

Entropy-Driven Prediction in Market-Neutral Trading: A New Statistical Signal

Bin Li
Jonathan Spring

May 3, 2025

Abstract

We propose a novel trading framework based on entropy-gradient prediction applied to market-neutral financial time series. Grounded in principles from statistical physics and information theory, our hypothesis asserts that in approximately closed financial systems, entropy tends to increase over time—implying a directional bias in asset price dynamics. We formalize this idea by modeling small perturbations in return distributions and comparing their impact on Shannon entropy. A large-scale empirical evaluation across 29 equity, sectoral, international, and fixed income ETF spreads reveals that the entropy-gradient signal performs consistently well in macro-hedged or structurally distinct pairings. Results show Sharpe ratios up to 0.86 and meaningful returns in semi-stable systems, even in the absence of transaction cost modeling or risk overlays. These findings highlight entropy as a second-order probabilistic signal and open avenues for entropy-aware modeling across finance, particularly in systematic strategies where conventional price-based signals fail.

1 Introduction

Market-neutral strategies aim to eliminate systematic market risk by taking offsetting long and short positions. Traditional approaches often rely on mean-reversion, cointegration, or machine learning. In this work, we introduce an orthogonal framework grounded in **entropy-based prediction**—drawing inspiration from statistical mechanics and information theory.

The conceptual foundation is the entropy principle: in approximately closed systems, configurations evolve toward maximum entropy [2]. Translating this to financial time series, we hypothesize that the spread between hedged asset pairs behaves analogously to a near-closed system. Therefore, the direction of maximal entropy increase in return distributions can serve as a probabilistic predictor of future movement. This yields a novel, information-theoretic directional signal distinct from traditional trend or mean-reversion mechanisms.

2 Theoretical Framework

Let X_t be a market-neutral time series, such as a normalized spread between two asset prices. Over a moving window $[t - w, t]$, define the return series R_t as:

$$R_t = \left\{ \frac{X_i - X_{i-1}}{X_{i-1}} \right\}_{i=t-w+1}^t \quad (1)$$

We define the Shannon entropy of the empirical distribution of R_t as [1]:

$$H(R_t) = - \sum_{i=1}^n p_i \log p_i \quad (2)$$

where p_i is the probability mass in histogram bin i .

Now suppose X_{t+1} undergoes a small hypothetical upward or downward perturbation, i.e.,

$$R_{\text{up}} = R_t \cup \left\{ \frac{X_t + \delta - X_t}{X_t} \right\}, \quad R_{\text{down}} = R_t \cup \left\{ \frac{X_t - \delta - X_t}{X_t} \right\} \quad (3)$$

The entropy differentials are:

$$\Delta H_{\text{up}} = H(R_{\text{up}}) - H(R_t) \quad (4)$$

$$\Delta H_{\text{down}} = H(R_{\text{down}}) - H(R_t) \quad (5)$$

We hypothesize that:

$$\text{If } \Delta H_{\text{up}} > \Delta H_{\text{down}}, \text{ then } X_{t+1} > X_t \text{ is more likely} \quad (6)$$

This can be interpreted as an entropy-gradient predictor analogous to a physical system's tendency toward thermodynamic equilibrium [3]. Importantly, this entropy signal operates as a *second-order effect*—its predictive power does not rely on direct price-level features like trend or reversion but instead emerges from the structural configuration of return distributions. In statistical physics, such entropy flows dominate once local forces equilibrate; likewise, in financial systems, we expect entropy-gradient directionality to become relevant when more dominant first-order signals (e.g., arbitrage, mean-reversion) are weak or neutralized. This makes the entropy signal both probabilistic and latent, functioning best in market-neutral spreads where directional drift is minimal and structure governs evolution.

3 Model Implementation

We constructed a normalized spread $S_t = P_t^{\text{DOW}}/P_0^{\text{DOW}} - P_t^{\text{SP500}}/P_0^{\text{SP500}}$, and applied the entropy-gradient method on its moving window of returns.

Key model components:

- **Entropy window size:** $w = 29$ days
- **Delta perturbation:** $\delta = 0.01$
- **Signal debounce:** trades must persist for ≥ 3 days
- **Volatility scaling:** position size scaled by rolling 20-day volatility
- **Return clipping:** max return per trade capped at 5%

A trade signal is generated if $|\Delta H| > 0.002$, and its direction is given by $\text{sign}(\Delta H_{\text{up}} - \Delta H_{\text{down}})$.

