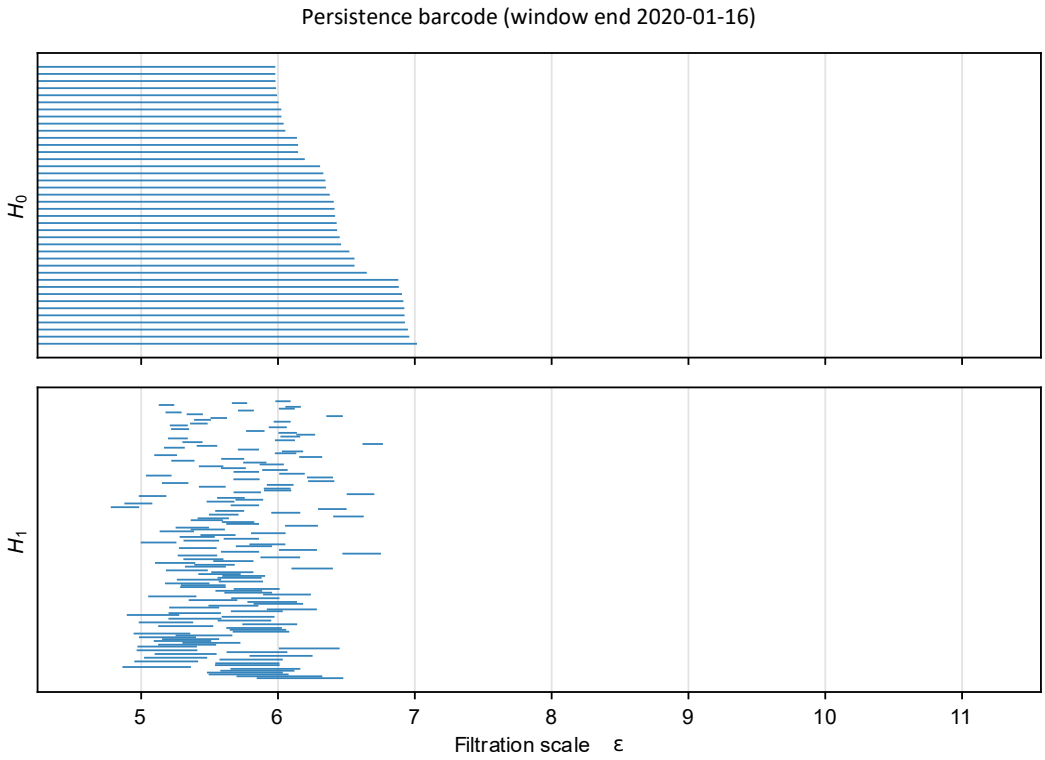

Fig. 3. Persistence barcode on the full-sample embedding, split by homology dimension

Source: own calculations in Python, using giotto-tda library (Tauzin et al., 2021); data from Stooq.pl.

3.2. Rolling Persistence Entropy Dynamics

Compute H_1 (for further analysis, H_0 is economically not interesting) persistence entropy within rolling windows of 252 trading days with a stride of 5 days. Figure 4 presents rolling H_1 entropy together with rolling volatility.

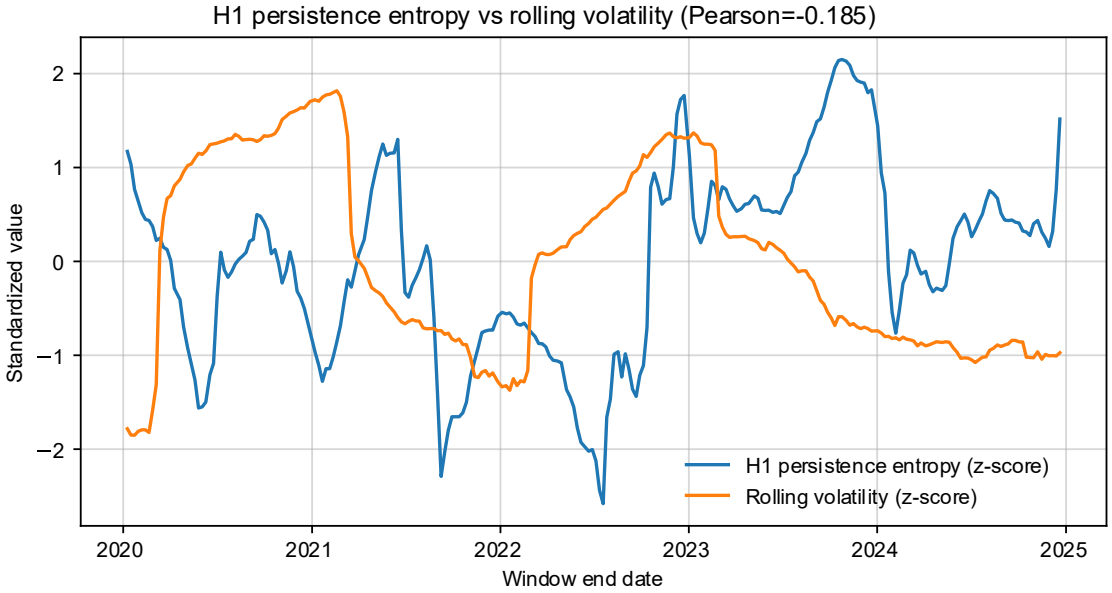

Fig. 4. Comparison of rolling volatility and H_1 persistence entropy in the researched period

Source: own calculations in Python, using giotto-tda library (Tauzin et al., 2021); data from Stooq.pl.

