FinEntropyBench: What Does Entropy Tell Us About Market Efficiency?

Idrees Aziz
Middle East Technical University
Ankara, Türkiye
idrees.mustafa.aziz@gmail.com

Abstract

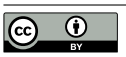
Entropy has become a standard instrument for comparing market efficiency, yet the statistics used under that name target different forms of uncertainty. We introduce FinEntropyBench, a finance-specific benchmark that validates an entropy measure before it enters an efficiency ranking. The benchmark combines analytic and oracle ground truth, finance-relevant stochastic processes, ordinal, marginal, conditional, and neural estimators, three return representations, established econometric tests, and independent walk-forward strategies. The controlled results distinguish estimand error from estimator error. For an AR(1) process with $\phi = 0.9$, a k-nearest-neighbor estimator accurately recovers marginal entropy (1.409 versus 1.419 nats) while missing the entropy rate (0.589 nats); a conditional Gaussian estimator returns 0.591. In matched simulations, sign dependence produces 0.0291 bits of out-of-sample code gain, volatility dependence produces 0.0353 magnitude bits, and oracle standardization removes the latter while preserving the former. Across 16 public daily market series, raw magnitude code gain averages 0.0150 bits per day, sign code gain averages -0.0009, and no sign result survives false-discovery control. BDS rejects independence for all raw series and one GARCH-filtered series; the heteroskedasticity-robust variance-ratio test rejects none. In 113 prospective asset-year tests, prior sign gain has Spearman correlation -0.039 with next-year net Sharpe (95% cluster-bootstrap interval $[-0.236, 0.092]$). The benchmark reports estimator validity, distributional regularity, directional efficiency, and economic efficiency as separate quantities.

CCS Concepts

• **Computing methodologies** → **Machine learning**; • **Applied computing** → **Economics**.

Keywords

entropy rate, market efficiency, financial time series, benchmark, neural cross-entropy, volatility



This work is licensed under a Creative Commons Attribution 4.0 International License.
ICAIF '26, Milan, Italy
© 2026 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-3087-0/2026/11
https://doi.org/10.1145/3858063.3858730

ACM Reference Format:
Idrees Aziz. 2026. FinEntropyBench: What Does Entropy Tell Us About Market Efficiency?. In *7th ACM International Conference on AI in Finance (ICAIF '26), November 14–17, 2026, Milan, Italy*. ACM, New York, NY, USA, 7 pages. https://doi.org/10.1145/3858063.3858730

1 Introduction

Entropy is widely used to compare financial market efficiency across countries, assets, sectors, frequencies, and crisis periods. The literature spans Shannon entropy, approximate and sample entropy, permutation entropy, Lempel–Ziv complexity, and entropy-derived predictability bounds [1, 5, 18, 21, 26, 27]. These methods share an intuitive premise: a more efficient return process leaves less structure for the past to explain.

The premise becomes operational only after the measured object is fixed. Marginal differential entropy describes one return and changes with scale. Permutation entropy describes ordinal motifs and omits amplitude. Sample entropy describes recurrence at a chosen tolerance. A symbolic block estimator targets a finite-alphabet entropy rate. Neural negative log-likelihood is a predictive code length whose gap above the entropy rate depends on model class, optimization, and data. A single label—entropy—therefore covers quantities with different units, invariances, and economic meanings.

Finance makes the distinction consequential. A GARCH process can have zero conditional mean and strongly predictable conditional variance [2, 7]. An amplitude-aware method will detect that structure even when return direction remains unpredictable. Weak-form efficiency, however, concerns information available from past prices and the performance of feasible decision rules [9, 25]. Distributional dependence, directional predictability, and profitable predictability form related but separate empirical questions.

The component limitations are established: finite-sample entropy bias [11, 17], ordinal amplitude blindness [8], embedding sensitivity [19], the difficulty of continuous entropy-rate estimation [6, 10], symbolization dependence [16], and market-microstructure and volatility effects [20]. We therefore evaluate these estimators in one protocol that also tests whether their outputs support the claimed economic interpretation.

