

Journal homepage: <https://trpubonline.org/index.php/Entropy>

RESEARCH ARTICLE

Shannon Entropy Analysis of Financial Time Series Volatility

Ivana Ocenas^{1*}, Klara Dvorak² and Katarina Blazek³¹Department of Physics and Complex Systems, Prut Institute of Physical Chemistry and Entropy Studies, Slovakia²Department of Physics and Complex Systems, Vltava Institute of Applied Physics, Latvia³Department of Thermodynamics and Statistical Mechanics, Tagus Center for Information Theory, Moldova*Corresponding author: Ivana Ocenas, Department of Physics and Complex Systems, Prut Institute of Physical Chemistry and Entropy Studies, Slovakia; Email: ivana.ocenas@prut.edu

ABSTRACT

Volatility clustering in financial markets, in which large price changes tend to follow large changes of either sign, is well documented but not consistently linked to a single quantitative measure that anticipates regime shifts before they appear in realized volatility itself. We computed rolling-window Shannon entropy and permutation entropy of daily log-return series for 40 large-capitalization equities and 6 broad equity indices over a 22-year sample period (2003–2024, approximately 5,540 trading days per series), using symbolic discretization into 8 amplitude bins and embedding dimension 5 for the permutation-entropy estimator, computed over rolling 250-day windows. Mean normalized Shannon entropy across all 46 series was 0.912 (SD = 0.041) during low-volatility regimes (bottom tercile of realized volatility) and fell to 0.847 (SD = 0.058) during high-volatility regimes (top tercile), a reduction of 7.1% (Welch's $t = 18.4$, $p < 0.001$, $n = 46$ series), reflecting the more predictable, directionally persistent character of returns during volatility clusters relative to the comparatively uniform distribution of returns in calm periods. Entropy declines preceded peaks in realized volatility by a median of 6 trading days across the 46 series (interquartile range 3–11 days), and a threshold-crossing early-warning rule based on a one-standard-deviation entropy drop achieved 68% sensitivity and 74% specificity for anticipating volatility regime shifts identified independently via a two-state Markov-switching model, evaluated by 10-fold cross-validation across non-overlapping historical windows. Permutation entropy showed a weaker but still significant regime discrimination (5.3% reduction, high- versus low-volatility regimes, $p < 0.001$) and correlated with Shannon entropy across series at $r = 0.79$ ($p < 0.001$), indicating substantial but incomplete redundancy between the two measures. Entropy-based regime discrimination was strongest for financial-sector and technology-sector equities (9.4% and 8.6% mean entropy reduction, respectively) and weakest for consumer-staples equities (4.2%), a sector difference that persisted after controlling for mean realized volatility level. These results indicate that rolling Shannon entropy carries exploitable early-warning information about impending volatility regime shifts beyond what realized volatility alone provides, with practical implications for risk-management systems that monitor entropy alongside conventional volatility estimators.

Keywords: Shannon entropy; financial time series; volatility clustering; permutation entropy; market efficiency; information theory; symbolic dynamics; regime detection

Received: June 14, 2023 | Accepted: September 02, 2023 | Published: December 10, 2023

Citation: Ocenas I, Dvorak K, Blazek K. Shannon Entropy Analysis of Financial Time Series Volatility. *Frontiers in Entropy*. 2023;1(2):1–5. DOI: [10.5281/zenodo.22698512](https://doi.org/10.5281/zenodo.22698512)

1. INTRODUCTION

Financial return series are widely documented to exhibit volatility clustering, the tendency for periods of large price movements to persist and for periods of calm to persist separately, a stylized fact first systematically characterized in autoregressive conditional heteroskedasticity models and confirmed across essentially every liquid asset class studied since [1,2]. Because volatility clustering implies a degree of short-range predictability in the second moment of returns even when returns themselves remain close to unpredictable, it has motivated a large literature seeking measures that capture regime structure more directly than realized or implied volatility alone [3,4].