The entropy values fluctuate over time instead of following a deterministic upward or downward trend. Periods of increased entropy alternate with periods of lower entropy. The Pearson correlation between entropy and rolling volatility was calculated as $\rho = -0.19$ (with $n = 250$). This weak (negative) association indicates that persistence entropy is not a simple (linear) transformation of volatility. It is also difficult to deduce from the plot any other trivial relationship that could describe the relationship of these two metrics.

3.3. Diagnostic Comparison

3.3.1. Linear Dependence

To evaluate whether persistence entropy captures information already embedded in standard financial risk measures, the author computed both Pearson and Spearman correlations between rolling H_1persistence entropy and a range of classical diagnostics. In addition to raw correlations, correlations were reported after removing the linear effect of rolling volatility (entropy residualised with respect to volatility). Table 2 shows the computed coefficients.

Table 2. Computed correlation coefficients between persistence entropy and selected standard risk metrics

Metric	Pearson (raw)	Spearman (raw)	Pearson (resid.)	Spearman (resid.)
Rolling volatility	-0.1854	-0.1374	2.8345e^-17	0.0662
Realised variance	-0.1877	-0.1385	-0.0031	0.0652
Downside semivariance	-0.2416	-0.1607	-0.0627	0.0344
Skewness	0.3576	0.3586	0.2227	0.2062
Excess kurtosis	-0.4372	-0.5434	-0.3209	-0.4327
5% VaR	0.1430	0.0917	-0.0385	-0.1049
Expected shortfall at 5%	0.3085	0.2789	0.1371	0.0886
First order autocorrelation	0.2729	0.2292	0.3014	0.2776

Source: own calculations performed in Python, using giotto-tda library (Tauzin et al., 2021); data from Stooq.pl.

Thus several patterns can be seen. First, the linear association between persistence entropy and volatility is weak (Pearson -0.1854), and vanishes entirely once entropy is residualised with respect to volatility. This confirms that persistence entropy is not a simple transformation of second-moment risk.

Second, entropy exhibits moderate association with higher-order distributional characteristics. In particular, excess kurtosis shows the strongest relationship (Pearson -0.4372; Spearman -0.5434). Even after removing the linear effect of volatility, a substantial association remains (Pearson -0.3209; Spearman -0.4327). This suggests that entropy is sensitive to tail thickness and distributional shape rather than merely dispersion.

Third, skewness and expected shortfall display moderate positive correlations with entropy. The relationship weakens after volatility adjustment but remains non-negligible. This could indicate partial independence from second-moment effects.

Fourth, persistence entropy shows moderate association with first-order autocorrelation, and this relationship slightly strengthens after volatility residualisation (Pearson 0.3014). This indicates that entropy may reflect temporal dependence structure in returns.

Finally, risk measures focused on lower-tail quantiles (VaR at 5%) show only weak association, particularly after volatility adjustment.

Overall, persistence entropy is weakly related to volatility, moderately related to higher-order moment structure (especially kurtosis), and partially associated with serial dependence. Therefore, since the correlation is not straightforwardly linearly dependent on a standard risk metric, it can be concluded that it captures structural characteristics of the embedded return geometry that extend beyond traditional variance-based statistics.

3.3.2. Nonlinear Dependence

In order to deepen the analysis, the author also checked for nonlinear dependencies, computing mutual information (MI) and distance correlation (dCor) between rolling H_1 persistence entropy and selected diagnostics. Mutual information captures general statistical dependence without assuming linearity, while distance correlation equals zero if and only if the variables are independent, thus also detecting nonlinear dependence structures.

Table 3. Computed nonlinear dependencies between persistence entropy and selected standard risk metrics.

Metric	Mutual information	Distance correlation
Volatility	0.8160	0.2401
Realised variance	0.8031	0.2396
Downside semivariance	0.5942	0.2811
Skewness	0.7394	0.4296
Excess kurtosis	0.7914	0.4851
5% VaR	0.6353	0.2201
Expected shortfall at 5%	0.7225	0.3460
First order autocorrelation	0.3371	0.2886

Source: own calculations performed in Python, using giotto-tda library (Tauzin et al., 2021); data from Stooq.pl.

First, mutual information values are relatively high across most diagnostics. However, MI is sensitive to marginal distributions and does not directly quantify strength of functional dependence. Distance correlation therefore provides a more interpretable measure of nonlinear association.

Second, distance correlations between persistence entropy and standard dispersion measures (volatility and realised variance) are modest (approximately 0.24). This confirms the earlier conclusion from linear analysis: persistence entropy is not strongly driven by second-moment risk.

Third, the strongest nonlinear association is observed with excess kurtosis ($dCor \approx 0.49$) and skewness ($dCor \approx 0.43$). This is consistent with the linear findings and suggests that persistence entropy is particularly sensitive to distributional shape and tail behaviour rather than variance alone.

Fourth, nonlinear dependence with downside risk measures (Expected Shortfall, VaR, semivariance) remains moderate but not dominant (distance correlations between 0.22 and 0.35), indicating partial, however not deterministic, linkage with tail risk.