FinEntropyBench evaluates an entropy measure before it is used to rank efficiency. It makes three contributions:

(1) **Estimand-aware measurement.** Every output records its target, representation, unit, uncertainty, and bound direction. Marginal entropy, conditional entropy rate, ordinal complexity, captured information, and predictive cross-entropy remain separate objects.
(2) **Finance-specific calibration.** Analytic and oracle processes test estimator behavior at sample sizes observed in financial data, with conditional heteroskedasticity as a central case.
(3) **Efficiency triangulation.** Sign and magnitude channels, volatility normalization, econometric diagnostics, a five-seed neural code, and independent walk-forward strategies connect distributional, directional, and economic efficiency.

2 Literature and Positioning

Entropy-based market studies established a flexible model-free language for regularity. Permutation-entropy work ranks national equity markets and specific U.S. money-, bond-, and stock-market instruments [21, 27]; multiscale and rolling measures examine horizons and crises [1, 18]; reviews document applications across foreign exchange, commodities, cryptocurrencies, and portfolio analysis [15, 26]. Entropy-rate work estimates serial dependence directly [11, 14], while Fano-based approaches convert symbolic rates into prediction-accuracy limits [5, 22]. Neural predictive loss supplies another route: held-out cross-entropy is an upper bound on the process entropy rate [23].

The closest corrections address individual stages. Weighted permutation entropy restores amplitude information omitted by ordinal patterns [8]. Olbryś and Komar compare financial symbolization schemes [16]. Shternshis et al. filter volatility, staleness, seasonality, and microstructure before interpreting high-frequency Shannon entropy [20]. Calcagnile et al. benchmark several estimators against certified dynamical entropy [4]. FinEntropyBench unifies ground-truth calibration, volatility decomposition, representation stability, neural coding, and independent economic validation in one finance-specific design.

3 FinEntropyBench

3.1 Targets and validity criteria

Let R_t be a continuous return and S_t a symbol derived from it. Marginal entropy measures unconditional uncertainty, while the entropy rate measures residual uncertainty after the past:

$$H(S_t) = -\sum_s p(s)\log p(s), \qquad h_\infty = \lim_{L\to\infty} H(S_t \mid S_{t-1:t-L}). \tag{1}$$

The continuous analogues use differential entropy. Ordinal and recurrence statistics are recorded as complexity measures in their native units.

For a predictive model q_θ, test cross-entropy obeys

$$\mathrm{CE}(q_\theta) = \mathbb{E}[-\log q_\theta(S_t \mid S_{<t})] = h_\infty + \mathbb{E}[D_{\mathrm{KL}}(p\|q_\theta)] \geq h_\infty. \tag{2}$$

We report the bound and the out-of-sample code gain

$$G = \mathrm{CE}_{\mathrm{marginal}} - \mathrm{CE}_{\mathrm{conditional}}, \tag{3}$$

measured in bits per observation. The marginal code is fit on training data and both codes are scored on the held-out block. With the true marginal and conditional distributions, G equals predictive mutual information. It also has an idealized Kelly log-growth interpretation when outcome probabilities are correct, admissible repeated bets are available, and odds satisfy the model's fair-odds conditions [12]. For fitted codes, positive G is captured predictive information; negative G means that estimation error makes the conditional code worse out of sample.

Four criteria govern an efficiency claim. **Conditionality** requires conditioning on the past. **Comparability** fixes scale, units, alphabet, and fitting window. **Calibration** tests the estimator at the observed sample length. **Task alignment** connects the measured channel to distributional, directional, or economic efficiency. These fields form an Entropy Estimate Card for every result.

3.2 Controlled processes and estimators

The calibration suite uses 25 independent replications at $N \in \{500, 1000, 2000, 5000\}$. It contains IID Gaussian and Student-t_3/t_5 returns; stationary Gaussian AR(1) processes with $\phi \in \{0.3, 0.6, 0.9\}$; GARCH(1,1) with $(\omega, \alpha, \beta) = (10^{-6}, 0.08, 0.90)$; and binary Markov chains with persistence $0.55, 0.65, 0.80$, and 0.95. Gaussian, Student-t, AR innovation, and Markov entropy rates are analytic. GARCH uses the oracle volatility path,

$$\bar{h} = \frac{1}{N}\sum_t \frac{1}{2}\log(2\pi e\sigma_t^2). \tag{4}$$

Bias and RMSE are computed only when estimator and reference share an estimand and units.