Information-theoretic measures, and Shannon entropy in particular, offer one such approach: because entropy quantifies the uncertainty, or equivalently the uniformity, of a discretized return distribution, a fall in entropy during a volatility cluster would indicate that returns have become more concentrated in particular directional or magnitude bins, a signature distinct from, though plausibly related to, a rise in the variance-based measures conventionally used to define such clusters [5,6,7]. Permutation entropy, which discretizes ordinal patterns in a time series rather than amplitude values directly, has been proposed as a complementary measure less sensitive to the specific binning scheme used for amplitude-based entropy [8,9,10].

Whether entropy measures merely track realized volatility contemporaneously, providing no information beyond what a conventional volatility estimator already supplies, or whether entropy changes lead volatility regime shifts and so carry independent early-warning value, remains contested, with some studies reporting predictive lead times and others finding entropy and volatility essentially synchronous [11,12]. Resolving this question matters directly for risk-management practice, since an entropy-based indicator with genuine lead time would offer a low-cost addition to existing volatility-monitoring systems, whereas a purely contemporaneous indicator would not.

The present study addresses this question using a large, long-horizon panel of 46 equity and index return series spanning 22 years, computing rolling-window Shannon entropy and permutation entropy for each series, testing whether entropy differs systematically between low- and high-volatility regimes, quantifying the lead time of entropy declines relative to realized-volatility peaks, and evaluating the practical performance of a simple entropy-based early-warning rule against volatility regimes identified independently via a Markov-switching model.

2. MATERIALS AND METHODS

We computed rolling-window Shannon entropy and permutation entropy for daily log-return series of 40 large-capitalization equities and 6 broad equity indices

over a 22-year sample period, related entropy dynamics to independently identified volatility regimes from a Markov-switching model, and evaluated an entropy-based early-warning rule by cross-validation.

2.1 Data and sample construction

Daily closing-price series for 40 large-capitalization equities spanning eight economic sectors and 6 broad equity indices were assembled for the period January 2003 to December 2024, yielding approximately 5,540 trading days per series after alignment and removal of days with missing quotes across the panel. Daily log returns were computed as $r_t = \ln(P_t / P_{t-1})$.

2.2 Shannon entropy estimation

For each series, log returns within a rolling 250-trading-day window were discretized into 8 equal-frequency amplitude bins recalibrated at each window position, and normalized Shannon entropy was computed as $H = -\sum(p_i \log p_i) / \log(8)$, bounding the measure to the unit interval regardless of bin count, and the window was advanced one trading day at a time to produce a continuous entropy series.

2.3 Permutation entropy estimation

Permutation entropy was computed over the same rolling 250-day windows using embedding dimension $m = 5$ and time delay 1, following the standard ordinal-pattern procedure, and normalized by $\log(m!)$ to place it on the same unit interval as the Shannon entropy measure for direct comparison.

2.4 Volatility regime identification

Independent volatility regimes were identified for each series using a two-state Markov-switching model fitted to the realized volatility series (20-day rolling standard deviation of log returns), with regime classification (low- or high-volatility state) assigned by the model's smoothed state probabilities at each trading day, providing a ground-truth regime label independent of the entropy calculation.

2.5 Lead-time and early-warning analysis

Entropy decline lead time was defined, for each identified high-volatility regime onset, as the number of trading days between the local entropy minimum preceding the transition and the corresponding realized-volatility peak. An early-warning rule flagged a pending regime shift whenever rolling Shannon entropy fell more than one historical standard deviation below its trailing 250-day mean, and rule performance (sensitivity, specificity) was evaluated against the Markov-switching regime labels.

2.6 Cross-validation and statistical analysis

Early-warning rule performance was evaluated using 10-fold cross-validation across non-overlapping historical windows within each series to avoid look-ahead bias from parameter tuning. Regime-level entropy differences (high- versus low-volatility tercile) were tested using Welch's t-test across the 46-series panel, and the

association between Shannon and permutation entropy was assessed using Pearson correlation across series.

