

What Is the Bid-Ask Spread?

The permanent gap between what buyers offer and sellers ask — how to measure it properly, the hidden reason it can never be zero, and why you pay it twice.

Alain AI Lab Research · Published July 25, 2026 · 10 min read

The bid-ask spread is the gap between the highest price anyone is currently willing to pay for an asset and the lowest price anyone is currently willing to accept for it. Those two numbers are never the same, and the space between them is the most reliable cost in trading — charged on every trade, in every market, whether or not you notice it. Most people glance at a single price and assume that is what they will pay. In fact there are always two prices, you get the worse one, and this report is about measuring that properly: how to express the spread, what actually sets its width, and why it costs you twice on every position you open and close.

AT A GLANCE

BID

Best price a buyer will pay

MEASURE IT AS

A percentage, not an amount

ASK

Best price a seller will take

YOU PAY IT

Twice — entry and exit

THE SPREAD

The gap between them

NEVER ZERO

Adverse selection sets a floor

01

Two prices, and the one you get

Every tradable market quotes two prices at once. The *bid* is the highest price a buyer is currently prepared to pay — so it is the price at which *you* can sell immediately. The *ask*, sometimes called the offer, is the lowest price a seller is prepared to accept — so it is the price at which you can buy immediately. Note the inversion that trips people up: you sell at

the bid and buy at the ask, which means you always transact at the less favourable of the two.

The single price displayed on a chart or a price site is usually the last traded price, or the midpoint between bid and ask — a useful summary and a number you can rarely actually trade at. This is why an asset can appear to be at one price while your buy fills slightly higher and your sale fills slightly lower. Nothing is wrong; you are simply seeing a convenient average of two real prices, only one of which applies to you at any moment. Where those two prices sit and how they are ranked is the structure covered in [how order books work](#).

02

Measuring it the right way

The spread is quoted two ways, and only one of them is useful for comparison. The *absolute* spread is simply the ask minus the bid, expressed in currency — easy to read and almost meaningless across assets, because the same absolute gap is trivial on an expensive asset and enormous on a cheap one.

The *relative* or percentage spread fixes this by dividing the gap by the midpoint price. Now you have a figure that compares fairly across any two assets regardless of their price level, which is the only way to answer the question that matters: *is this market expensive to trade?* Professionals often express the same thing in basis points, where one basis point is a hundredth of a percent, purely because the numbers involved are small enough that percentages become awkward. The practical instruction is simple — when you look at a spread, convert it to a percentage of the price. An asset whose spread is a small fraction of a percent is cheap to trade; one whose spread runs to several percent is expensive, and no amount of enthusiasm about the project changes that arithmetic.

03

The spread you see and the spread you pay

There is a further distinction that matters once your orders are larger than the smallest quoted size. The *quoted* spread is the gap between the best bid and best ask — the headline figure. The *effective* spread is what you actually paid, measured as the distance between your execution price and the midpoint at the moment you traded.

These diverge for a simple reason: the best quoted prices are only good for a limited quantity. An order larger than that reaches into worse prices to complete, so your effective

cost exceeds the quoted spread — the point at which spread and price impact blur into one another. The practical consequence is that the quoted spread is a *best case*, achievable only for small orders, and the number worth tracking is the one you actually realise. A market can advertise an attractive quoted spread and deliver a much worse effective one to anybody trading meaningful size, which is why depth and spread must always be read together rather than either alone.

04

Why it can never be zero

Here is the part rarely explained. It is tempting to think competition should eventually drive the spread to nothing. It cannot, and the reason is *adverse selection* — the risk that the person trading against you knows something you do not.

Consider the position of whoever is quoting both prices. Most of the people who trade against their quotes are ordinary participants with no special information, and serving those trades profitably is straightforward. But some are trading precisely because they know something — news not yet reflected in the price, a large order about to arrive, an insight the market has not absorbed. Against those counterparties, the quoter systematically loses: they buy just before the price falls and sell just before it rises. The spread is the charge levied on everyone to cover the losses inflicted by the informed few. This is why spreads widen when information is uncertain, why they blow out before major announcements, and why they can never compress to zero no matter how fierce the competition. Underneath the spread sits a permanent, irreducible cost of not knowing who is on the other side — a cost that sits alongside the inventory risk described in [what is a market maker](#).

The spread is not a fee somebody decided to charge. It is the price of two genuine risks: holding an asset whose value moves, and trading against people who may know more than you. Neither ever goes away, so neither does the spread.