The suite evaluates k-nearest-neighbor marginal entropy, Gaussian autoregressive conditional coding, ordinary and weighted permutation entropy, and Markov plug-in rates. Real-series measurement adds sample entropy. Permutation statistics use $m = 5, \tau = 1$ in the main comparison; sensitivity grids are retained in the result schema.

A second controlled experiment isolates channels. Four processes combine predictable or unpredictable signs with clustered or IID magnitudes while preserving matched marginals. One hundred replications of length 5,000 are evaluated before and after oracle volatility standardization.

3.3 Market data and representations

The market layer contains daily OHLCV series from a public repository frozen at commit `5fa48d39fd20858d6591db25e6eeee5de45ce6fd` [24]: AAPL, NFLX, and TSLA; S&P 500, Dow 30, Nasdaq 100, DAX 30, and FTSE 100 index CFDs; EUR/USD, GBP/USD, and USD/JPY; gold, silver,

Table 1: Representative methodological coverage. Entries describe components explicitly implemented and evaluated in each study.

Work	Primary target	Ground-truth calibration	Volatility treatment	Representation protocol	Predictive bound	Independent downstream validation
Zunino et al. [27]	Ordinal complexity and market ranking	–	–	Fixed ordinal patterns	–	–
Gao et al. [11]	Binary entropy-rate estimation	Known simulations	–	Fixed binary alphabet	–	–
Takahashi and Tanaka-Ishii [23]	Neural entropy-rate bound	–	–	Multiple contexts	Held-out neural cross-entropy	–
Olbryś and Komar [16]	Symbolization sensitivity	–	–	Multiple threshold encodings	–	–
Chen et al. [5]	Lempel–Ziv rate and stock-price accuracy limit	–	–	Fixed tick states	Fano relation	Markov and kernel forecasts
Shternshis et al. [20]	High-frequency Shannon entropy	–	Empirical filtering	Staleness and microstructure controls	–	–
Calcagnile et al. [4]	Estimator accuracy for deterministic dynamics	Certified entropy	–	Numerical and symbolic trajectories	–	–
FinEntropyBench	Estimand-aware market-efficiency measurement	Analytic and oracle targets	Controlled GARCH and causal scaling	Walk-forward, train-fit symbols	Five-seed neural NLL	Prospective forecasts and net trading tests

and Brent; and BTC and ETH. Histories contain 1,657–5,171 returns and end in September 2023. A seventeenth crypto series is excluded by the prespecified $N \geq 1{,}500$ rule. Apple's 4:1 split and Tesla's 5:1 and 3:1 splits are adjusted before returns are formed; a post-adjustment audit finds no absolute log return above 0.59. The repository declares an MIT license for its contents, but it does not certify the upstream vendor/exchange provenance or redistribution rights of the underlying records; the empirical layer therefore tests measurement stability across public series rather than exchange-level price discovery.

Three representations test where information resides: raw returns, causal EWMA-standardized returns with $\lambda = 0.94$, and GARCH-standardized residuals whose parameters are estimated on the training window. Direction uses $S_t = \mathbb{1}(R_t \geq 0)$. Magnitude uses four train-fitted quantile bins of $|R_t|$. Markov orders 1–3 are selected on validation data separately for each channel. All thresholds and scale models are frozen before test evaluation.

3.4 Predictive and economic protocol

Full-series experiments use chronological 60/20/20 train, validation, and test blocks. A pooled LSTM predicts return signs across all 16 assets using a 20-day context, an 8-dimensional symbol embedding, a 4-dimensional asset embedding, one 24-unit recurrent layer, and five seeds. Validation NLL selects the epoch, with a maximum of eight.

Independent return forecasters use 20 lags plus causal rolling means, volatility, and absolute-return features at 5, 10, 20, and 60 days. In the full-series test, validation net Sharpe selects among prespecified ridge, random-forest, histogram-gradient-boosting, and momentum candidates. The annual panel uses ridge and momentum candidates to preserve estimation depth within each three-year training window. The selected specification is refit on train plus validation and evaluated once on test. Strategies take the sign of the forecast and pay 5 basis points per unit turnover.

External validity uses 113 prospective asset-year tests from 2011–2023. Each test year follows three training years and one validation year. Cluster bootstrap intervals resample assets; fixed-effects regressions include asset and year indicators. Benjamini–Hochberg correction controls each family of asset-level tests.

