

How to Ladder Buy Orders

The order-book mechanics of a staged entry — setting the range, spacing the rungs, sizing each one, and managing a ladder that is half filled when everything changes.

Alain AI Lab Research · Published July 24, 2026 · 10 min read

A buy ladder is a set of limit orders resting at successively lower prices, each one a rung, together making up a single intended position. Instead of choosing one entry price and hoping, you cover a range and let the market decide how much of your order it wants to fill. It is the most common way a staged entry actually gets implemented — and it is where the plan meets the order book, which is a less forgiving place than a spreadsheet. This report is about the mechanics: how to place the rungs, space them, size them, and manage them once they are live. The strategic question of whether to stage an entry at all, and how to budget its risk, is covered in [how to scale into a position](#).

AT A GLANCE

WHAT IT IS

Resting limit orders at set prices

SPACING

Percentage beats fixed steps

YOU CONTROL

Price — not whether you fill

FULL FILL MEANS

Price went through your range

DEFINE FIRST

Top rung, bottom rung, total

CRYPTO TWIST

Wicks, minimums, fee per rung

01

What a ladder is, mechanically

A limit buy is an instruction to purchase at a specified price *or better*, and it sits in the order book until it fills, expires, or is cancelled. A ladder is simply several of these placed at once at different prices below the current market, dividing one intended position across a range instead of concentrating it at a point.

The defining property follows directly from the order type: a limit order gives you control over price and no control over execution. A market order is the reverse — it fills immediately at whatever the book offers. So a ladder trades certainty of getting in for certainty of what you pay. If price falls through your entire range you own the full position with a better average than a single top-of-range buy would have given you; if price never reaches your first rung you own nothing at all, having been exactly right about the asset and entirely unfilled. Those two outcomes are not edge cases. They are the two ends of what you are choosing between.

02

The range is the real decision

Everything else in a ladder is arithmetic; the range is judgement. Two prices define it, and each carries a distinct meaning.

The top rung is where you begin to want the asset — the price at which you would be content to own some, knowing it may go lower. Place it too far below the market and the ladder never engages during a shallow pullback; place it at the market and you have simply bought, with rungs below as an afterthought. The bottom rung matters more, because it is a statement about where your interest ends. It should sit above the level at which your reason for buying would be broken, never below it. A ladder whose lowest rung sits under your invalidation is not an entry plan — it is an instruction to keep buying something after you have agreed it is no longer worth owning, which is the most expensive way to automate a mistake. Set the bottom rung first, then the top, and let the distance between them tell you how wide a range you are actually underwriting.

03

Spacing the rungs

With a range fixed, spacing decides where the individual orders sit inside it. There are two conventions and they behave differently.

Linear spacing places rungs an equal number of price units apart. It is simple and reads cleanly on a chart, but it distorts as price falls: a fixed step is a small proportion of a high price and a large proportion of a low one, so lower rungs end up bunched in percentage terms. *Geometric* or percentage spacing places each rung a constant percentage below the last, which keeps the economic distance between rungs constant all the way down. For

volatile assets, percentage spacing is the more coherent choice, because your exposure between rungs is what you actually care about, not the absolute price gap.

Two further points govern the choice. The spacing should be wide relative to the asset's ordinary noise — rungs packed inside the range price wanders through on a quiet afternoon will all fill on nothing, converting a considered plan into a market buy executed slowly. And the number of rungs should stay small; each one is a separate trade with its own cost, and beyond a handful the marginal improvement in average price is quickly eaten by fees and complexity.

04

Sizing each rung

The rungs must sum to the position you already decided to hold — not to a total discovered afterwards by adding up what happened to fill. Within that fixed total, three distributions are common. *Flat* puts identical size on each rung: neutral, easy to reason about, and the sensible default. *Front-weighted* puts more at the top: you own more if the dip is shallow, less if it is deep. *Back-weighted* puts more at the bottom: it produces the best average price when the range fills completely, at the cost of holding very little through a shallow retracement.

Back-weighting is the most seductive and the most dangerous, because the scenario it optimises for — price sweeping the whole range — is also the scenario in which something may genuinely be wrong. Committing your largest tranche to the lowest rung means committing most heavily exactly when the market disagrees with you most strongly. That may be correct, but it should be a decision you made deliberately rather than a by-product of wanting a flattering average entry.

A ladder that fills completely is not a ladder that worked. It is the market telling you it was willing to sell you everything you asked for, at every price you named.

05

Managing a live ladder

Once the orders are resting, the work shifts to maintenance, and three situations recur.

