

The Drawdown Ladder

Can a 1-2-3 recovery ladder profitably buy the pullbacks of mature \$1M+ memecoins? The answer is yes — but only after three evidence-driven corrections to the entry rules. As originally specced, the ladder loses money on every parameter tested.

THE VERDICT

The raw ladder (buy the dip, double down, triple down, tight stop) backtests at **-12 to -14% per \$100 deployed** — the martingale tail is real: half of all campaigns filled the full ladder and the tight stop was triggered by normal volatility, not failure.

Three corrections flip it: **wait for a bounce confirmation before the first buy, only enter below the average holder cost, and use a wide close-based stop**. The corrected ladder backtests at **+24.6% per \$100 deployed** (58% wins, 7% stop-outs), with a steadier variant at **83% wins** and +14%.

Your two core instincts survived contact with the data: **deeper entries genuinely have better bounce odds**, and **buying below the crowd's cost basis genuinely works** — it was the knife-catching entry timing and the too-tight stop that were bleeding the strategy.

208

tokens that crossed \$1M market cap, fully analyzed (of 309 collected)

63%

went on to print a HIGHER all-time high after their first major top

-69%

median maximum drawdown from the first \$1M+ top (quartiles -54% / -85%)

79%

chance of a +25% bounce within 14 days when entering at -80% from the top

01 What was studied

The cohort is the 309 contract addresses you collected from pump.fun's explore tables — the exact population this strategy trades. For each token we pulled its hourly candles (≈3 weeks) and 4-hour candles (≈12 weeks) — 210,953 OHLC candles total. **208**

tokens crossed \$1M market cap with a confirmed top (a peak followed by a $\geq 30\%$ drawdown) and form the study set.

Every number below is computed, not estimated: tops and bottoms from the candle series, bounce odds from forward windows after simulated entries, and the ladder results from bar-by-bar simulation of the full buy/stop/take-profit sequence on 1-hour candles.

02 The shape of a mature token's life

63% print a higher ATH after their first major top. This is the most surprising number in the study and it is good news for the strategy: these are not one-pump corpses. A majority of tokens that reach \$1M and survive their first dump eventually trade above the old high — the dip you are buying usually is a dip, not a funeral.


But the dips are violent. The median token falls **-69%** from its first big top before bottoming; a quarter fall **-85% or more**. This immediately explains why first entries at -40% or -50% "miss the bottom" so often — statistically, they almost always do. The bottom usually forms in the -60% to -85% zone, which is where the ladder's later rungs belong.

03 Bounce odds by entry depth

The deeper the entry below the top, the better every single bounce statistic — monotonically. And from *any* entry, the median token still falls another 21–29% before turning: that further-fall band is the number that must set your DCA spacing and your stop, and it is exactly why a 15% stop kept getting hit.

Probability of a +25% bounce within 14 days, by entry depth below the top

208 qualified tokens · entries simulated at each depth the token actually reached



ENTRY DEPTH	TOKENS REACHING IT	+15% IN 7D	+25% IN 7D	+25% IN 14D	+50% IN 14D
-40%	184	67%	51%	61%	41%
-50%	166	71%	59%	65%	44%
-60%	142	77%	67%	74%	56%
-70%	97	80%	76%	77%	62%
-80%	68	81%	76%	79%	76%

Read the last column carefully: it is not "how far the losers fall" — it is the median further drop across ALL entries at that depth, winners included. The bounce usually comes, but it usually comes from ~20–29% below wherever you got in. Plan for that; don't stop out inside it.

04 Your cost-basis instinct, tested

You look for final entries below the top holders' average buy price. We tested a computable proxy — the token's volume-weighted average price (VWAP), the average cost of all money that ever entered — and the mechanism is real:

ENTRY AT -40% FROM TOP	SAMPLE	+25% BOUNCE WITHIN 14D
Below the crowd's average cost	108	70%
Above the crowd's average cost	76	47%

A **23-point edge** from one filter. The mechanism: above the crowd's cost, every rally hits a wall of underwater holders selling at breakeven; below it, that overhead supply thins out. The filter matters most at shallow entries — by -70% almost everything is below VWAP anyway. This validates your rule and upgrades it: **don't just place the last buy below the crowd's cost — refuse to START a campaign above it.**

05 Backtesting the ladder — and fixing it

The full 1-2-3 ladder was simulated bar-by-bar on hourly candles: first buy at -X% below the rolling 7-day peak, second buy (2×) lower, third buy (3×) lower still, stop below the last rung, take-profit on the blended cost basis, 14-day time limit. Three versions:

Return per \$100 deployed — the ladder's three generations

best parameter set of each generation · same 208-token history



Fix 1 — the stop was inside the noise. A 15% stop sits inside the normal 21–29% post-entry chop, so it converted ordinary volatility into realized losses (40–52% of all campaigns stopped). Widened to 30% and triggered on a candle *close*, not a wick.