Finally, dependence with first-order autocorrelation ($dCor \approx 0.29$) suggests that persistence entropy may encode information related to short-term serial dependence in returns.

Nonlinear diagnostics reinforce the earlier conclusion: persistence entropy is not merely a nonlinear transformation of volatility, as it appears to reflect higher-order structural properties of the return distribution (particularly kurtosis and skewness) along with elements of temporal dependence.

Linking with the prior linear results, the evidence suggests that persistence entropy captures aspects of the geometric organization of embedded returns that extend beyond classical variance-based and tail-risk statistics.

3.4. Regime Comparison

3.4.1. Volatility-based Regimes

To investigate whether persistence entropy differs systematically across market regimes, the author first considered a classification based on volatility intensity. Rolling windows were partitioned according to quartiles of rolling volatility:

- **high-volatility regime**: top quartile of rolling volatility,
- **low-volatility regime**: bottom quartile of rolling volatility.

For each regime the mean persistence entropy was computed and compared using both parametric (Welch t-test) and nonparametric (Mann–Whitney U test) procedures (West, 2021).

Table 4. Mean persistence entropy by volatility regime

Regime	Mean entropy
High-volatility	6.8468
Low-volatility	6.7733

Source: own calculations based on rolling-window estimates; data from Stooq.pl.

The mean entropy is slightly lower during high-volatility periods, with a difference of approximately -0.0735.

Table 5. Statistical tests for equality of entropy distributions across regimes

Statistic	Value
Welch t-statistic	-0.9209
T-test p-value	0.3590
Mann–Whitney p-value	0.2495
Cohen's d	-0.1641

Source: own calculations based on rolling-window estimates; data from Stooq.pl.

Both the Welch t-test and the Mann–Whitney test failed to reject equality of means (or distributions) at conventional significance levels. The estimated effect size ($\text{Cohen's d} = -0.1641$) indicates a small magnitude difference (Hedges, 2024).

Thus, while persistence entropy exhibits a mild tendency to decrease during high-volatility periods, the evidence does not support statistically robust regime separation when regimes are defined purely by volatility magnitude. This result reinforces the earlier conclusion that persistence entropy is not primarily determined by second-moment risk levels.

3.4.2. Structural Regime Split: Pre and Post-2022

Since persistence entropy is intended to capture the intrinsic geometric organization of embedded return trajectories, it was natural to examine whether it responds to major structural changes in the economic environment. In contrast to volatility-based regime classification which focuses on the magnitude of short-term fluctuations, a structural split allows to assess whether broader macroeconomic transformations alter the underlying geometry of returns.

The Polish economy experienced a pronounced structural shift following February 2022, when Russia launched its full-scale invasion of Ukraine. For Poland, this event marked a transition from relatively stable and expansionary economic conditions to a regime characterised by geopolitical uncertainty, elevated inflationary pressure, energy-market disruptions, and a large-scale inflow of refugees. This transition did not correspond to a single, well-defined stock market crash, but rather it represented a broader reorganization of economic conditions that plausibly affected corporate expectations, investors' sentiments, sectoral dynamics, fiscal and monetary responses, and cross-border capital flows.

If persistence entropy reflects structural properties of return trajectories rather than merely volatility intensity (which these conclusions support so far), one would expect it to be sensitive to such a macro-level transformation.

Accordingly, rolling windows were divided into:

- **pre-2022 period** (before February 2022),
- **post-2022 period** (February 2022 onward).

The results are reported in Tables 6 and 7.

Table 6. Persistence entropy before and after the Russian full-scale invasion of Ukraine

Period	N	Mean entropy	Standard deviation
Pre-invasion	108	6.6321	0.4788
Post-invasion	142	6.9677	0.6526

Source: own calculations based on rolling-window estimates; data from Stooq.pl.

The post-2022 period exhibits a higher average persistence entropy, with a difference of approximately 0.336.

Table 7. Statistical tests for structural regime comparison

Statistic	Value
Welch t-statistic	-4.6896
T-test p-value	4.53×10^{-6}
Mann–Whitney p-value	7.16×10^{-7}
Cohen's d	−0.575

Source: own calculations based on rolling-window estimates; data from Stooq.pl.

Both the parametric and nonparametric tests indicate a statistically significant difference between the two structural periods. The estimated effect size ($|d| \approx 0.58$) suggests a moderate magnitude separation.

Unlike the volatility-based regime classification, which showed no robust separation, the structural split indicated a clear shift in the geometric complexity of the embedded WIG20 return trajectory. This suggests that persistence entropy is responsive to broader structural transformations in market dynamics rather than merely to fluctuations in volatility, which further supports the study's findings. For a complete view see Figure 5 presenting the temporal trajectory of $H1$ persistent homology, with the time split in February 2022.