Metric	Value
Final Return	74.8%
Sharpe Ratio	0.45
Max Drawdown	-42.1%
Win Rate	50.9%
Trades Executed	119

Table 1: Performance metrics for robust entropy strategy, 2019–2025.




Figure 1: Cumulative return of the entropy strategy over time using the S&P 500 and Dow Jones Industrial Average (Dow 30) spread.

4 Results

4.1 Quantitative Performance

4.2 Discussion

Despite a near-random win rate, the strategy achieves meaningful returns due to volatility-adjusted position sizing and the probabilistic nature of the entropy gradient signal. Importantly, trades are constructed to be net-neutral across the pair, avoiding persistent directional bias and mitigating long-term divergence effects. This design ensures that observed returns emerge from distributional asymmetries rather than directional drift, trend-following, or persistent exposure to a dominant asset.

The entropy-based signal does not rely on explicit mean-reversion or momentum assumptions. Instead, it captures weak directional pressure derived from differences in expected entropy under small perturbations—a second-order effect that becomes salient in structured, market-neutral environments. While raw drawdowns remain significant, they reflect the unfiltered nature of the entropy signal and can be managed with capital overlays, stop-loss mechanisms, or entropy-aware drawdown constraints.

Overall, the empirical asymmetries in ΔH distributions lend support to the theoretical claim that financial systems exhibit local entropy gradients, which, when carefully extracted, can inform position-taking even in the absence of traditional price signals.

5 Theoretical Significance and Generalization

This model suggests a new class of predictors based on entropy gradients in financial signals. Its potential generalizations include:

- Currency spreads (e.g., USD/JPY vs. EUR/USD)
- Volatility indices (e.g., VIX vs. MOVE)
- Sector pairs (e.g., tech vs. utilities)
- Implied vs. realized vol spread dynamics

From a theoretical lens, it hints at deeper links between information theory and market microstructure, where entropy evolution governs expectation formation among agents.

6 Extended Results Across 29 ETF Pairs

To examine the generalizability of the entropy-gradient trading strategy, we extended our empirical study to a comprehensive set of 29 ETF pairs covering equity indices, sector rotations, global markets, style factors, credit spreads, and real assets. These pairs were carefully chosen to span distinct macroeconomic dimensions and asset class behaviors.

6.1 Performance Summary


Table 2 and Table 3 present summary metrics across all 29 ETF pairs. The strategy generally maintains win rates close to 50–52%, with Sharpe ratios ranging from approximately -1.0 to 0.86. Pairs with strong macro-structural distinctions yielded notably higher Sharpe ratios and smoother drawdowns.

Pair	ETFs	Final Return	Sharpe	Max DD	Win Rate	Trades
DJI–GSPC	Dow 30 vs. S&P 500	1.75	0.45	-0.42	0.509	119
SPY–QQQ	S&P 500 vs. Nasdaq 100	0.47	-0.26	-0.72	0.505	110
IWM–SPY	Russell 2000 vs. S&P 500	2.18	0.56	-0.54	0.525	111
VOO–SPY	Vanguard vs. SPDR S&P 500	0.82	0.03	-0.49	0.501	128
DIA–SPY	Dow ETF vs. S&P 500	1.30	0.29	-0.51	0.508	118
XLFX–XLK	Financials vs. Technology	3.54	0.83	-0.26	0.522	106
XLY–XLP	Discretionary vs. Staples	3.72	0.86	-0.34	0.522	118
XLI–XLU	Industrials vs. Utilities	1.53	0.38	-0.55	0.493	136
XLE–XLV	Energy vs. Health Care	0.24	-0.63	-0.85	0.497	109
XLB–XLRE	Materials vs. Real Estate	2.57	0.66	-0.45	0.516	111
VUG–VTV	Growth vs. Value (Vanguard)	0.21	-0.66	-0.80	0.473	115
IWF–IWD	Growth vs. Value (Russell 1000)	0.33	-0.44	-0.83	0.485	97
MGK–MGV	Mega Cap Growth vs. Value	0.68	-0.04	-0.58	0.491	101
SPYG–SPYV	S&P 500 Growth vs. Value	0.32	-0.47	-0.73	0.502	120
RPG–RPV	S&P Pure Growth vs. Value	1.55	0.39	-0.42	0.501	128

Table 2: Entropy strategy results: Index, sector, and style pairs (1–15 of 29).

