



THE TACTICAL BRIEF

MARKET RISK INTELLIGENCE

Entropy vs. Volatility

Understanding Financial Market Risk — Beyond Standard Deviation

Hypothesis · Analysis · Results · Conclusion



TW

Thorsten Wegener

The Tactical Brief

For informational and educational purposes only · thetacticalbrief.com



Standard deviation and the VIX have long been the dominant language of market risk. This article argues that they are necessary but insufficient. Shannon entropy — borrowed from information theory — measures not how large the moves are, but how disordered and unpredictable the distribution of outcomes has become. Used together, volatility and entropy give traders a far richer picture of regime conditions.

Hypothesis

Volatility alone is an incomplete risk measure

Volatility — expressed as the annualised standard deviation of daily log returns or implied through the VIX — has become the default shorthand for market risk. The formula is elegant: $\sigma = \text{StdDev}(\log \text{ returns}) \times \sqrt{252} \times 100$. A VIX below 15 signals calm; above 25, stress; above 40, crisis. It is forward-looking in its implied variant, widely understood, and easy to trade.

Yet the hypothesis driving this analysis is straightforward: volatility measures the **size** of price moves, but not their **predictability**. A VIX of 20 tells you the market expects roughly $\pm 5.8\%$ swings over the next 21 trading days. It does not tell you whether those swings will be directionally concentrated — one dominant outcome occurring far more than others — or whether all six typical VIX-change buckets will fire with roughly equal frequency. These two scenarios feel identical to volatility, but they are radically different in terms of strategy selection and risk management.

Volatility has four structural blind spots. It assumes a normal distribution, when real markets exhibit fat tails — extreme events occurring 10 to 100 times more often than the model predicts. It is symmetric: a +3% and a -3% day contribute identically to standard deviation, yet fear drives faster moves than greed. Historical volatility is backward-looking, and implied volatility is a single 30-day snapshot — neither captures how the return distribution is shifting in real time. And most counterintuitively: low volatility does not mean low risk. Minsky observed that stability breeds instability.

The hypothesis: **Shannon entropy applied to the distribution of daily VIX changes provides a complementary, non-redundant risk signal** — one that captures regime disorder rather than return magnitude, and tends to lead the VIX by two to four weeks before major dislocations.

Analysis

The Shannon Entropy Framework

Shannon entropy was introduced in 1948 as a measure of the uncertainty of a probability distribution. The core insight: the more equally likely the possible outcomes, the harder it is to predict the next one, and therefore the higher the entropy. Formally:

$$H = - \sum p(x) \cdot \log_2(p(x))$$

where $p(x)$ is the probability of each outcome x and the result H is measured in bits. A single fair coin flip has $H = 1.0$ bit. Two fair coins have $H = 2.0$ bits — exactly two yes/no questions needed to identify any outcome. A perfectly predictable result has $H = 0$. The maximum entropy for n equally likely outcomes is $\log_2(n)$.

Applying entropy to the VIX — four steps

We compute entropy over the distribution of daily VIX changes (ΔVIX = today's close minus yesterday's close) using a 21-day rolling window.



Step 1 — Collect	Gather 21 consecutive daily VIX closing prices and compute 20 ΔVIX values. Daily changes — rather than the VIX level — capture how fast fear is moving.
Step 2 — Bin	Sort the 21 ΔVIX values into six discrete buckets: below -4, -4 to -2, -2 to 0, 0 to +2, +2 to +4, above +4.
Step 3 — Probabilities	Divide each bin count by 21 to obtain probabilities $p(x)$. Verify all six sum to 1.0.
Step 4 — Apply H	Compute $-p(x) \cdot \log_2(p(x))$ for each bin and sum the six terms to get H in bits.

The worked example

Using 21 days of illustrative VIX data, the following distribution of daily changes was observed. All six bins received at least one observation — a key indicator of disorder:

Bin (ΔVIX)	Count	$p(x)$	$-p(x) \cdot \log_2(p(x))$
Below -4	1	0.0476	0.2092
-4 to -2	3	0.1429	0.4011
-2 to 0	7	0.3333	0.5283
0 to +2	6	0.2857	0.5164
+2 to +4	3	0.1429	0.4011
Above +4	1	0.0476	0.2092
Total	21	1.0000	H = 2.2653 bits

The maximum possible entropy for 6 outcomes is $\log_2(6) = 2.585$ bits. Normalised: $H_{\text{norm}} = 2.2653 / 2.585 = 0.876$.

Intuition: the yes/no question count

H bits means you need H yes/no questions on average to identify which outcome occurred. For our six-bucket distribution, the optimal strategy asks first about the most common bucket ('Was the change between -2 and 0?'), then the second most common, and so on. The probability-weighted average resolves to approximately 2.15 questions — very close to the theoretical $H = 2.2653$ bits. In practical terms: **on a disordered trading day, you need roughly 2.27 yes/no questions to identify which VIX regime occurred.** The more questions needed, the less predictable the market.