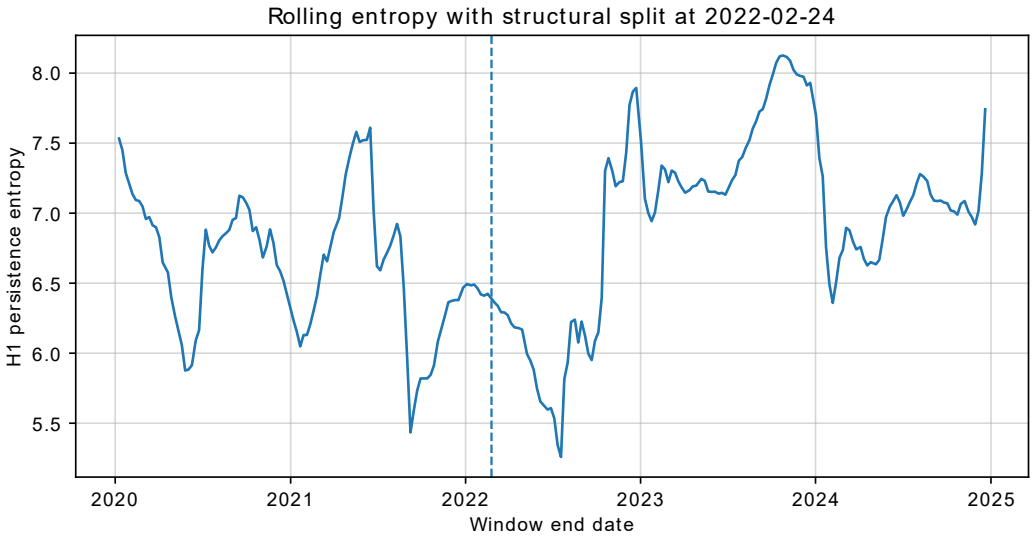

Fig. 5. Rolling H1 persistence entropy with structural split in February 2022

Source: own calculation based on rolling-window estimates; data from Stooq.pl.

3.5. Surrogate Test of Dependence on Temporal Organization

The previous analyses suggest that persistence entropy is not primarily driven by volatility levels and is only partially related to higher-order distributional characteristics. However, an important conceptual question remains: does persistence entropy merely reflect static properties of the return distribution, or does it depend on the temporal organization of returns?

Persistent homology is applied to delay-embedded time series, meaning that geometric structure emerges from the trajectory of returns in phase space. If entropy were determined solely by marginal distributional features (such as variance or kurtosis), then randomly permuting returns within each window, thereby preserving the distribution but destroying temporal dependence, should not materially affect entropy values. Conversely, if persistence entropy reflects dynamic organization, shuffling should alter the geometric structure and therefore change entropy.

To test this, the author performed a shuffle-based surrogate experiment. Within each selected rolling window, the return sequence was randomly permuted. This preserves:

- the unconditional distribution of returns,
- volatility magnitude,
- skewness and kurtosis,

but eliminates:

- serial dependence,
- volatility clustering,
- any higher-order temporal structure.

Persistence entropy was then recomputed for each shuffled window, and compared with entropy from the original ordering.

The results are reported in Table 8.

Table 8. Original vs shuffled persistence entropy

Statistic	Value
Number of windows	60
Mean entropy (original)	6.8267
Mean entropy (shuffled)	7.4192
Mean difference (orig − shuf)	−0.5926
Paired t-statistic	−5.06
T-test p-value	5.06×10^{-6}
Wilcoxon p-value	2.90×10^{-5}
Cohen's d (paired)	−0.6481

Source: own calculations based on rolling-window estimates; data from Stooq.pl.

Both parametric and nonparametric tests indicate a statistically significant difference between entropy computed on original returns and entropy computed on shuffled surrogates. The estimated effect size ($|d| \approx 0.65$)corresponds to a moderate-to-large magnitude difference (Hedges, 2024).

Entropy was also systematically higher for the shuffled series. Since shuffling preserves the unconditional distribution of returns but eliminates temporal dependence, this result implies that persistence entropy depends materially on return ordering. Temporal organization appears to constrain the geometric complexity of the embedded trajectory, leading to lower entropy in the original data relative to its randomly permuted counterpart.

Figure 6 illustrates the standardised entropy time series for the original and shuffled windows.

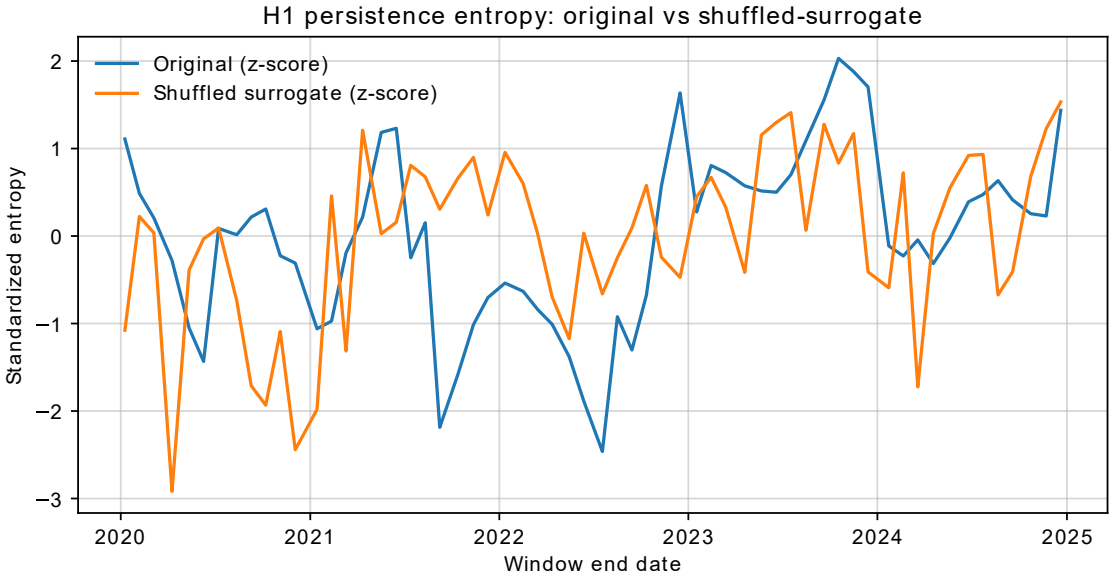

Fig. 6. H_1 persistence entropy for original vs. shuffled timeseries

Source: own calculations based on rolling-window estimates; data from Stooq.pl.

