What Is Risk-Adjusted Return?

Why the size of a return means little until you know the risk that bought it — and how the ratios that measure it behave in crypto.

Alain AI Lab Research · Published July 23, 2026 · 10 min read

A return quoted on its own is half a sentence. Two assets can both have doubled over a year, yet one climbed steadily while the other lurched through stomach-churning swings and a brutal mid-year collapse. They are not equally good investments, and raw return cannot tell them apart. Risk-adjusted return is the missing half: it asks not just how much you made, but how much risk you carried to make it. It is the difference between admiring a number and understanding it — a distinct idea from the per-trade risk-reward ratio, which weighs a single trade's payoff, not a whole track record's efficiency.

AT A GLANCE

DEFINITION	CLASSIC MEASURE	SHARPE FORMULA
Return per unit of risk	**Sharpe ratio (1966)**	**(Return − risk-free) ÷ volatility**
DOWNSIDE VARIANT	**DRAWDOWN VARIANT**	**CRYPTO CAVEAT**
Sortino ratio	**Calmar ratio**	**Fat tails & 24/7 annualising**

01

Return per unit of risk

The core idea is a division. Take the return an investment earned, divide it by some measure of the risk it took, and you have a figure that lets you compare unlike things on even terms. A modest return earned calmly can be "better," on this basis, than a larger return earned through wild volatility, because the first delivered more reward for each unit of risk endured. Institutional investors reach for these figures precisely because they must compare a bond fund, an equity strategy, and a crypto position on one page, and only a risk-adjusted number puts them on a common scale.

Everything then turns on how you define the risk in the denominator. The most common choice is *volatility* — the standard deviation of returns, a measure of how widely they scatter around their average. But volatility is only one definition of risk, and much of the subject is really an argument about which denominator to use. The measures that follow are, at heart, the same numerator over different notions of risk.

02

The Sharpe ratio

The benchmark measure is the Sharpe ratio, introduced by William F. Sharpe in 1966 — he called it the “reward-to-variability ratio,” and the eponymous name came later. It divides *excess* return by total volatility: Sharpe = $(R_p - R_f) \div \sigma_p$, where R_p is the return, R_f is the risk-free rate, and σ_p is the standard deviation of returns. The higher the number, the more return the investment produced for each unit of total risk.

Two design choices are worth pausing on. The numerator is *excess* return — return above the risk-free rate, usually proxied by short-term government bills — because you deserve credit only for what you earned beyond what a near-riskless asset would have paid anyway. The denominator is total volatility, which treats every deviation from the average, up or down, as risk. That second choice is the ratio’s great strength and, as we will see, the source of most of its critics. (Sharpe later shared the 1990 Nobel Memorial Prize in Economics, though for his work on the Capital Asset Pricing Model, not for the ratio itself.)

03

Changing the denominator

If volatility is a contestable definition of risk, the natural response is to swap it out. Two well-known measures do exactly that. The *Sortino ratio*, developed by Frank Sortino in the early 1980s, replaces total volatility with *downside deviation* — the scatter of returns that fall below a target, ignoring upside swings entirely. Its logic is intuitive: an investor does not lie awake over gains, so why should a big winning month count as “risk”? It divides return above a target by downside deviation alone.

The *Treynor ratio*, from Jack Treynor’s mid-1960s work, takes a different turn. Its denominator is *beta* — sensitivity to the broad market — rather than total volatility, so it measures excess return per unit of *market* risk, on the assumption that asset-specific risk has already been diversified away. Where Sharpe judges a standalone holding, Treynor judges a piece of a diversified whole. The *information ratio* completes the set, dividing return above a benchmark by tracking error, to gauge how consistently a manager beats the index they are measured against. No single figure is credited with inventing it; it is a field-standard statistic of active management rather than one person's named creation.

04

Beyond volatility: drawdown

Volatility and its cousins all describe how returns *scatter*. They say nothing about the single figure that most often decides whether an investor actually stays invested: the *maximum drawdown*, the largest peak-to-trough fall before a new high. Drawdown is a different animal — it is loss-only, and it is *path-dependent*, because it depends on the order in which returns arrive, not merely their spread. A portfolio can be modestly volatile yet inflict a savage drawdown if its losses happen to cluster.

Because drawdown captures felt pain and capital impairment, a family of ratios is built on it. The best known is the *Calmar ratio* — annualised return divided by maximum drawdown, conventionally over about three years — generally credited to Terry Young around 1991, its name a contraction of his California Managed Account Reports newsletter. The closely related *MAR ratio* divides compound growth by the worst drawdown over a full track record. Both answer a question Sharpe cannot: how much return did you earn per unit of worst-case loss?

The distinction is not academic. Two funds can post the same Sharpe ratio while one recovered from a shallow dip and the other clawed back from a near-halving; the drawdown ratios separate them, the volatility ratios do not. This is why they are staples of the managed-futures and hedge-fund world, where an investor's tolerance for a deep hole — and the temptation to redeem at the bottom — is often the binding constraint, not the year-end return.