05

What sets the width

Four forces determine how wide a spread runs, and they move together. *Volatility* is the most immediate: the faster a price moves, the more risk in standing ready to trade, so quotes widen defensively the moment a market becomes turbulent. *Activity* works the other

way — heavily traded assets let a quoter turn over inventory quickly and recover costs across many transactions, so busy markets support tighter spreads than quiet ones.

Competition matters as much as either: several parties quoting the same asset undercut each other toward the narrowest sustainable gap, while a market served by one participant has no such pressure. And the asset's own character sets a floor — a large, well-understood asset with widely agreed value carries less adverse-selection risk than an obscure token whose fair price is genuinely uncertain. Taken together these explain the pattern you will observe in practice: spreads are tightest on the biggest assets during active hours, and widest on small assets during quiet periods or violent moves — which is precisely when an impatient trader is most likely to be crossing them.

06

You pay it twice

The most consequential fact about the spread is the one least often stated: a complete trade crosses it twice. You buy at the ask, and to realise anything you must eventually sell at the bid — so the round trip costs you the full spread, before fees and before the asset has moved at all.

This sets a hurdle. The moment you enter a position you are behind by the spread, and the price must move in your favour by at least that much before you are level. On a deeply liquid asset the hurdle is negligible. On a thin one it can be several percent, meaning a trade must be substantially right just to break even — and that hurdle is paid *every time* you trade, so a strategy that turns over frequently pays it repeatedly, compounding into a serious drag that never appears on any statement. This is the arithmetic reason that frequent trading in illiquid assets is so reliably unprofitable, and it belongs in any assessment of whether a return justifies its risks, as discussed in [risk-adjusted return](#).

07

Using it as a signal

Beyond being a cost, the spread is a free and instant diagnostic, and learning to read it takes seconds. A tight spread signals an actively traded asset with several parties competing to serve it — a market in good health. A wide one signals the reverse: few participants, uncertain valuation, or elevated risk, and it is often the single quickest tell that an asset is thinner than its marketing suggests.

Two habits make this useful. First, check the spread *before* you buy anything, not after — it takes a moment and immediately tells you the minimum cost of the round trip you are contemplating, which is information you cannot get from a price chart. Second, watch how it changes: a spread that widens sharply is telling you that whoever was providing liquidity has become nervous, which is often the earliest visible sign of stress in an asset, well before it shows in the price. The broader use of these signals is covered in [market liquidity](#), and where you trade affects what you are quoted, as our note on [choosing an exchange](#) discusses.

08

Spreads in crypto, and on-chain

Crypto exhibits the widest range of spreads of any market most people will encounter. The largest assets on major venues trade at spreads comparable to serious traditional markets; small tokens can quote gaps of several percent or worse. Because the market never closes, spreads also vary by time of day even without a formal session — widening during hours when fewer participants are active, which is a real cost for anyone who trades at inconvenient times. And because the same asset trades on many venues at once, spreads differ across them simultaneously, so the venue you choose changes the price you get.

On decentralised exchanges the concept transforms. An automated market maker has no bids and no asks — it prices from a formula against a pool, so there is no quoted spread in the traditional sense. What plays the equivalent role is the pool's trading fee combined with the price impact of your trade against the curve, described in [automated market makers](#). The economic function is identical — a cost you pay for immediacy, wider on thin pools than deep ones — but the mechanism is entirely different, so comparing an order-book spread with an AMM fee tier requires care. Whichever venue you use, the discipline is the same: find out what the round trip costs before you commit to it.

“It is naught, it is naught, saith the buyer: but when he is gone his way, then he boasteth.”

PROVERBS 20:14

This report explains a market concept and how to measure it; it contains no spread figures, basis-point values, asset names, or venue names, all of which vary continuously by market and moment — observe current quotes directly. The definitions used are standard: the bid as the highest standing buy price and the ask as the lowest standing sell price, with the trader selling at the bid and buying at the ask; the absolute spread as their difference and the relative spread as that difference divided by the midpoint; and the effective spread as the realised distance between execution price and prevailing midpoint, which exceeds the quoted spread for orders larger than the quoted size. Adverse selection is presented as a long-established explanation in market microstructure for why spreads carry an irreducible floor: quoting parties lose systematically to better-informed counterparties and recover those losses across all trades. The four drivers of spread width — volatility, activity, competition, and asset uncertainty — are described directionally, not quantified. The comparison with automated market makers is structural: AMMs quote no bid or ask, and the economically equivalent cost is the pool fee combined with price impact against the pricing curve. Nothing here is a recommendation regarding any asset, venue, or trade.

Alain AI Lab — independent crypto & AI research.

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

© 2026 Alain AI Lab. All rights reserved.