Computational assistance and data scope. The study uses public market records and simulated processes and contains no human-subject or personal data. ChatGPT drafted benchmark scripts and was used during manuscript refinement. The benchmark and experiments were designed by the authors, who debugged and ran the code, checked its outputs, and interpreted the results.

4 Results

4.1 Calibration separates estimands

Figure 1 shows three distinct outcomes. First, conditional Gaussian coding tracks Gaussian entropy rates. At $N = 5{,}000$, the AR(1) process with $\phi = 0.9$ has true rate 0.5886 nats and estimate 0.5908. The binary Markov plug-in estimate at persistence 0.8 is 0.7160 bits against 0.7219 truth.

Second, a marginal estimator can be accurate and still fail as an efficiency measure. The kNN estimate for AR(1) $\phi = 0.9$ is 1.4091 nats, close to the true marginal entropy 1.4189 and 0.8205 above the entropy rate. The error comes from substituting one estimand for another.

Third, ordinary ordinal patterns leave volatility structure nearly invisible. Normalized permutation entropy is 0.9976 for both IID Gaussian and GARCH returns; weighted permutation entropy is 0.9965 and 0.9948. The statistic remains useful for ordinal organization, while its near-maximal value does not establish conditional distributional efficiency.

The neural code passes the bound sanity check. On binary IID data, five-seed LSTM test NLL is 1.0008 ± 0.0009 bits against a one-bit rate. On a Markov process with persistence 0.8, it is 0.7938 ± 0.0226 against 0.7219 truth. The positive

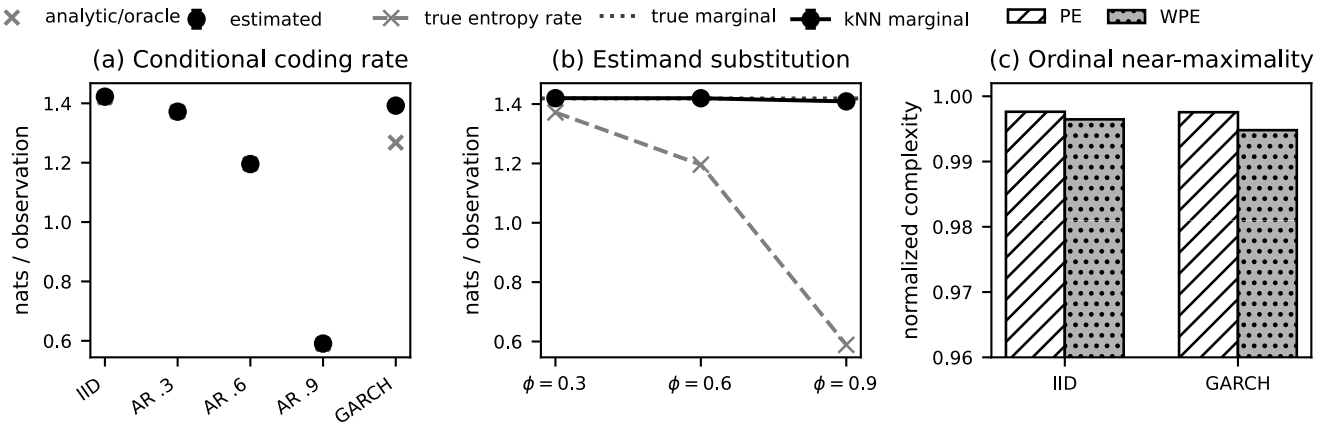


Figure 1: Ground-truth calibration at $N = 5{,}000$. Error bars are 95% Monte Carlo intervals. The panels keep conditional rates, marginal entropy, and normalized ordinal complexity on separate scales.

0.0718-bit gap quantifies model looseness. On the 16 real series, pooled test NLL is 1.0008 bits and sign accuracy is 51.28%.

4.2 Return regularity lies in magnitude

The controlled channel experiment recovers the source of information (Figure 2). Predictable signs yield 0.0291 bits of code gain whether volatility is raw or standardized. Volatility clustering yields 0.0353 magnitude bits in raw returns and no positive gain after oracle standardization. The combined process contains both channels. The construction therefore distinguishes conditional-mean information from conditional-scale information.