Pair	ETFs	Final Return	Sharpe	Max DD	Win Rate	Trades
EEM-EFA	Emerging vs. Developed Mkts	1.92	0.51	-0.38	0.501	141
EWJ-EWC	Japan vs. Canada	0.12	-1.02	-0.90	0.477	133
EWZ-EWW	Brazil vs. Mexico	0.59	-0.13	-0.66	0.479	117
FXI-EWY	China vs. South Korea	1.20	0.25	-0.75	0.499	133
VGK-VPL	Europe vs. Pacific	0.80	0.04	-0.63	0.511	149
HYG-LQD	High Yield vs. IG Bonds	0.51	-0.20	-0.72	0.510	145
TLT-IEF	Long vs. Intermediate Treasuries	2.10	0.55	-0.60	0.507	147
SHY-TLT	Short vs. Long Treasuries	2.28	0.60	-0.55	0.504	134
TIP-LQD	TIPS vs. Investment Grade	2.64	0.66	-0.42	0.522	129
GLD-SLV	Gold vs. Silver	0.68	-0.07	-0.62	0.493	133
DBC-USO	Commodity Index vs. Crude Oil	0.57	-0.15	-0.70	0.500	97
GDX-GDXJ	Gold Miners vs. Junior Miners	1.22	0.26	-0.57	0.497	132
VNQ-IYR	REITs vs. Real Estate Sector	2.29	0.59	-0.30	0.507	150
XOP-OIH	Exploration vs. Oil Services	1.20	0.25	-0.37	0.508	126

Table 3: Entropy strategy results: Global, fixed income, and commodity pairs (16–29 of 29).



(a) XLY-XLP: Strong, stable upward trend; the entropy signal aligns well with macro-cyclical regime shifts (cyclical vs. defensive sectors).

(b) EWJ-EWC: Poor and unstable performance; regional pair with no persistent structure and high noise undermines entropy signal efficacy.

Figure 2: Comparison of best (XLY-XLP) and worst (EWJ-EWC) performing entropy strategy pairs. Performance strongly depends on structural hedging and macro asymmetry between assets.

6.2 Insights by Category

Sector Rotations: Pairs such as XLY--XLP and XLF--XLK achieved the highest Sharpe ratios (0.86 and 0.83 respectively), showing that entropy gradients are especially effective when applied to spreads reflecting cyclical versus defensive exposures.

Growth vs. Value: Style factor pairs including VUG--VTI, IWF--IWD, and SPYG--SPYV underperformed, with negative Sharpe ratios and large drawdowns. These results suggest entropy signals struggle in environments with episodic regime shifts or overlapping exposures.

International Equity Spreads: Mixed results were observed. EEM--EFA produced a strong return and Sharpe (1.92, 0.51), whereas others like EWJ--EWC suffered due to regional noise and structural inefficiencies.

Bond and Credit Spreads: Treasury curve trades (TLT--IEF, SHY--TLT) and inflation vs. credit (TIP--LQD) performed well, validating entropy's ability to detect directional pressure in fixed income regimes. Conversely, credit-risk spreads like HYG--LQD proved noisy and unstable.

Real Assets and Commodities: The entropy model performed moderately well for VNQ--IYR and GDX--GDXJ, but failed for highly correlated metal pairs (GLD--SLV) or reactive energy spreads

(DBC--US0).

7 Discussion and Summary

7.1 Empirical Insights

Entropy-gradient strategies are most effective when applied to semi-closed systems with persistent macro asymmetries. Sector rotation pairs, bond curve spreads, and regional macro divergences offer clearer entropy profiles, while highly correlated or volatile asset pairs undermine predictability. These results support the hypothesis that entropy gradients capture weak but consistent directional bias in structured financial subsystems.

7.2 Purity of Signal and Experimental Controls

It is important to note that our empirical tests did not include transaction costs, which for institutional investors typically amount to approximately 0.2% round-trip per trade, including bid/ask spreads and fees. Additionally, no form of risk management, capital allocation constraints, or stop-loss overlays were applied—by design—to isolate and evaluate the pure predictive effect of the entropy gradient signal. Despite this, several pairs exhibited strong net performance, suggesting intrinsic signal efficacy.