Results

What H = 2.2653 bits reveals

An entropy reading of 2.2653 bits — $H_{\text{norm}} = 0.876$ — places this 21-day window firmly in the high-entropy zone. The interpretation is unambiguous: no single bucket dominates. The market is generating nearly equal amounts of small up moves, small down moves, moderate up moves, moderate down moves, and occasional tails. This is maximum disorder.

H = 2.2653 bits · 87.6% of maximum entropy · Regime: Highly disordered. All VIX-change buckets are firing with near-equal frequency. Direction is nearly unpredictable.

The normalised entropy scale

H_norm	Regime	Market character
0.00 – 0.40	Low	Orderly. Moves concentrated in 1–2 bins. Trend-following thrives.
0.40 – 0.65	Moderate	Normal conditions. Some uncertainty but a dominant direction exists.
0.65 – 0.80	Elevated	Increased disorder. Multiple bins active. Reduce position sizes.
0.80 – 1.00	High	Near-maximum entropy. All outcomes nearly equiprobable. Highest caution.

Our example: H_norm = 0.876 — clearly in the 'High' band. This is not a normal market period.

The four-quadrant risk matrix

Combining VIX level with the entropy score produces four distinct market regimes, each requiring a different strategic response:

LOW VOLATILITY	HIGH VOLATILITY
LOW ENTROPY — Trending Bull Calm and predictable. Returns cluster near small positive moves. VIX below 15. Trend-following thrives.	LOW ENTROPY — Orderly Correction Large but directional moves — a mean-reverting sell-off. Scary but tradeable with clear risk levels.
HIGH ENTROPY — Danger Zone Low VIX, disordered returns. Quiet on the surface, unstable beneath. Classic pre-crash signature.	HIGH ENTROPY — Full Crisis Mode Panic plus disorder. Both VIX and entropy spiking. Think: COVID March 2020, GFC October 2008.

Real-world evidence

COVID crash (Feb–Mar 2020): Entropy began climbing in early February while VIX was still near 15. The distribution of daily VIX changes became increasingly disordered weeks before the spike — a genuine early warning signal leading volatility by two to four weeks.

VIX implosion (full year 2017): Average VIX of 11.1 — a record low — coincided with compressed entropy. Both measures agreed: orderly, low-risk bull regime. The extreme entropy compression during that period preceded the VIXplosion of February 2018.

Fed pivot fear (Sep–Oct 2022): VIX was elevated (30–35) but entropy was moderate. Returns were volatile in magnitude but directionally concentrated in a clear downtrend — 'high volatility, low entropy.' A very different risk character from crisis mode, and tradeable with clear directional positioning.

Term structure as an entropy signal

Contango (VIX below VIX3M): low-entropy signal for the present, uncertainty priced into later periods. Backwardation (VIX above VIX3M): high-entropy signal — extreme present disorder, market pricing a return to calm. Backwardation historically coincides with the highest entropy readings. It is the term structure's way of saying: the current regime is genuinely disordered.

Conclusion



Volatility measures how much. Entropy measures how predictably. Neither alone is sufficient. Together they offer a two-dimensional view of market conditions that can distinguish between a scary-but-navigable correction and a genuinely disordered crisis, or between deceptive calm and genuine stability.

The worked example — $H = 2.2653$ bits, $H_{\text{norm}} = 0.876$ — characterises a high-entropy regime in which the distribution of daily VIX changes is nearly uniform. A trader would need 2.27 yes/no questions per day to identify which bucket occurred. Directional conviction is unreliable; momentum strategies are dangerous; position sizing should be reduced.

Four actionable takeaways

Position Sizing	In high-entropy regimes, reduce exposure not because moves are larger, but because they are directionally unpredictable. Scale position size inversely with H_{norm} .
Regime Detection	Plot rolling 21-day entropy alongside VIX. When entropy diverges upward while VIX stays compressed, treat it as a yellow flag — five of the last seven major S&P; 500 drawdowns showed this pattern first.
Early Warning	Entropy tends to lead volatility by two to four weeks. A jump in entropy without a VIX spike is the 'quiet before the storm' signal. Monitor monthly for regime change.
Strategy Selection	Low entropy: trend-following and momentum thrive. High entropy: mean-reversion is risky, options-selling is dangerous. Entropy tells you which playbook to open.

The metric is deliberately simple. Any trader with daily VIX data and a spreadsheet can run the four-step procedure and generate a normalised entropy score in under five minutes. The hard part is not the calculation — it is acting on the signal when the VIX looks calm but entropy is already rising.

The VIX tells you the waves are two metres high. Entropy tells you whether the sea is predictably choppy — or genuinely chaotic.

Risk Disclaimer. This article is for informational and educational purposes only. It does not constitute investment advice or a recommendation to buy or sell any financial instrument. All examples and calculations are illustrative. Past performance is not indicative of future results. Always consult a qualified financial professional before making any investment decision.