The market data follow the volatility case. Raw magnitude code gain averages 0.01498 bits per day (median 0.01569), while raw sign gain averages −0.00090 bits and no sign test survives false-discovery control. Twelve of 16 assets have positive raw magnitude gain; three survive correction. EWMA and GARCH standardization eliminate the positive cross-asset mean. The paired median raw-minus-EWMA difference is 0.01691 bits ($q = 0.029$); raw-minus-GARCH is 0.01537 ($q = 0.015$). Sign differences remain negligible.

4.3 Econometric tests locate the same structure

The heteroskedasticity-robust Lo–MacKinlay variance-ratio implementation is validated before application [13]. At $N = 1{,}000$, its 5% rejection rates are 6.0% for IID data and 4.5% for a GARCH martingale difference, while power is 95% for AR(1) $\phi = 0.15$ and 100% for $\phi = 0.30$. The BDS implementation rejects 10% of IID samples, 96% of GARCH samples, and 98% of AR(1) $\phi = 0.30$ samples at the same length [3].

Across the 16 market series, variance-ratio tests at horizons 2, 5, and 10 reject none after false-discovery control. BDS rejects all raw series and one GARCH-standardized series. Ljung–Box rejects zero autocorrelation in squared returns for 15 series and in returns for 10. These diagnostics target different restrictions: nonlinear and variance dependence is pervasive, while long-horizon variance-ratio evidence is absent. Their joint pattern matches the entropy decomposition.

Table 2: Efficiency triangulation after Benjamini–Hochberg correction.

Test or outcome	Result
VR rejections, $q = 2/5/10$	0/0/0 of 16
BDS rejections, raw / GARCH	16/1 of 16
Ljung–Box, returns / squared	10/15 of 16
Raw sign-gain rejections	0 of 16
Raw magnitude-gain rejections	3 of 16
Pooled LSTM test NLL	1.0008 bits
Independent test accuracy, mean	49.76%
Independent net Sharpe, median	−0.067

4.4 Entropy rankings do not determine economic efficiency

The 113 asset-year panel tests whether code gain measured before a year predicts performance during that year. Prior sign gain has Spearman correlation −0.039 with net Sharpe, −0.099 with directional accuracy, and −0.116 with information coefficient; every cluster-bootstrap interval crosses zero (Figure 4). Prior magnitude gain has correlation −0.146 with net Sharpe, also crossing zero. Its reliable relation is negative correlation with return-prediction R^2 (−0.218, interval $[-0.392, -0.062]$), consistent with volatility structure complicating conditional-mean prediction.

Asset and year fixed-effects regressions produce the same conclusion. Adding sign and magnitude information to variance ratio and Ljung–Box diagnostics changes net-Sharpe R^2

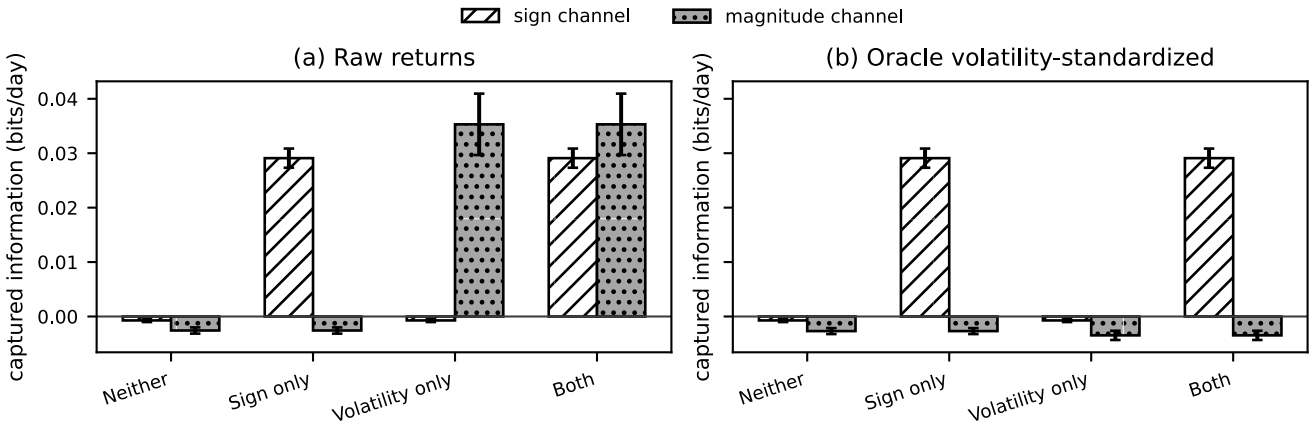


Figure 2: Matched synthetic channel isolation. Bars show mean out-of-sample code gain and 95% Monte Carlo intervals over 100 replications.