Our position sizing methodology ensured that trades were on average net-neutral, i.e., not persistently long one asset and short the other. This was critical to avoid skewing results due to long-term divergence in the relative value of instruments within a pair. The observed returns, therefore, reflect entropy-driven directional asymmetries rather than mean reversion, drift, or persistent beta effects.

7.3 Theoretical Interpretation of Entropy as a Second-Order Effect

Entropy-based signals can be viewed as second-order effects: weak but persistent directional pressures that emerge when dominant first-order market dynamics—such as trend, arbitrage, or mean-reversion—cancel or reach equilibrium. In this regime, entropy gradients provide a statistical preference for future configurations, much like in thermodynamic systems where entropy governs macroscopic flows once local constraints are averaged out.

This second-order character explains the modest but consistent predictive performance of the strategy. Entropy gradients are not sharp arbitrage opportunities, but probabilistic indicators of directional imbalance embedded in the distributional structure of returns.

7.4 Interaction with Market Efficiency

If markets are highly efficient and traders actively exploit entropy signals, then the strength of these signals may diminish over time as their informational value is arbitrated away. However, because entropy gradients are not directly observable from price or volume alone, and because they operate on distributional rather than transactional structures, complete arbitrage is unlikely.

Conversely, if market participants remain unaware of entropy-based mechanisms, such signals may persist as residual inefficiencies, especially in complex spreads where informational saturation is lower. This places entropy-gradient strategies in a unique position: weak enough to coexist with efficient markets, yet distinct enough to offer sustainable edge in certain structured subsystems.

7.5 Conclusion of Findings

The cumulative evidence from our empirical tests and theoretical framing suggests that entropy-aware trading strategies represent a novel and potentially generalizable class of predictive models. Their low correlation with traditional signal types makes them strong candidates for inclusion in diversified systematic frameworks, particularly in market-neutral or volatility-hedged contexts. Future research should investigate their robustness under slippage, regime shifts, and their potential synergy with learning-based signal enhancers.

8 Conclusion

Entropy-based prediction offers a novel lens on market-neutral dynamics by leveraging directional pressure inferred from information-theoretic principles. Our extended empirical study across 30 diverse ETF pairs demonstrates that entropy-gradient signals are not only viable but particularly effective in semi-structured financial subsystems—especially macro-sectoral spreads, fixed income curve trades, and volatility-hedged pairings.

While raw trading performance varies, the method consistently identifies weak directional biases in systems characterized by persistent regime asymmetries. In contrast, entropy signals tend to underperform in highly correlated or episodically volatile pairs where information entropy offers limited predictive separation.

Future work should formalize the entropy gradient as a probabilistic indicator, extend the framework to multivariate and non-linear entropy dynamics, and explore integration with reinforcement learning and Bayesian signal updating. Additionally, cross-validation under varying volatility and liquidity conditions will be essential to assessing robustness and scalability. Importantly, extending the entropy-based framework to options markets and derivative-hedged portfolios represents a promising research direction. Such instruments introduce nonlinear payoffs and state-contingent asymmetries, offering a richer landscape in which entropy gradients may reveal directional or volatility-based pressure more effectively.

These results contribute to a growing intersection of statistical physics, information theory, and financial market modeling, suggesting entropy-based approaches as a promising paradigm for data-driven trading strategies and risk diagnostics.

References

- [1] C. E. Shannon. A mathematical theory of communication. *Bell System Technical Journal*, 27(3):379–423, 1948.
- [2] E. T. Jaynes. Information theory and statistical mechanics. *Physical Review*, 106(4):620–630, 1957.
- [3] L. E. Reichl. *A Modern Course in Statistical Physics*. Wiley-VCH, 2016.
- [4] E. Gatev, W. N. Goetzmann, and K. G. Rouwenhorst. Pairs trading: Performance of a relative-value arbitrage rule. *The Review of Financial Studies*, 19(3):797–827, 2006.
- [5] M. Avellaneda and J. Lee. Statistical arbitrage in the US equities market. *Quantitative Finance*, 10(7):761–782, 2010.