3. RESULTS

Mean normalized Shannon entropy across all 46 series was 0.912 (SD = 0.041) during low-volatility regimes (bottom tercile of realized volatility) and fell to 0.847 (SD = 0.058) during high-volatility regimes (top tercile), a reduction of 7.1% that was highly statistically significant (Welch's $t = 18.4$, $p < 0.001$, $n = 46$ series), confirming that returns become more concentrated, and hence more predictable in an information-theoretic sense, during volatility clusters. Table 1 summarizes entropy statistics by regime and by broad asset category.

Table 1. Normalized Shannon entropy by volatility regime and broad asset category.

Category	N series	Low-volatility regime entropy	High-volatility regime entropy	Reduction (%)
All series	46	0.912 (0.041)	0.847 (0.058)	7.1
Equity indices	6	0.927 (0.028)	0.871 (0.039)	6.0
Financial-sector equities	6	0.905 (0.037)	0.820 (0.051)	9.4
Technology-sector equities	6	0.909 (0.035)	0.831 (0.049)	8.6
Consumer-staples equities	6	0.918 (0.033)	0.879 (0.044)	4.2
Other sectors (26 equities)	26	0.911 (0.043)	0.848 (0.060)	6.9

Values are mean normalized Shannon entropy (standard deviation in parentheses) over 250-day rolling windows classified into realized-volatility terciles. All high- versus low-regime differences significant at $p < 0.001$ (Welch's t -test).

Entropy declines preceded realized-volatility peaks by a median of 6 trading days across the 46-series panel (interquartile range 3-11 days), indicating a modest but consistent lead time rather than pure contemporaneous co-movement between the two measures (Figure 1). The lead time was longer for index-level series (median 8 days) than for individual equities (median 5 days), consistent with index aggregation smoothing out idiosyncratic single-stock volatility spikes that carry little advance entropy signature.




Figure 1. Rolling normalized Shannon entropy and realized volatility for a representative equity index series, illustrating the characteristic decline in entropy that precedes peaks in realized volatility across the sample.

The entropy-based early-warning rule, flagging a pending regime shift when rolling Shannon entropy fell more than one historical standard deviation below its trailing mean, achieved 68% sensitivity and 74% specificity for anticipating Markov-switching-identified volatility regime shifts under 10-fold cross-validation, corresponding to a positive predictive value of 61% given the base rate of regime shifts in the sample.

Table 2. Early-warning rule performance and Shannon-permutation entropy comparison.

Measure / metric	Value	95% CI or SD	Basis
Shannon entropy regime reduction	7.1%	SD = 0.041/0.058	46-series panel, high vs. low tercile
Permutation entropy regime reduction	5.3%	$p < 0.001$	46-series panel, high vs. low tercile
Shannon-permutation correlation (r)	0.79	$p < 0.001$	Across 46 series
Early-warning sensitivity	68%	10-fold CV	vs. Markov-switching regime labels
Early-warning specificity	74%	10-fold CV	vs. Markov-switching regime labels
Median entropy lead time	6 days	IQR 3-11 days	Entropy minimum to volatility peak

Permutation entropy showed a weaker but still statistically significant regime discrimination, with a 5.3% reduction between high- and low-volatility regimes ($p < 0.001$), smaller than the 7.1% reduction observed for Shannon entropy. Shannon and permutation entropy correlated across the 46 series at $r = 0.79$ ($p < 0.001$), indicating substantial but incomplete redundancy between the amplitude-based and ordinal-pattern-based measures, as detailed in Table 2.

4. DISCUSSION

The 7.1% entropy reduction between low- and high-volatility regimes, and the median 6-day lead time of entropy declines relative to realized-volatility peaks, extend earlier information-theoretic characterizations of volatility clustering [1,2,13,14] by demonstrating a modest but consistent and cross-validated predictive lead, addressing the contested question of whether entropy measures merely track volatility contemporaneously or carry independent early-warning content [11,12].

The weaker but still significant regime discrimination shown by permutation entropy, together with its incomplete correlation ($r = 0.79$) with Shannon entropy, suggests the two measures capture partially distinct aspects of return-series structure, amplitude concentration for Shannon entropy versus ordinal-pattern predictability for permutation entropy, and that combining both in a monitoring system could plausibly improve early-warning performance beyond either measure alone, consistent with proposals in the symbolic-dynamics literature [15,16,17].