Figure 7 presents the distribution of paired entropy differences.

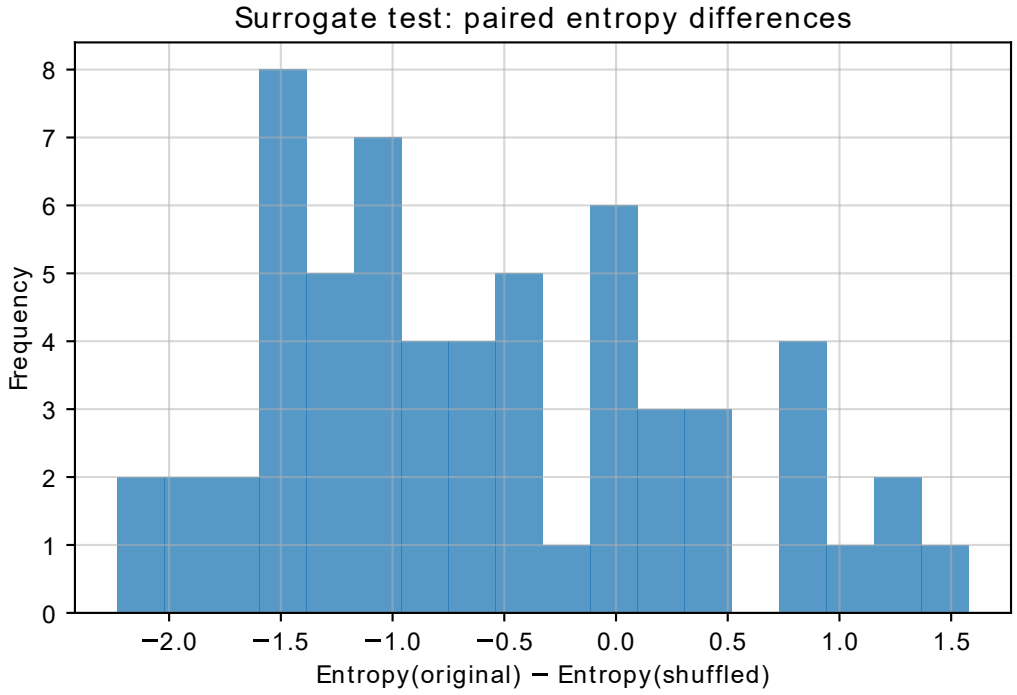

Fig. 7. Paired H_1 persistence entropy differences distribution

Source: own calculations based on rolling-window estimates; data from Stooq.pl.

The predominance of negative differences ($\text{original} < \text{shuffled}$) also confirms that destroying temporal structure systematically alters the geometric organization, which persistent homology is designed to capture. The presented surrogate experiment once again confirms that persistent entropy

is fundamentally and structurally different from traditional risk measures, as it is dependent structurally on temporal structure, the intrinsic geometry of the series etc., which – for example – volatility does not take into account at all.

4. Discussion and Conclusions

This study examined whether persistent homology and persistence entropy capture the structural properties of financial return dynamics that extend beyond traditional variance-based risk measures. The findings can be interpreted directly in light of the research questions posed earlier.

Q1. Does the delay-embedded WIG20 return trajectory exhibit persistent multiscale cyclic structure?

The static topological analysis indicates that H_1 features (loops) are present only at very small filtration scales and are mainly short-lived. No long-persistent multiscale cyclic structure was detected. This suggests that the embedded WIG20 return cloud does not exhibit globally persistent cyclic organization. Instead, loop structures appear local and scale-limited, therefore persistent homology does not reveal strong global periodic geometry in the investigated period.

Q2. Is persistence entropy dependent on classical risk measures?

Both linear and nonlinear diagnostics show that persistence entropy is only weakly related to volatility and realised variance. After residualising entropy with respect to volatility, linear associations vanish almost entirely. Nonlinear dependence, as measured by distance correlation, remains modest.

Stronger associations are observed with higher-order distributional characteristics, particularly kurtosis and skewness, and moderate association appears with first-order autocorrelation. However, these relationships remain partial rather than deterministic. Thus, persistence entropy is not reducible to a linear or nonlinear transformation of second-moment risk.

Q3. Does entropy differentiate volatility-defined regimes?

When regimes are defined purely by volatility quartiles, entropy does not display statistically significant separation. Parametric and nonparametric tests failed to reject equality of entropy distributions across both high and low-volatility regimes.

This result is crucial: persistence entropy does not simply increase during volatile periods. It does not behave as an alternative volatility proxy. This directly supports the central thesis of the paper.

Q4. Is entropy sensitive to structural macroeconomic shifts?