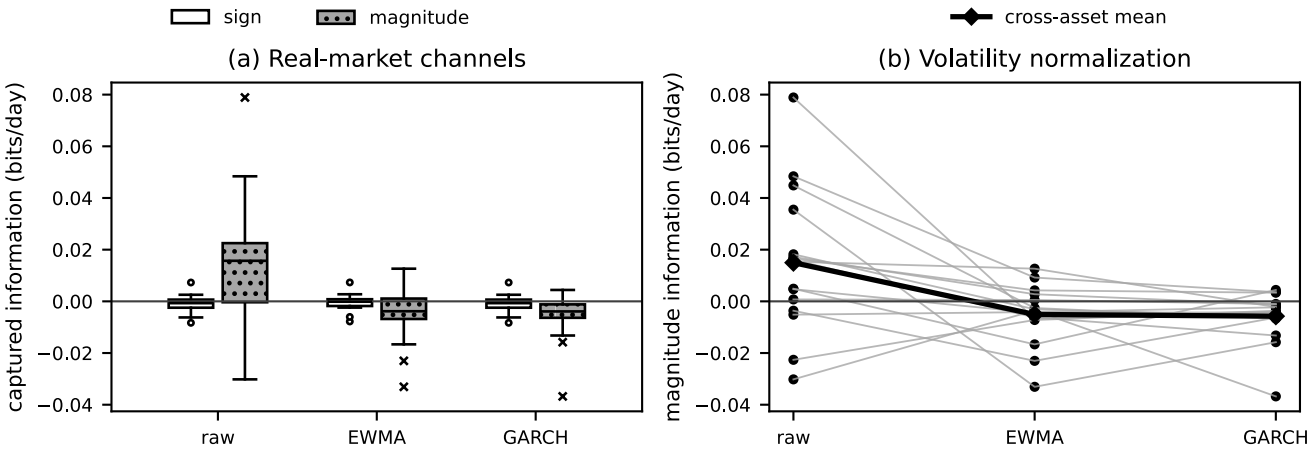


Figure 3: Real-market code-gain channels. Negative values mean the validation-selected conditional code underperforms the training-fit unconditional code on the test block.

from 0.535 to 0.540, raises AIC by 2.94, and yields entropy-channel p-values of 0.454 and 0.596. The independent full-sample strategies average 49.76% directional accuracy and median net Sharpe −0.067 after costs.

Representation also changes the ordering. Rank agreement is 0.821 between EWMA and GARCH channels, but 0.447 between EWMA and raw returns and 0.406 between GARCH and raw returns. A market-efficiency ranking therefore inherits the representation used to construct it.

5 Reporting Standard and Implications

FinEntropyBench yields five reporting rules.

Name the measured object. Distinguish marginal entropy, conditional entropy rate, ordinal or recurrence complexity, captured information, and predictive cross-entropy.

Preserve units and bound direction. Keep bits, nats, normalized scores, estimates, and bounds distinct. A neural NLL coding bound does not become a minimum error rate when inserted into a Fano relation.

Calibrate at the available sample length. Report process-specific bias, variance, and failures, including heteroskedastic cases that separate risk from direction.

Carry distinct representations in parallel. Raw, EWMA-standardized, and GARCH-standardized returns answer different questions; report their rank agreement.

Validate economic interpretations prospectively. Separate estimator and forecaster families, validate before test, report costs, and control multiplicity.

