RESEARCH LIBRARY · ON-CHAIN ANALYSIS

What Is the MVRV Z-Score?

The valuation gauge that standardises Bitcoin's market price against what holders actually paid — how it is built, how to read it, and where it quietly misleads.

Alain AI Lab · Research Desk

Published July 3, 2026

The MVRV Z-Score is a single number that tries to answer a deceptively hard question: is Bitcoin expensive or cheap right now, relative to its own history? It takes the gap between what the whole market is worth today and what everyone collectively paid for their coins, then scales that gap by how volatile the market has been, producing a smoothed oscillator, a measure that swings between a high "overvalued" extreme and a low "undervalued" extreme. It is closely related to plain MVRV, but the Z-Score version rescales the data so tops and bottoms line up on a consistent, comparable band.

THE Z-SCORE IN FOUR NUMBERS

>7	≤0	2018	~3–4
HISTORIC TOP BAND	HISTORIC BOTTOM BAND	METRIC INTRODUCED	CYCLES OF DATA

01

The two building blocks

Everything in the Z-Score rests on two figures. The first is market value, also called market capitalisation, which is simply Bitcoin's current price multiplied by the number of coins in circulation — the market's live opinion of what the whole network is worth. The second is realized value, also called realized capitalisation, a cost-basis measure that values each coin not at today's price but at the price it last moved on the blockchain, then sums those figures across every coin.

That second idea is the clever part. Because realized value only updates when a coin actually changes hands, it behaves like the network's aggregate purchase price — a rough record of what holders truly paid. Divide realized value by the coin supply and you get realized price, the average cost basis per coin. When market value sits far above realized value, holders are sitting on large paper gains; when it sinks toward or below realized value, the average holder is breaking even or losing money.

Because coins that never move keep their old cost basis, realized value changes slowly and calmly compared with the frantic swings of price. That smoothing is exactly why it makes such a useful anchor: it strips out the day-to-day noise of speculation and leaves behind a steadier picture of what the network genuinely paid, giving the Z-Score a stable baseline to measure against.

02

The formula, plainly

The Z-Score takes the difference between those two figures and divides it by a measure of how spread-out the market has been, the standard deviation, a statistic describing how far values typically stray from their own average.

MVRV Z-Score = (Market Value – Realized Value) ÷ standard deviation of Market Value

One detail is worth pinning down because it is frequently stated wrongly: the denominator, the number you divide by, is the standard deviation of market value, not of realized value and not of the gap between them. That distinction matters for anyone trying to reproduce the metric. It is also worth noting, in the spirit of honesty, that this is not a textbook statistical z-score drawn from a normal distribution; the metric's own author cautioned that the standard deviation here is used loosely as a scaling device, so treating the output as rigorous statistics overstates its precision.

03

How to read it: the red and green zones

The Z-Score is displayed as a line that oscillates over years. Analysts shade two bands onto it. When the score climbs into a high band, historically around 7 and above, Bitcoin has been richly valued relative to its cost basis — the market-top or red zone. When it falls into a low

band, historically near zero or slightly negative, the network has been cheap relative to what holders paid — the bottom or green zone.

These bands have coincided with major turning points: elevated readings clustered near the 2013, 2017 and 2021 peaks, while readings at or below zero appeared near the 2018 and 2022 lows. But the exact band edges are heuristics, useful rules of thumb rather than laws, and they are drawn slightly differently by each data provider. No specific value is a mechanical "sell" or "buy" button.

The 2021 cycle offered a textbook warning against reading the bands too literally. Bitcoin's April peak printed a higher Z-Score than its November peak, even though November reached a higher price — a divergence, meaning the valuation gauge quietly weakened while price pushed on. That kind of disagreement, where the score fails to confirm a new price high, has historically carried more information than any single threshold crossing.

04

Why bother with a Z-Score at all?

If plain MVRV — market value divided by realized value — already shows the same over- and undervaluation, why add the statistics? Because the raw ratio's extremes are hard to compare across a decade in which Bitcoin grew from a toy to a trillion-dollar asset. Standardising the gap by the standard deviation of market value pulls the outliers onto a common scale, so a peak in one cycle can be visually compared with a peak in another. The Z-Score, in short, is a normalised version of the same information, engineered to make cycle extremes legible at a glance.

There is a subtle trade-off in that engineering. The standard-deviation term is calculated over Bitcoin's whole price history, so it keeps drifting as new data arrives; the yardstick itself slowly stretches. This means the Z-Score is not a pure, fixed transformation of the MVRV ratio — the two can diverge modestly over long spans — and it is the main reason absolute Z-values from one charting provider should never be compared directly with those from another.

05

The honest caveat: it is not an independent signal

Here is the point most beginners miss, and it is the most important one in this report. The MVRV Z-Score, the plain MVRV ratio, realized price, and NUPL — Net Unrealized

Profit/Loss, a gauge of whether the network sits in aggregate profit or loss — are all built from the same two inputs: market value and realized value. They are re-expressions of a single underlying spread, not four separate opinions. In fact NUPL and MVRV are linked by exact arithmetic: NUPL equals one minus one divided by MVRV.