Fix 2 — never start a campaign above the crowd's cost. The VWAP filter from section 04, applied as an entry precondition. Cuts campaign count ~25% and removes the worst cohort.

Fix 3 — the big one: stop catching knives. As specced, buy #1 fires the instant price touches $-X\%$. But the median token falls another quarter after that touch. v3 waits: after the level is touched, the first buy only fires once price *closes 5% above its running low* — you buy the confirmed turn, not the fall. This one change collapsed the full-ladder-stop-out tail from ~50% of campaigns to 7%, because most campaigns now never need the deep rungs at all.

ENTRY	TAKE PROFIT	CAMPAIGNS	WINS	STOPS	TIME-OUTS	FULL LADDER
-50% + confirm	+60%	100	58%	7%	35%	
-60% + confirm	+60%	103	63%	7%	30%	
-60% + confirm	+40%	113	75%	4%	20%	
-60% + confirm	+25%	130	83%	3%	14%	
-60% raw (v1 best)	+25%	289	53%	40%	7%	

All v3 rows use: DCA spacing 25% then 30% below the prior rung, 30% close-based stop under rung 3, VWAP entry filter, 5% bounce confirmation, 14-day limit. "Time-outs" exit at market on day 14 — their (usually small) win or loss is included in the EV.

06 The playbook

Translated back into your workflow, at your \$400-scale example (ladder units of \$65/\$130/\$195, ~\$390 max deployed — in practice most campaigns only ever use the first unit):

1. **Universe:** mature tokens that have printed $\geq \$1\text{M}$ market cap and survived a confirmed top. Remember the base rate is on your side here — 63% eventually print a higher high.
2. **Arm a campaign** when price reaches **-50% to -60% below the recent major top** — and only if that price is **below the token's average holder cost** (GMGN's top-100 avg-cost line is your live proxy for this). Above it: no campaign, walk away.
3. **Do not buy the touch.** Wait for the turn: on the 1h chart, price must close **~5% above its low** since the level was hit. Then buy unit 1. This single habit is worth more than every other parameter combined.
4. **Recovery buys:** unit 2 (2×) at **-25%** below your first fill; unit 3 (3×) at **-30%** below that. These are standing bids — the 21–29% further-fall band means they fill often and that's fine; it's what they're for.
5. **Stop:** a 1h *close* **30% below rung 3**. It feels wide; the data says anything tighter sells the bottom to someone else. It fires in only ~7% of campaigns.
6. **Exit:** take profit at **+40% to +60% on your blended basis** for max EV, or +25% if you prefer the 83%-win rhythm. If 14 days pass without either, close it and move on — dead time is its own cost.

HONEST LIMITS OF THIS STUDY

Survivorship lean: the cohort came from explore tables, which skew toward tokens alive when collected. Bounce odds on a fully unbiased cohort (including every token that died post-ATH) would be somewhat worse — treat the EVs as optimistic bounds, and the RELATIVE findings (deeper is better, below-cost is better, confirmation is better) as the durable conclusions.

No fees or slippage modeled: on thin books, entering and exiting costs real percentage points. The +24.6% EV has headroom for it; the +14% variant less so.

History window: 3 weeks of hourly / 12 weeks of 4-hour data per token, one market regime (late-Aug 2026, mixed-to-bullish). Rerun this study quarterly — the pipeline is one command now.

Fills assume patience: limit fills at exact rung prices, no front-running of your own size. At \$400-unit scale on \$1M+ tokens this is realistic; at 10× that it starts not to be.

Method: pump.fun swap-api OHLC candles · tops = first $\geq$ \$1M running-max high confirmed by a $\geq$ 30% drawdown · bounce odds from forward windows after simulated entries · ladder backtested bar-by-bar on 1h candles with 1-day cooldown between campaigns · VWAP = volume-weighted average close over full candle history. Data & code: research.db, research-analyze.mjs in tradewatch2.