The first is the partial fill: some rungs execute and the rest do not, leaving a position smaller than intended with orders still open beneath it. This is the normal state of a working ladder, and the discipline is to treat the unfilled rungs as still-valid commitments rather than as a decision to revisit each time price moves. The second is a change in the underlying reason for the trade. If the thesis breaks while the ladder is live, the correct action is to cancel the remaining rungs immediately — resting orders do not know your news, and they will keep buying an asset you no longer want. The third is expiry and hygiene: orders left open for weeks can fill during a move you were not watching and no longer agree with, so ladders need a lifespan, a review date, or both. On most venues an open limit order also reserves the funds behind it, so a stale ladder quietly locks up capital you may want elsewhere.

06

Fees, minimums, and the cost of rungs

Splitting one purchase into several changes what you pay in ways worth planning around. On most venues a resting limit order that adds liquidity is charged the maker fee, which is typically lower than the taker fee paid by a market order — so a ladder often costs less per unit traded than an immediate buy, one of the genuine advantages of the technique. Many exchanges also offer a post-only flag that cancels the order rather than letting it execute as a taker, which preserves that treatment.

Working the other way are the minimums. Every venue enforces a smallest permissible order size, and dividing a modest position into many rungs can push individual orders below it, so the ladder simply cannot be placed as designed. Fees also have fixed components on some venues, and each rung pays them separately. The practical rule is to choose the rung count from what the position and the venue can actually support, rather than picking a satisfying number and discovering afterwards that the bottom third of the ladder is unplaceable.

07

Automating it

Ladders are well suited to automation because they are entirely rule-based. Most centralised exchanges support placing all the rungs manually as ordinary limit orders, and many now offer grid tools that maintain a lattice of buys and sells across a range automatically. A grid is a ladder that also sells — a strategy for ranging markets rather than a staged entry — and the distinction matters: a grid bot is designed to trade a range repeatedly, not to build one

position and stop. If your goal is a single accumulated holding, the grid's sell side is working against you.

On-chain the picture differs, because automated market makers execute against a curve rather than an order book and have no native concept of a resting limit order. What exists instead are protocol- or keeper-based limit orders that fill when someone triggers them, and range-based liquidity provision, which behaves somewhat like a ladder in that it accumulates the asset as price moves through a chosen band — with different mechanics and its own exposures, as covered in our note on [automated market makers](#). The other constraint on-chain is gas: every rung is a transaction, so a many-runged ladder on a congested network can lose more to network fees than it gains in average price.

08

Where ladders go wrong in crypto

The failure modes are specific and mostly avoidable. Sharp wicks are the most common: a violent spike downward can sweep an entire ladder in seconds and reverse just as fast, leaving you fully positioned at an average you would not have chosen — a reason to keep the bottom rung meaningful rather than parked at a fantasy price you would regret being filled at. Thin books do the opposite damage, filling your lower rungs on trades so small they say nothing about genuine demand. Because the market runs continuously, both happen while you sleep, which is the point of resting orders but also the reason a ladder must be one you would accept waking up fully filled into.

Three quieter costs deserve mention. Each rung is a separate tax lot with its own cost basis and acquisition date, which complicates record-keeping and later disposals in ways that vary by jurisdiction. Capital reserved behind unfilled orders is capital not available elsewhere. And laddering into an asset in genuine structural decline does not rescue the decision — a better average price on something that keeps falling is still a loss, more efficiently acquired. The ladder is an execution tool. It improves how you enter; it cannot improve what you chose to enter, which remains a question of [sizing and selection](#) rather than order placement.

“Wealth gotten by vanity shall be diminished: but he that gathereth by labour shall increase.”

PROVERBS 13:11

Methodology & Sources

This report covers order mechanics; it contains no rung counts, spacing percentages, range widths, fee schedules, or gas figures, all of which depend on venue, asset, and market conditions and are stated only in direction. Maker and taker fee treatment, post-only behaviour, minimum order sizes, and fund reservation on open orders are standard exchange conventions that vary by platform — check your venue's own schedule rather than assuming. The distinction drawn between a buy ladder and a grid is functional: the former builds one position, the latter trades a range in both directions. On-chain, automated market makers have no native resting limit orders; keeper-executed limit orders and range liquidity are described as different mechanisms with their own risks, not equivalents. Tax lot treatment varies by jurisdiction. For the strategy this executes, see [how to scale into a position](#); for scheduled, price-indifferent accumulation, see [dollar-cost averaging in crypto](#).

Alain AI Lab — independent crypto & AI research.

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

© 2026 Alain AI Lab. All rights reserved.