What Is the Kelly Criterion?

The formula that sizes a bet to maximise long-run growth — and why crypto traders should treat it with caution.

Alain AI Lab Research · Published July 20, 2026 · 10 min read

The Kelly criterion is a formula that answers a single question with unusual precision: given an edge and the odds, what fraction of your capital should you stake on a bet to grow your money as fast as possible over the long run? It is one of the few tools in finance that produces an exact number rather than a rule of thumb, and that precision is both its appeal and its trap. Used carefully it is the mathematical ceiling on rational bet size; used naively — especially in a market as noisy as crypto — it is a fast route to a blown account.

This report explains where the formula comes from, what it actually optimises, how to compute it, and why almost no serious practitioner bets the full amount it prescribes. It closes with why crypto's particular features quietly break the assumptions the formula depends on. It sits alongside our broader risk-management strategies for crypto investors, which frame sizing as one control among several.

AT A GLANCE

INTRODUCED BY	WHAT IT MAXIMISES	TRADING FORMULA
John L. Kelly Jr., 1956	**Long-run growth (log wealth)**	**$f^* = W - (1-W) \div R$**
REQUIRES	**COMMON PRACTICE**	**KEY DANGER**
A real, positive edge	**Half or quarter Kelly**	**Overbetting → ruin**

01

Where the formula came from

The criterion was introduced by John L. Kelly Jr., a researcher at Bell Labs, in a 1956 paper in the *Bell System Technical Journal* titled “A New Interpretation of Information Rate.”

Kelly derived it from the information theory of his colleague Claude Shannon, reframing the problem of transmitting signals over a noisy channel as the problem of a gambler with private information deciding how much to wager. Shannon inspired the work but did not co-author it; the formula is Kelly's.

What began as an argument about information rates became, in the hands of later practitioners, a general theory of capital growth. The insight that survived the translation is simple to state and hard to internalise: the right bet size is not the one that makes the most money on the next trade, but the one that compounds fastest across thousands of them.

02

What it actually maximises

The subtle point that defeats most newcomers is what Kelly optimises. It does not maximise your expected wealth — the simple average of outcomes — because that objective would tell you to bet everything on any favourable edge, which guarantees eventual ruin the first time you lose. Instead, Kelly maximises the expected *logarithm* of wealth, which is mathematically equivalent to maximising the long-run geometric growth rate of your capital. Geometric growth is the rate that compounding actually follows, and it punishes large losses far more heavily than an arithmetic average does.

This is why the formula refuses to let you bet the whole bankroll even on a strong edge. Log-wealth maximisation builds in a deep respect for the fact that a 50% loss requires a 100% gain merely to break even, and that a 90% loss demands a 900% gain that almost never arrives. The criterion is, at heart, a formal statement that surviving to keep compounding matters more than winning big once, because a single ruinous bet ends the sequence of trades on which all future growth depends.

03

The formula, stated plainly

In its original gambling form the criterion is $\mathbf{f^* = (bp - q) \div b}$, where f^* is the fraction of the bankroll to stake, $\mathbf{p}$ is the probability of winning, $\mathbf{q}$ is the probability of losing $(1 - p)$, and $\mathbf{b}$ is the *net* odds — the profit received per unit staked, not the gross payout. That distinction matters: reading b as the total return rather than the net gain is a common and expensive error.

Traders more often use an equivalent restatement, $\mathbf{f^* = W - (1 - W) \div R}$, where **W** is the win probability and **R** is the payoff ratio — the average win divided by the average loss. For an even-money bet the whole thing collapses to $f^* = 2p - 1$. Two cautions travel with the trading form: it is exact only for a clean, binary win-or-lose bet, and it becomes an approximation once R is an average of variable real-world outcomes rather than a fixed payout.

04

A worked example

Suppose a strategy wins 60% of the time and its average win equals its average loss, so $W = 0.6$ and $R = 1$. The formula gives $f^* = 0.6 - (0.4 \div 1) = 0.2$ — Kelly says risk 20% of the bankroll. Now suppose a more selective strategy wins only 55% of the time but its winners are twice its losers, so $W = 0.55$ and $R = 2$. Then $f^* = 0.55 - (0.45 \div 2) = 0.325$, or 32.5%. A lower win rate paired with a larger payoff earns a larger prescribed bet.

The formula also enforces its own precondition. If your edge is zero or negative, f* comes out at zero or below, which the criterion reads as an instruction to bet nothing at all. Kelly never recommends a position without a genuine, positive expectancy behind it — a discipline that already rules out most impulses to trade.

> *Bet more than Kelly and you take on more risk to earn less growth. Overbetting is the rare position in finance that is strictly worse in both directions at once — the surest way to turn a real edge into a losing strategy.*

05

Why almost no one bets full Kelly

The full Kelly fraction is growth-optimal in theory and brutal in practice. Because it maximises growth without regard to short-term comfort, it produces enormous equity swings. A well-known property of the strategy, described in the academic literature by MacLean, Thorp, and Ziemba, captures the point vividly: a full-Kelly bettor faces roughly a 40% chance that their capital will at some stage fall to 40% of its starting value, and more generally about a one-in-two-fifths kind of exposure to deep temporary declines. Few humans can watch that happen and hold the line.