05

Alpha, beta, and skill

One more distinction underlies the whole field: the split between the return the market handed you and the return you actually added. *Beta* measures how much an asset moves with the market — a beta of one tracks it, above one amplifies it, below one dampens it — and represents systematic risk that diversification cannot remove. *Alpha* is what is left over: the return above what beta alone would predict, the part attributable to skill or selection rather than simply riding the market.

The formal version is *Jensen's alpha*, set out by Michael Jensen in 1968, which measures a portfolio's return against what the Capital Asset Pricing Model says it should have earned for its level of market risk. A positive alpha is genuine outperformance for the risk taken; a high raw return that is merely high beta is not skill at all, only leverage on the market's own movement dressed up as talent.

> *Every risk-adjusted measure is a single numerator — return — divided by a different answer to one question: what, exactly, is risk? Volatility, downside deviation, market beta, worst drawdown. Choose the denominator and you have chosen what you are afraid of.*

06

Where the Sharpe ratio misleads

The Sharpe ratio's dominance invites scrutiny, and four criticisms are well established. First, by using *total* volatility it penalises upside as harshly as downside: a strategy that occasionally posts huge gains looks "riskier" for doing so — the very objection that motivates Sortino. Second, standard deviation captures risk cleanly only when returns are roughly normally distributed; when they are skewed and fat-tailed, prone to rare, violent losses, a single deviation number understates the danger and the ratio flatters the asset.

Third, the ratio can be gamed. Strategies with smooth, stale, or infrequently priced returns show artificially low volatility and so an inflated Sharpe — a documented effect in illiquid assets — and payoffs that sell hidden optionality can look serene right up until they crash. Fourth, and most philosophical: volatility is not the same as the risk of permanent loss. Value investors have long argued that the fluctuation of the path matters far less than whether capital comes back at all, and by that lights a smooth ride to zero would score well on Sharpe until the very end.

07

Annualising, and the 24/7 problem

These ratios are usually quoted on an annual basis, which hides a convention that quietly breaks comparisons. Volatility scales with the square root of time — annual volatility is the per-period figure times the square root of the number of periods in a year — while return scales with the number of periods. To annualise, you must decide how many periods a year holds.

For stocks the answer is settled: about 252 trading days. Crypto has no such convention, because it never closes. Practitioners commonly annualise with 365 days — some use every hour of a 24/7 market — and the choice is not cosmetic. Because volatility scales by a square root, moving from 252 to 365 days multiplies an annualised Sharpe by roughly 1.2 on the scaling alone, so two "annual Sharpe ratios" computed under different conventions simply are not comparable. Worse, the square-root rule is exact only for independent returns; crypto's volatility clusters, so even the scaling is an approximation. The honest practice is to state the convention and the sample window every time.

08

Reading risk-adjusted return in crypto

Crypto is where these measures matter most and mislead most easily. Its assets have delivered large raw returns alongside very high volatility and deep drawdowns, so raw return badly overstates their appeal — risk adjustment is the honest comparison, both among coins and against traditional assets. But the high volatility sits in the denominator, so a spectacular return can still yield only a middling Sharpe: the ratio taxes the turbulence. And because crypto returns are markedly fat-tailed and skewed, standard deviation understates their tail risk, which is exactly why downside- and drawdown-based measures like Sortino and Calmar are often preferred here.

Several traps deserve naming. A ratio measured over a bull run looks nothing like the same ratio through a bear market, so the sample period is doing much of the work. Thinly traded tokens with stale prices show deceptively low volatility and inflated Sharpe ratios. There is no single consolidated price feed, so the same coin's numbers differ across venues. And samples that quietly drop dead tokens flatter the survivors. The disciplined reading, set out in our risk-management strategies for crypto investors, is to treat no single ratio as a verdict — read Sharpe, Sortino, and maximum drawdown together, disclose the convention and the window, and be most suspicious of the cleanest-looking number.

"Prove all things; hold fast that which is good."

Methodology & Sources

This report explains definitions and their properties, not the performance of any asset; it contains no Sharpe, volatility, return, or drawdown figures, all of which depend on the data source, sample window, and annualisation convention used. Formulas and attributions are given as documented: the Sharpe ratio to William F. Sharpe (1966), the Sortino ratio to Frank Sortino, the Treynor ratio to Jack Treynor, and Jensen's alpha to Michael Jensen (1968). The Calmar ratio is credited to Terry Young (c. 1991) on the accepted but lightly documented naming account. The four Sharpe critiques — upside penalty, non-normality, gameability under illiquidity, and volatility versus permanent loss — are standard in the literature. A clarification: the Sharpe ratio can always be computed; normality is what makes volatility a *sufficient* risk measure, so the fat-tail objection is that standard deviation understates tail risk, not that the ratio is uncalculable. Related reading: position sizing and risk-adjusted entry tiers.

Alain AI Lab — independent crypto & AI research.

intelligencecrypto.org · Research is educational and analytical, not financial advice.