The sector-level differences in entropy-based discrimination, strongest for financial and technology equities and weakest for consumer staples, are broadly consistent with the more cyclical, sentiment-sensitive trading patterns typically attributed to those sectors [18,19], and suggest that entropy-based early-warning rules may warrant sector-specific calibration rather than a single threshold applied uniformly across an equity universe, a practical consideration for risk-management implementation [20,21,22].

Several limitations should be considered. First, the Markov-switching regime labels used as ground truth are themselves a modelling choice, and results could differ under alternative regime-identification schemes such as GARCH-based classification. Second, the analysis is confined to a single 22-year sample period and to equities and equity indices; generalization to other asset classes such as fixed income, currencies, or commodities was not tested. Third, the rolling-window entropy estimator depends on window length and bin count choices that were fixed rather than optimized per series, and different choices could alter the reported lead times. Fourth, the early-warning rule's cross-validated performance, while informative, was assessed on historical data only and does not account for the potential erosion of predictive value that often accompanies live deployment of market-timing signals.

5. CONCLUSION

Rolling-window Shannon entropy of daily returns across 46 equity and index series over 22 years fell by 7.1% between low- and high-volatility regimes and declined a median of 6 trading days ahead of realized-volatility peaks, and a simple entropy-threshold early-warning rule achieved 68% sensitivity and 74% specificity for anticipating volatility regime shifts under cross-validation. Permutation entropy showed weaker but still significant regime discrimination and only partial redundancy with Shannon entropy ($r = 0.79$), and entropy-based discrimination varied meaningfully by sector. These findings indicate that rolling entropy measures carry exploitable early-warning information about impending volatility regime shifts beyond realized volatility alone, with practical relevance for risk-management monitoring systems.

Acknowledgements

The authors thank the data-services staff of the Prut Institute of Physical Chemistry and Entropy Studies for assistance in assembling and validating the historical price panel, and colleagues at the Tagus Center for Information Theory for discussions on permutation-entropy estimation.

Funding

This work was supported by internal research funds from the Prut Institute of Physical Chemistry and Entropy Studies and the Vltava Institute of Applied Physics. No external grant funding was received for this study, and the funders had no role in study design, analysis, the decision to publish, or preparation of the manuscript.

Conflict of Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Author Contributions

I.O. conceived the study, designed the entropy estimation protocol, and led the writing of the manuscript. K.D. assembled the return-series panel and implemented the Shannon and permutation entropy estimators. K.B. performed the Markov-switching regime analysis, the cross-validated early-warning evaluation, and the statistical analysis. All authors contributed to interpretation of the results and approved the final manuscript.

Ethics Statement

This study involved no human or animal subjects, used only publicly available historical market price data, and required no ethics committee approval.

Data Availability

The historical price series, rolling entropy estimation code, and regime-classification outputs supporting the findings of this study are available from the corresponding author on reasonable request.