The drawdowns are not a sign the formula is wrong; they are the price of maximal growth. But the price is one most traders and investors are unwilling — and often financially unable — to pay, because a strategy that is mathematically optimal is worthless if you abandon it at the bottom of a decline. The behaviour the number demands, in other words, is one that human psychology and real capital constraints rarely permit. This gap between theoretical optimality and human tolerance is precisely why the full number is treated as a ceiling rather than a target.

06

Fractional Kelly and the overbetting trap

The standard response is to bet a fraction of the Kelly amount — commonly half or a quarter. The trade-off is favourable because the growth curve is nearly flat around its peak: half Kelly is often summarised as retaining roughly three-quarters of the long-run growth while cutting the volatility by about half, though these figures are approximate and depend on the specific payoffs. You give up a little compounding for a large reduction in the swings, which is a bargain most people will actually stick to.

Overbetting runs the other way and is unforgiving. In the simple symmetric case, staking twice the Kelly fraction drives your excess long-run growth to zero — you carry all the risk and compound at nothing — and beyond that point growth turns negative and ruin becomes near-certain despite a real edge. The curve is asymmetric, so overshooting is far more destructive than the same-sized undershoot. That asymmetry has a practical consequence: since your true edge is always estimated, and overestimating it pushes you into the destructive zone, the safer error is always to bet too little.

07

Thorp, Samuelson, and the long debate

The criterion left the engineering journals largely through Edward O. Thorp, the mathematician who used it to size bets while card-counting at blackjack — work he published in *Beat the Dealer* in 1962 — after Claude Shannon pointed him to Kelly's paper in 1960. Thorp later carried the same logic into markets through his hedge fund, renamed Princeton/Newport Partners, whose long run of strong reported returns became part of the formula's legend. Kelly is now a recognised tool across gambling, blackjack, and investment position sizing alike.

It has never been free of serious objection. The economist Paul Samuelson argued repeatedly — in papers in 1971 and, famously, in 1979 using almost entirely one-syllable words — that maximising the geometric mean is not a universal rule: it is optimal only for an investor whose preferences happen to match log utility, and anyone more risk-averse than that will find full Kelly overbets. The dispute is genuine and unresolved, which is the honest frame for the formula: a powerful growth criterion, not a law that every rational actor must obey.

08

Does Kelly work in crypto?

Kelly's guarantees rest on assumptions that crypto markets strain or break. The formula needs a true, stable win probability and payoff ratio; crypto edges are hard to estimate and shift with the regime, so the inputs are unreliable and the ever-present risk is that overestimating your edge quietly pushes you into overbetting. Crypto returns are also famously fat-tailed and far from the clean win-or-lose model the formula assumes, which makes a naively computed fraction less trustworthy than its decimal precision suggests.

Leverage is the sharpest problem. A core Kelly assumption is that you cannot lose more than you stake, but a leveraged position can be liquidated — force-closed at the maintenance-margin level against the exchange's mark price — and with slippage or gaps the realised loss can exceed the margin posted, violating the assumption outright. Correlation compounds the danger: sizing several alts by single-bet Kelly ignores that they tend to move together, so the real aggregate bet is far larger than any one line suggests, and trending or mean-reverting price action violates the independence the simple formula assumes. The reasonable practice, and the common view among disciplined traders, is to treat Kelly as a conceptual upper bound — a discipline against overbetting — and in practice to use a heavily fractional version or a flat, small fixed percentage of capital per trade, the same ground covered in our note on position sizing and risk-adjusted entry tiers. Overbetting, whatever its source, remains one of the most common crypto mistakes. This material is educational and analytical, not financial advice.

> *"Wealth gotten by vanity shall be diminished: but he that gathereth by labour shall increase."*
>
> PROVERBS 13:11

Methodology & Sources

This report draws on the primary source — J. L. Kelly Jr., “A New Interpretation of Information Rate,” *Bell System Technical Journal* (1956) — and on the standard practitioner and academic literature, including MacLean, Thorp, and Ziemba on the good and bad properties of the Kelly criterion, and Paul Samuelson’s 1971 and 1979 objections to geometric-mean maximisation. Formulas are stated in their classical forms; the trading restatement $f^* = W - (1-W) \div R$ is exact for a binary bet and approximate for variable real-world returns. Figures such as “half Kelly retains roughly three-quarters of growth” are well-known approximations, not universal constants, and are labelled as such. Crypto-specific limitations — unstable edge, fat tails, liquidation breaking the bounded-loss assumption, and correlation across assets — are described qualitatively; no specific volatility, correlation, funding, or fund-performance figures are asserted. Related reading: risk-management strategies, position sizing and entry tiers, and common crypto mistakes. Research is educational and analytical, not financial advice.