Interpretation has three levels: *distributional efficiency* concerns the conditional return distribution, *directional efficiency* its sign and mean, and *economic efficiency* estimation error, turnover, costs, and utility. The reliable daily structure is mainly volatility, not direction or net strategy

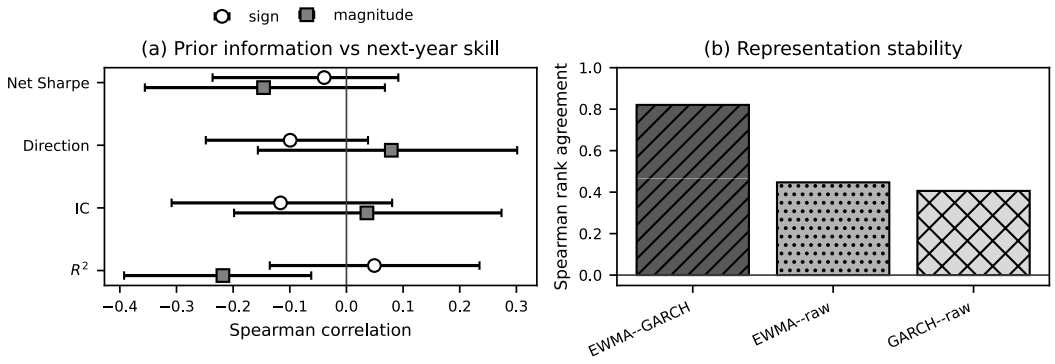


Figure 4: External validity and measurement stability. Horizontal bars are 95% asset-cluster bootstrap intervals.

performance. Without consolidated quotes and exchange-grade volume, liquidity remains out of scope. The pooled LSTM is an upper bound with a measured approximation gap, not an exact entropy-rate estimate.

6 Conclusion

Measured entropy does not provide a representation-free ordering of market efficiency. Ground-truth calibration exposes estimand substitution; matched simulations separate sign from volatility; and real-market tests place most positive code gain in magnitude without stable directional or economic advantage. Entropy rankings therefore require an explicit object, representation, calibrated estimator, and economic validation.

References

[1] José Alvarez-Ramirez, Eduardo Rodriguez, and Jesus Alvarez. 2012. A Multiscale Entropy Approach for Market Efficiency. *International Review of Financial Analysis* 21 (2012), 64–69. doi:10.1016/j.irfa.2011.12.001

[2] Tim Bollerslev. 1986. Generalized Autoregressive Conditional Heteroskedasticity. *Journal of Econometrics* 31, 3 (1986), 307–327. doi:10.1016/0304-4076(86)90063-1

[3] William A. Brock, W. Davis Dechert, Jose A. Scheinkman, and Blake LeBaron. 1996. A Test for Independence Based on the Correlation Dimension. *Econometric Reviews* 15, 3 (1996), 197–235. doi:10.1080/07474939608800353

[4] Lucio M. Calcagnile, Angelo Di Garbo, and Stefano Galatolo. 2026. A Benchmark for Entropy Estimators. *Entropy* 28, 3 (2026), 311. doi:10.3390/e28030311

[5] Yueshan Chen, Xingyu Xu, Tian Lan, and Sihai Zhang. 2024. The Predictability of Stock Price: Empirical Study on Tick Data in Chinese Stock Market. *Big Data Research* 35 (2024), 100414. doi:10.1016/j.bdr.2023.100414

[6] David Darmon. 2016. Specific Differential Entropy Rate Estimation for Continuous-Valued Time Series. *Entropy* 18, 5 (2016), 190. doi:10.3390/e18050190

[7] Robert F. Engle. 1982. Autoregressive Conditional Heteroscedasticity with Estimates of the Variance of United Kingdom Inflation. *Econometrica* 50, 4 (1982), 987–1007. doi:10.2307/1912773

[8] Bilal Fadlallah, Badong Chen, Andreas Keil, and José Príncipe. 2013. Weighted-Permutation Entropy: A Complexity Measure for Time Series Incorporating Amplitude Information. *Physical Review E* 87, 2 (2013), 022911. doi:10.1103/PhysRevE.87.022911

[9] Eugene F. Fama. 1970. Efficient Capital Markets: A Review of Theory and Empirical Work. *Journal of Finance* 25, 2 (1970), 383–417. doi:10.1111/j.1540-6261.1970.tb00518.x

[10] Andrew Feutrill and Matthew Roughan. 2021. A Review of Shannon and Differential Entropy Rate Estimation. *Entropy* 23, 8 (2021), 1046. doi:10.3390/e23081046

[11] Yun Gao, Ioannis Kontoyiannis, and Elie Bienenstock. 2008. Estimating the Entropy of Binary Time Series: Methodology, Some Theory and a Simulation Study. *Entropy* 10, 2 (2008), 71–99. doi:10.3390/entropy-e10020071