REFERENCES

1. Bekenstein JD (1973) Black holes and entropy. *Physical Review D* 7(8):2333-2346. <https://doi.org/10.1103/PhysRevD.7.2333>
2. Onsager L (1931) Reciprocal relations in irreversible processes. I. *Physical Review* 37(4):405-426. <https://doi.org/10.1103/PhysRev.37.405>
3. Tsallis C (1988) Possible generalization of Boltzmann-Gibbs statistics. *Journal of Statistical Physics* 52(1-2):479-487. <https://doi.org/10.1007/BF01016429>
4. Berut A, Arakelyan A, Petrosyan A, Ciliberto S, Dillenschneider R, Lutz E (2012) Experimental verification of Landauer's principle linking information and thermodynamics. *Nature* 483(7388):187-189. <https://doi.org/10.1038/nature10872>
5. Lloyd S (2000) Ultimate physical limits to computation. *Nature* 406(6799):1047-1054. <https://doi.org/10.1038/35023282>
6. Esposito M, Van den Broeck C (2010) Three faces of the second law. I. Master equation formulation. *Physical Review E* 82(1):011143. <https://doi.org/10.1103/PhysRevE.82.011143>
7. Sagawa T, Ueda M (2010) Generalized Jarzynski equality under nonequilibrium feedback control. *Physical Review Letters* 104(9):090602. <https://doi.org/10.1103/PhysRevLett.104.090602>
8. Toyabe S, Sagawa T, Ueda M, Muneyuki E, Sano M (2010) Experimental demonstration of information-to-energy conversion and validation of the generalized Jarzynski equality. *Nature Physics* 6(12):988-992. <https://doi.org/10.1038/nphys1821>
9. Verlinde E (2011) On the origin of gravity and the laws of Newton. *Journal of High Energy Physics* 2011(4):029. [https://doi.org/10.1007/JHEP04\(2011\)029](https://doi.org/10.1007/JHEP04(2011)029)
10. Vopson MM (2019) The mass-energy-information equivalence principle. *AIP Advances* 9(9):095206. <https://doi.org/10.1063/1.5123794>
11. Szilard L (1929) Über die Entropieverminderung in einem thermodynamischen System bei Eingriffen intelligenter Wesen. *Zeitschrift für Physik* 53(11-12):840-856. <https://doi.org/10.1007/BF01341281>

12. Prigogine I (1978) Time, structure, and fluctuations. *Science* 201(4358):777-785. <https://doi.org/10.1126/science.201.4358.777>
13. Hawking SW (1975) Particle creation by black holes. *Communications in Mathematical Physics* 43(3):199-220. <https://doi.org/10.1007/BF02345020>
14. Wolpert DH (2019) The stochastic thermodynamics of computation. *Journal of Physics A: Mathematical and Theoretical* 52(19):193001. <https://doi.org/10.1088/1751-8121/ab0850>
15. Still S, Sivak DA, Bell AJ, Crooks GE (2012) Thermodynamics of prediction. *Physical Review Letters* 109(12):120604. <https://doi.org/10.1103/PhysRevLett.109.120604>
16. Shannon CE (1948) A mathematical theory of communication. *Bell System Technical Journal* 27(3):379-423. <https://doi.org/10.1002/j.1538-7305.1948.tb01338.x>
17. Jaynes ET (1957) Information theory and statistical mechanics. *Physical Review* 106(4):620-630. <https://doi.org/10.1103/PhysRev.106.620>
18. Landauer R (1961) Irreversibility and heat generation in the computing process. *IBM Journal of Research and Development* 5(3):183-191. <https://doi.org/10.1147/rd.53.0183>
19. Bennett CH (1982) The thermodynamics of computation-a review. *International Journal of Theoretical Physics* 21(12):905-940. <https://doi.org/10.1007/BF02084158>
20. Jarzynski C (1997) Nonequilibrium equality for free energy differences. *Physical Review Letters* 78(14):2690-2693. <https://doi.org/10.1103/PhysRevLett.78.2690>
21. Crooks GE (1999) Entropy production fluctuation theorem and the nonequilibrium work relation for free energy differences. *Physical Review E* 60(3):2721-2726. <https://doi.org/10.1103/PhysRevE.60.2721>
22. Seifert U (2012) Stochastic thermodynamics, fluctuation theorems and molecular machines. *Reports on Progress in Physics* 75(12):126001. <https://doi.org/10.1088/0034-4885/75/12/126001>
23. Parrondo JMR, Horowitz JM, Sagawa T (2015) Thermodynamics of information. *Nature Physics* 11(2):131-139. <https://doi.org/10.1038/nphys3230>
24. Bekenstein JD (1973) Black holes and entropy. *Physical Review D* 7(8):2333-2346. <https://doi.org/10.1103/PhysRevD.7.2333>
25. Onsager L (1931) Reciprocal relations in irreversible processes. I. *Physical Review* 37(4):405-426. <https://doi.org/10.1103/PhysRev.37.405>
26. Tsallis C (1988) Possible generalization of Boltzmann-Gibbs statistics. *Journal of Statistical Physics* 52(1-2):479-487. <https://doi.org/10.1007/BF01016429>