In contrast to volatility-based classification, the structural split at February 2022 reveals a statistically significant increase in entropy in the post-2022 period. The effect size is moderate, and both parametric and nonparametric tests confirm the difference.

This suggests that persistence entropy responds to broader structural transformations in market organization rather than to short-term fluctuation intensity. The macroeconomic shift associated with geopolitical disruption appears to have altered the geometric organization of embedded returns.

Q5. Does entropy depend on temporal ordering?

The surrogate shuffle experiment provides the strongest structural evidence. By randomly permuting returns within rolling windows, the study preserved marginal distributions while destroying temporal dependence.

Entropy systematically increased under shuffling, and the difference is statistically significant with moderate-to-large effect size. This demonstrates that persistence entropy depends materially on temporal organization of returns, and it is not merely a function of distributional shape.

4.1. Overall Conclusions

Taken holistically, the findings indicate that (H_1) persistence entropy captures mainly intrinsic geometric and temporal organization of embedded return trajectories. In order to answer the hypothesis from the title of this article – it is neither a linear nor nonlinear proxy for volatility (or any other standard risk measure, in fact). It does not reflect dispersion nor tail risk, but instead it appears to encode how returns are arranged in phase space, which is a much richer conclusion. This sensitivity becomes especially evident under macro-level regime shifts and temporal disruption. From a financial viewpoint, this may suggest that topological descriptors have the ability to complement traditional econometric risk measures (instead of replacing them) by providing information about the structure of market dynamics not captured by variance-based statistics.

4.2. Limitations and Future Research Directions

While the application of persistent homology to the WIG20 index has yielded structurally meaningful insights, several limitations should be acknowledged.

First, interpretability remains a central challenge. Although persistent homology provides mathematically rigorous summaries of geometric structure, the economic meaning of specific topological features, such as short-lived loops or fluctuations in persistence entropy, is not always straightforward. Unlike volatility, autocorrelation, or tail risk measures, which have well-established interpretations in financial economics, topological descriptors operate at a more abstract level of geometric organization. Translating geometric complexity into economically interpretable mechanisms remains an open problem.

Second, computational complexity constitutes a practical limitation. Persistent homology of high-dimensional embeddings, especially in rolling-window settings, is resource intensive. Although recent algorithmic advances have improved scalability, large-scale or real-time applications remain challenging. Future work may explore approximation techniques, sparse filtrations, or streaming implementations of persistent homology to facilitate broader empirical use.

Third, the presented study is structural and diagnostic in nature. While it demonstrates that persistence entropy is not a volatility proxy, responds to macro-structural shifts, and depends materially on temporal ordering, it does not assess predictive performance. Whether topological features can enhance forecasting models (for example, in volatility prediction, regime classification, or systemic risk monitoring) remains an open research direction. Integrating persistent homology with econometric or machine learning frameworks could clarify its incremental predictive value.

Finally, the empirical scope is limited to a single index and a defined time span. Extending the analysis to cross-market comparisons, sectoral indices or multi-asset systems, could reveal whether the structural properties identified by the study are specific to WIG20 or reflect broader market regularities.

Then, there is the problem of interpretability – looking at any time series from the geometric view is still relatively new, compared to the rich accomplishments of econometrics, statistics and quantitative finance. As it is mainly mathematicians who develop TDA, there are more tools to recover the insights from the data than ways of interpreting the formulated insights. This is where economists can play a relevant and vital role.

References

Aguilar, A. & Ensor, K. (2020). Topology Data Analysis Using Mean Persistence Landscapes in Financial Crashes. *Journal of Mathematical Finance*, *10*(4), 648-678. https://doi.org/10.4236/jmf.2020.104038

Akingbade, S. W., Gidea, M., Manzi, M., & Nateghi, V. (2024). Why topological data analysis detects financial bubbles? *Communications in Nonlinear Science and Numerical Simulation*, *128*, 107665. https://doi.org/10.1016/j.cnsns.2023.107665

Barański, K., Gutman, Y., & Śpiewak, A. (2024). Prediction of dynamical systems from time-delayed measurements with self-intersections. *Journal de Mathématiques Pures et Appliquées*, *186*, 103-149. https://doi.org/10.1016/j.matpur.2024.04.001

Bois, A. (2024). *Topological data analysis for time series* (Doctoral dissertation, Université Paris-Saclay). HAL Thèses. https://theses.hal.science/tel-05005755

Carlsson, G. (2009). Topology and data. *Bulletin of the American Mathematical Society* 46(2), 255–308. http://dx.doi.org/10.1090/S0273-0979-09-01249-X

Chaudhari, S., & Singh, S. K. (2024). A Review on Topological Data Analysis in Time Series. In Mahajan, V., Chowdhury, A., Singh, S.N., Shahidehpour, M. (eds) *Emerging Technologies in Electrical Engineering for Reliable Green Intelligence. ICSTACE 2023*. Lecture Notes in Electrical Engineering, vol 1117. Springer, Singapore. https://doi.org/10.1007/978-981-99-9235-5_36

Chazal, F., de Silva, V., Glisse, M., & Oudot, S. (2016). The Structure and Stability of Persistence Modules. *SpringerBriefs in Mathematics*. Springer, Cham. https://doi.org/10.1007/978-3-319-42545-0

de Jesus, L. C., Fernández-Navarro, F., & Carbonero-Ruz, M. (2025). Enhancing financial time series forecasting through topological data analysis. *Neural Computing & Applications,* 37, 6527-6545. https://doi.org/10.1007/s00521-024-10787-x

El-Yaagoubi, A. B., Chung, M. K., & Ombao, H. (2023). Topological Data Analysis for Multivariate Time Series Data. Entropy, 25(11), 1509. https://doi.org/10.3390/e25111509

El-Yaagoubi, A., Freyermuth, J.-M., & Ombao, H. (2025). A Robust Topological Framework for Detecting Regime Changes in Multi-Trial Experiments With Application to Predictive Maintenance. *Journal of Time Series Analysis*, Early View, 1–19. https://doi.org/10.1111/jtsa.70032.