[12] John L. Kelly, Jr. 1956. A New Interpretation of Information Rate. *Bell System Technical Journal* 35, 4 (1956), 917–926. doi:10.1002/j.1538-7305.1956.tb03809.x

[13] Andrew W. Lo and A. Craig MacKinlay. 1988. Stock Market Prices Do Not Follow Random Walks: Evidence from a Simple Specification Test. *Review of Financial Studies* 1, 1 (1988), 41–66. doi:10.1093/rfs/1.1.41

[14] Esfandiar Maasoumi and Jeffrey Racine. 2002. Entropy and Predictability of Stock Market Returns. *Journal of Econometrics* 107, 1-2 (2002), 291–312. doi:10.1016/S0304-4076(01)00125-7

[15] Joanna Olbryś. 2026. Symbolic Time Series Analysis: A Systematic Review with Entropy-Based Applications in Finance. *Information* 17, 5 (2026), 423. doi:10.3390/info17050423

[16] Joanna Olbryś and Natalia Komar. 2023. Symbolic Encoding Methods with Entropy-Based Applications to Financial Time Series Analyses. *Entropy* 25, 7 (2023), 1009. doi:10.3390/e25071009

[17] Liam Paninski. 2003. Estimation of Entropy and Mutual Information. *Neural Computation* 15, 6 (2003), 1191–1253. doi:10.1162/089976603321780272

[18] Daniel Papla and Rafał Siedlecki. 2024. Entropy as a Tool for the Analysis of Stock Market Efficiency During Periods of Crisis. *Entropy* 26, 12 (2024), 1079. doi:10.3390/e26121079

[19] Maik Riedl, Andreas Müller, and Niels Wessel. 2013. Practical Considerations of Permutation Entropy. *European Physical Journal Special Topics* 222, 2 (2013), 249–262. doi:10.1140/epjst/e2013-01862-7

[20] Andrey Shternshis, Piero Mazzarisi, and Stefano Marmi. 2022. Measuring Market Efficiency: The Shannon Entropy of High-Frequency Financial Time Series. *Chaos, Solitons & Fractals* 162 (2022), 112403. doi:10.1016/j.chaos.2022.112403

[21] Fotios M. Siokis. 2018. Credit Market Jitters in the Course of the Financial Crisis: A Permutation Entropy Approach in Measuring Informational Efficiency in Financial Assets. *Physica A* 499 (2018), 266–275. doi:10.1016/j.physa.2018.02.005

[22] Chaoming Song, Zehui Qu, Nicholas Blumm, and Albert-László Barabási. 2010. Limits of Predictability in Human Mobility. *Science* 327, 5968 (2010), 1018–1021. doi:10.1126/science.1177170

[23] Shuntaro Takahashi and Kumiko Tanaka-Ishii. 2018. Cross Entropy of Neural Language Models at Infinity—A New Bound of the Entropy Rate. *Entropy* 20, 11 (2018), 839. doi:10.3390/e20110839

[24] TheSnowGuru. 2024. Stocks, Futures and Financial Time-Series Tick/Bar Data. GitHub repository, https://github.com/TheSnowGuru/Stocks-Futures-Financial-Time-series-Tick-Bar-Data/tree/5fa48d39fd20858d6591db25e6eeee5de45ce6fd. Repository snapshot at commit 5fa48d39fd20858d6591db25e6eeee5de45ce6fd, accessed 29 July 2026; the repository declares an MIT license

but does not certify the exchange provenance or redistribution rights of the underlying market records.

[25] Allan Timmermann and Clive W. J. Granger. 2004. Efficient Market Hypothesis and Forecasting. *International Journal of Forecasting* 20, 1 (2004), 15–27. doi:10.1016/S0169-2070(03)00012-8

[26] Rongxi Zhou, Ru Cai, and Guanqun Tong. 2013. Applications of Entropy in Finance: A Review. *Entropy* 15, 11 (2013), 4909–4931. doi:10.3390/e15114909

[27] Luciano Zunino, Massimiliano Zanin, Benjamin M. Tabak, Darío G. Pérez, and Osvaldo A. Rosso. 2009. Forbidden Patterns, Permutation Entropy and Stock Market Inefficiency. *Physica A* 388, 14 (2009), 2854–2864. doi:10.1016/j.physa.2009.03.042