> **Why this matters.** If you see MVRV, the Z-Score, and NUPL all flashing "undervalued" and treat that as three confirmations, you are really seeing the same signal three times. Genuine confluence — multiple independent measures agreeing — requires pairing this family with something built from different data, such as miner activity, exchange flows, or price structure.

The Z-Score does add one ingredient the others lack — the standard-deviation term — so it is a rescaled cousin rather than an identical twin. But it is close enough that it should never be counted as an independent vote alongside MVRV or NUPL.

A simple analogy helps. Measuring the same room in feet, in metres, and in paces gives you three numbers, but only one measurement — converting between units does not make the wall longer. MVRV, its Z-Score, and NUPL are units of the same wall. The mistake is stacking them into a false sense of agreement, then sizing a position as if three separate witnesses had spoken when really one witness repeated itself three times.

06

The fading signal: diminishing amplitude

A second caveat undercuts naive threshold-watching. The Z-Score's peaks have shrunk with every cycle. Early cycles topped above 8; the 2021 peaks printed lower, in the high single digits; and through 2024–2025 the metric spent long stretches compressed in a low single-digit range rather than spiking. As Bitcoin's realized value grows enormous, the same dollar inflows move the score less, so a fixed rule such as "sell when it hits 7" grows less reliable over time.

This feeds a live debate. A growing body of analyst commentary argues that steady buying from exchange-traded funds — regulated funds that hold Bitcoin on investors' behalf — and from corporate treasuries may be permanently compressing the metric's amplitude, so this cycle could top well below the legacy red zone. Treat that as a widely discussed hypothesis, not settled fact; specific forward predictions built on it remain unproven.

A quieter distortion sits underneath the whole metric: lost coins. Bitcoin whose keys are gone forever still count in realized value at whatever price they last moved, often years ago at

trivial levels. That permanently understated cost basis nudges the Z-Score, and no one can measure the effect precisely, which is one more reason to read the number as directional rather than exact.

07

Who built it — the honest lineage

The Z-Score has a genuine pedigree, often garbled online. The foundation, realized capitalisation, was devised by Antoine Le Calvez and presented by Nic Carter of Coin Metrics in 2018. Building on it, analysts David Puell and Murad Mahmudov introduced the MVRV ratio that October. Days later, a pseudonymous researcher publishing as "Awe & Wonder" adapted it into the Z-Score. It was then popularised through Philip Swift's Look Into Bitcoin charts, which is why many wrongly credit him as the inventor. The original 2018 write-up carried a promotional claim of "90%+ accuracy" at calling tops; that figure is a backward-looking marketing line, not a validated forward statistic, and should not be repeated as fact.

08

How to actually use it

Reading	Rough meaning	What to remember
High band (>7)	Historically overvalued	Peaks shrinking each cycle
Mid-range (1–4)	Neither extreme	Where most time is spent
Low band (≤0)	Historically undervalued	Can stay low for months
Any reading	Cycle context, not a trigger	Not independent of MVRV/NUPL

Used well, the MVRV Z-Score is a cycle-position compass: it tells you roughly where the market sits between euphoria and exhaustion, not what will happen next week. It can be stretched to Ethereum and other on-chain assets, but their shorter histories and thinner liquidity make the bands far less stable, so Bitcoin's thresholds must never be copied onto an altcoin. Read it as one contextual input, pair it with genuinely different metrics, respect the shrinking amplitude, and it earns its place; treat it as a mechanical buy-and-sell switch, and it will eventually mislead you.

Perhaps the healthiest way to hold the metric is as a thermometer for crowd psychology. A deep-green reading tends to appear precisely when headlines are darkest and conviction is

thinnest, while a red reading arrives amid euphoria and certainty that “this time is different.” The number does not predict the future, but it does quantify how stretched sentiment has become — and that, used with discipline, is genuinely valuable context.

“A just balance and scales are the LORD’s; all the weights in the bag are his work.”

PROVERBS 16:11

METHODOLOGY & SOURCES

The formula, the market-value/realized-value definitions, and the standard-deviation denominator follow the canonical metric guides published by Glassnode. Attribution of realized capitalisation (Antoine Le Calvez & Nic Carter, Coin Metrics, 2018), the MVRV ratio (David Puell & Murad Mahmudov, October 2018), and the Z-Score adaptation (the pseudonymous “Awe & Wonder,” October 2018) reflects the documented public record; the identity behind that pseudonym is not established, and it is credited to popularisation via Philip Swift’s Look Into Bitcoin. Threshold bands (~7 and above; ~0 and below) are heuristic conventions that differ by provider and are stated as ranges, not fixed triggers. Historical extremes are described qualitatively because absolute Z-values are not comparable across providers with different standard-deviation windows, and the cycle sample (three to four completed cycles) is too small for statistical certainty. The “90%+ accuracy” framing from the original article is marketing, not a validated statistic. This report is educational research, not investment advice.