Emrani, S., Gentimis, T., & Krim, H. (2014). Persistent-homology of delay embeddings and its application to wheeze detection. *IEEE Signal Processing Letters*, *21*(4), 459-463. https://doi.org/10.1109/LSP.2014.2305700

Ghrist, R. (2008). Barcodes: The persistent topology of data. Bulletin of the American Mathematical Society, 45(1), 61–75. http://dx.doi.org/10.1090/S0273-0979-07-01191-3

Gidea, M. (2017). Topological data analysis of critical transitions in financial networks. In E. Shmueli, B. Barzel, & R. Puzis (Eds.) *3rd International Winter School and Conference on Network Science (NetSci-X 2017)*. Springer Proceedings in Complexity. Springer. https://doi.org/10.1007/978-3-319-55471-6_5

Gidea, M., Goldsmith, D., Katz, Y., Roldan, P., & Shmalo, Y. (2020). Topological recognition of critical transitions in time series of cryptocurrencies. *Physica A: Statistical Mechanics and its Applications 548*, 123843. https://doi.org/10.1016/j.physa.2019.123843

Gidea, M., & Katz, Y. (2018). Topological data analysis of financial time series: Landscapes of crashes. Physica A: Statistical Mechanics and its Applications 491, 820–834. https://doi.org/10.1016/j.physa.2017.09.028

Girish, N., & Shahabudheen, A. (2024). Analysing Indian Stock Market Crashes Through a Topological Lens. In B. B. Mallik, K. Deyasi, S. Das, S. Ghosh, & S. Jana(Eds.) *Proceedings of the Fifth International Conference on Emerging Trends in Mathematical Sciences & Computing (IEMSC-24)*. Information Systems Engineering and Management, vol 10. Springer. https://doi.org/10.1007/978-3-031-71125-1_33

Hedges, L. V. (2024). Interpretation of the Standardized Mean Difference Effect Size When Distributions Are Not Normal or Homoscedastic. *Educational and Psychological Measurement*, *85*(2), 245-257. https://doi.org/10.1177/00131644241278928

Ichinomiya, T. (2023). Time series analysis using persistent homology of distance matrix. *Nonlinear Theory and Its Applications, IEICE* 14(2), 79–91. https://doi.org/10.1587/nolta.14.79

Islambekov, U., Pathirana, H., Khormali, O., Akçora, C., & Smirnova, E. (2024). A topological approach for capturing high-order interactions in graph data with applications to anomaly detection in time-varying cryptocurrency transaction graphs. *Foundations of Data Science 6*(4), 492–513. https://doi.org/10.3934/fods.2024024

Koltai, P., & Kunde, P. (2024). A Koopman–Takens theorem: Linear least squares prediction of nonlinear time series. *Communications in Mathematical Physics, 405*(5), 120. https://doi.org/10.1007/s00220-024-05004-8

Leaverton, L. A. (2020). *Analysis of Financial Time Series using TDA: Theoretical and Empirical Results* [Thesis, Universitat de Barcelona]. https://diposit.ub.edu/items/5f76bb5f-e836-4120-95cc-ae96201e3d3d

Li, Y., Li, X., Zhang, Y., & Cao, Y. (2023). A topological based feature extraction method for the stock market. *Data Science in Finance and Economics 3*(3), 208–229. https://doi.org/10.3934/DSFE.2023013

Manziy, O., Senyk, A., Sheremeta, D., & Senyk, Y. (2024). Information aspects of the manifestation of irrational influence of financial bubbles on financial markets. *Galician Economic Journal, 91*, 20-28. http://dx.doi.org/10.33108/galicianvisnyk_tntu2024.06.020

Noakes, L. (1991). The Takens Embedding Theorem. *International Journal of Bifurcation and Chaos, 1*(4), 867-872. https://doi.org/10.1142/S0218127491000634

Perea, J.A., Harer, J. (2015) Sliding Windows and Persistence: An Application of Topological Methods to Signal Analysis. *Foundations of Computational Mathematics, 15,* 799-838. https://doi.org/10.1007/s10208-014-9206-z

Rai, A., Nath Sharma, B., Rabindrajit Luwang, S., Nurujjaman, M., & Majhi, S. (2024). Identifying extreme events in the stock market: A topological data analysis. *Chaos: An Interdisciplinary Journal of Nonlinear Science 34*(10), 103106. https://doi.org/10.1063/5.0220424

Ravishanker, N., & Chen, R. (2021). An introduction to persistent homology for time series. *WIREs Computational Statistics.* 13, e1548. https://doi.org/10.1002/wics.1548

Robinson, J. C. (2010). The Takens Time-Delay Embedding Theorem.*Dimensions, Embeddings, and Attractors*. Cambridge University Press. https://doi.org/10.1017/CBO9780511933912

Rudkin, S., Rudkin, W., & Dłotko, P. (2023). On the topology of cryptocurrency markets. *International Review of Financial Analysis, 89*, 102759. https://doi.org/10.1016/j.irfa.2023.102759

Seversky, L. M., Davis, C., & Berger, M. (2016). On time-series topological data analysis: New data and opportunities. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, 59–67. https://doi.org/10.1109/CVPRW.2016.131

Song, S., & Li, H. (2025). Can topological transitions in cryptocurrency systems serve as early warning signals for extreme fluctuations in traditional markets? *Physica A: Statistical Mechanics and its Applications, 657*, 130194. https://doi.org/10.1016/j.physa.2024.130194

Sugathadasa, H., Erfani, S.M., Leckie, C. (2025). *S-CPD*: Topological Smoothing-Based Change Point Detection. In: Wu, X., et al. Data Science: Foundations and Applications. PAKDD 2025. Lecture Notes in Computer Science(), vol 15875. Springer, Singapore. https://doi.org/10.1007/978-981-96-8295-9_10

Takens, F. (1981). Detecting strange attractors in turbulence. In D. Rand, & L. S. Young (Eds.), *Dynamical Systems and Turbulence, Warwick 1980*. Lecture Notes in Mathematics, Vol. 898, pp. 366-381. Springer. https://doi.org/10.1007/BFb0091924

Tauzin, G., Lupo, U., Tunstall, L., Pasrez, J. B., Caorsi, M., Medina-Mardones, A. M., Dassatti, A., & Hess, K. (2021). giotto-tda: A topological data analysis toolkit for machine learning and data exploration. *Journal of Machine Learning Research*, *22*(39), 1-6. http://jmlr.org/papers/v22/20-325.html

West, R. M. (2021). Best practice in statistics: Use the Welch t-test when testing the difference between two groups. *Annals of Clinical Biochemistry: International Journal of Laboratory Medicine*, *58*(4), 267-269. https://doi.org/10.1177/0004563221992088

Yao, J., Li, J., Wu, J., Yang, M., & Wang, X. (2025). Change Point Detection in Financial Market Using Topological Data Analysis. *Systems*, 13(10), 875. https://doi.org/10.3390/systems13100875

Yen, P. T.-W., & Cheong, S. A. (2021). Using topological data analysis (TDA) and persistent homology to analyze the stock markets in Singapore and Taiwan. *Frontiers in Physics*, 9. https://doi.org/10.3389/fphy.2021.572216

Zabaleta-Ortega, A., Masoller, C., & Guzmán-Vargas, L. (2023). Topological data analysis of the synchronization of a network of Rössler chaotic electronic oscillators. *Chaos: An Interdisciplinary Journal of Nonlinear Science* 3(11), 113110. https://doi.org/10.1063/5.0167523

Zheng, J., Feng, Z., & Ekstrom, A. D. (2024). Towards analysis of multivariate time series using topological data analysis. *Mathematics 12*(11). https://doi.org/10.3390/math12111727

Zheng, X., Mak, S., Xie, L., & Xie, Y. (2023). PERCEPT: A New Online Change-Point Detection Method using Topological Data Analysis. *Technometrics*, *65*(2), 162–178. https://doi.org/10.1080/00401706.2022.2124312

Czy trwała entropia sprowadza się do zmienności? Wnioski z analizy indeksu WIG20

Streszczenie

Cel: Celem artykułu jest zbadanie, czy trwała homologia i entropia trwała uchwytują strukturalne właściwości szeregów czasowych wykraczające poza klasyczne miary ryzyka oparte na wariancji. Na przykładzie indeksu WIG20 (2019–2024) analizuje się, czy deskryptory topologiczne odzwierciedlają wewnętrzną organizację geometryczną i czasową stóp zwrotu, a nie jedynie poziom zmienności.

Metodyka: Logarytmiczne stopy zwrotu osadzono metodą okna przesuwnego, a następnie przeanalizowano z wykorzystaniem homologii trwałej Vietorisa–Ripsa i obliczono deskryptory topologiczne. Zależności między entropią a klasycznymi miarami ryzyka oceniono przy użyciu miar zależności liniowej i nieliniowej, analizy reżimów oraz testów z losowym przetasowaniem obserwacji.

Wyniki: Stwierdzono słabą zależność entropii trwałej od zmienności i miar opartych na drugich momentach. Silniejsze związki zaobserwowano z wyższymi momentami rozkładu, zwłaszcza skośnością i kurtozą. Podział próby według poziomu zmienności nie prowadzi do istotnej separacji entropii, natomiast podział związany z szokiem geopolitycznym w 2022 roku ujawnia istotny wzrost entropii. Testy z przetasowaniem potwierdzają jej zależność od porządku czasowego danych.

Implikacje: Entropia trwała opisuje strukturalną i czasową organizację stóp zwrotu i może stanowić uzupełnienie klasycznych miar ekonometrycznych ryzyka.

Oryginalność/Wartość: Wykazano, że trwała homologia odzwierciedla strukturę dynamiki rynku, a nie jedynie poziom zmienności.

Słowa kluczowe: topologiczna analiza danych, homologia trwała, WIG20, indeksy rynkowe