

Leverage Tokens Without Decay

The short paper

Xatarrer

Decayless

October 2026

Status. This is the short paper, one of three. The full paper has the mathematics, with each result tagged as proved, checked numerically by a script, or still open, and the ELI5 paper explains the design from scratch. The result numbers here come from the same verification scripts and the same model as the full paper, and we typed in no result by hand.

1 Summary

If a leverage token is a claim on a shared reserve of collateral instead of a portfolio that gets rebalanced, it has no decay, no funding and no liquidation, as long as its pool has headroom.

This paper introduces the Decayless leverage token, a token that gives you constant leverage on the asset you chose, without decay, for as long as you hold it and its pool has headroom. **No decay. No funding. No liquidations.** One fee when you open, one when you close, nothing while you hold. Full leverage while the pool has headroom; if price runs past the pool's headroom, your leverage compresses toward 1x, and it never liquidates.

That last sentence names the principal trade-off of the design, and most of what follows is about it. The other limits, among them the price feed and the integer arithmetic, are listed in Section 10. You probably know leveraged tokens already, and their best-known flaw, the bleed that eats a position in a choppy market. The Decayless token is built to remove that bleed.

We removed it at a price, a limit on how far the price can run before the token stops being a full 2x. We call that limit headroom. The app shows it at all times, and it grows when liquidity providers deposit.

Crossing it is gentle. The token keeps gaining, at a shrinking multiple, and returns to full leverage when the price comes back.

We use pictures and worked examples throughout. Section 3 shows how the token works, and Section 4 how far the price can run and what happens past that point. Section 5 puts four leverages on one pool, and Section 6 follows them when several run out of headroom together. Then come the fees (Section 7), the other side of the trade as a liquidity provider sees it (Section 8), why trading at a fixed price cannot extract value from the reserve (Section 9), and the limits of our claims (Section 10). A glossary closes the paper (Section 11).

How to read it. If you trade, read Sections 2 to 4, 6, 7 and 10. If you provide liquidity, add Sections 5 and 8. If you review the design, read Sections 9 and 10 and then the full paper, which tags each result as proved, checked numerically, or still open, and traces the numbers here to the script that produced them.

2 Why do leveraged tokens bleed?

Suppose you want twice the move of ETH and you do not want a margin account. The product you probably reach for is a leveraged token. Exchanges and index protocols issue them, each

with a target leverage, or a target range, on the asset it tracks. To stay near its target, the token rebalances, daily or after a large move.

Where the target is a multiple of the daily return, the promise holds over a single day. The trouble is the rebalance itself. In a market that chops sideways, the token buys more after each rise and sells after each fall. That is a systematic loss, and it can leave the token a month lower than it started while the asset it tracks ends higher.

This is what traders call decay, or bleed. The rebalancing causes it, and leverage on its own does not.

Stock markets have the same product in leveraged exchange-traded funds, and the effect has a public record. In 2009 the SEC and FINRA published an investor alert on these funds with an example that ran from December 2008 to April 2009 ([U.S. Securities and Exchange Commission and Financial Industry Regulatory Authority, 2009](#)). Over those months an index gained 2%, while a fund seeking twice its daily return fell 6%. In the same period a different index gained about 8%, and a fund seeking three times its daily return fell 53%.

Crypto exchanges and protocols built the same product with different rules for when to rebalance. The leveraged tokens of FTX rebalanced once a day, and also during the day once leverage had risen 33% above its target ([FTX, 2020](#)). Binance gave its leveraged tokens a target range rather than a single target, between 1.5x and 3x ([Binance, 2020](#)). On chain, Index Coop launched ETH2x-FLI with a 2x target ([Index Coop, 2021](#)). Among its initial parameters, a rebalance never leaves leverage below 1.7x or above 2.3x. Between rebalances its leverage drifts with the price, which makes that range the place a rebalance returns it to rather than a hard band.

A perpetual future takes the other route. It does not rebalance. It has funding instead, a periodic payment between the two sides of the contract, which a holder can pay or receive and which longs usually pay when demand to be long is high. It also has a liquidation price. If the price dips below that level, even for a minute, the position is closed at a loss, even if the price recovers the next day.

So each existing way to hold leverage without a margin account has a cost. The token bleeds in a choppy market, and the perpetual can charge you funding for every period you hold it.

We take a third way. Suppose the token is a claim on collateral that already sits in a reserve. The collateral is not borrowed and no target has to be held, and the token has no reason to rebalance. Figure 1 shows what that means on one illustrative price path. It runs for 60 days of coin-flip moves of up to $\pm 10\%$ a day. At one point ETH is 29% above the start, at another 27% below it, and it ends +20% higher.

Suppose you put \$1,000 into each product on day zero. Holding ETH turns it into \$1200. A 2x token that rebalances daily turns it into \$1185, less than simply holding ETH, because the chop has eaten what the leverage should have added and more besides.

The Decayless 2x token turns it into \$1440, which is the starting amount times the square of ETH's price ratio (its end price over its start price). It would have reached the same amount on any other path that ended at the same price, on one condition, which Section 4 explains: the market must stay inside its headroom at every trade along the way, as it did on this path.

3 A reserve with two seats

Take ETH priced in dollars. The reserve holds ETH alone. It is not borrowed or posted in a position on an exchange, and no debt priced by an oracle stands against it. Two kinds of claim sit on the reserve (Figure 2). Traders hold *leverage tokens*, and liquidity providers (LPs) hold *LP tokens*, which together make up the *pool*. A token is worth the share of the ETH in the reserve assigned to it, at the current price.

Traders open and close tokens against the reserve, and they pay a fee each time. LPs deposit into the pool and withdraw from it. Half of each fee is added to the pool as ETH, and the other half leaves for the protocol's auctions, which are described elsewhere.




Figure 1: One illustrative choppy path, with \$1,000 in each product on day zero. ETH ends +20% higher. The rebalancing 2x token ends at \$1185 and the Decayless 2x token at \$1440, which is the starting amount times the square of ETH’s end price over its start price, with nothing subtracted for the route. The path is a seeded coin-flip sequence chosen for illustration. It is not a market forecast. As long as the market stays inside its headroom throughout, the Decayless token ends at the same value on any path that ends at the same price (Section 4).

You do not have to think in ETH at all. You can pay in ETH or in dollars when you enter, and take ETH or dollars when you leave, because the protocol swaps through the market’s own price pool inside the same transaction. The unit you think in is your choice, and we show values in both.

A swap has its own price impact, which the trade details show before you confirm. Between entering and leaving, though, your position is a claim on the ETH in the reserve, whichever asset you used at the door.

The split between the seats moves with the price. As ETH rises, a rule assigns more of the ETH in the reserve to the trader seat and less to the pool, and as ETH falls, it moves ETH back.

For small moves, a 2x token here behaves like the familiar 2x. Suppose ETH rises 10%. The 2x token rises 21%, and if ETH falls 10% instead, the token falls 19%.

But large moves compound, in both directions. If ETH multiplies by ten and the pool has headroom all the way, the token multiplies by 100, and if ETH halves, the token keeps 25% of its value. Table 1 gives the same picture for all four leverages.

Behind both is the rule. How fast the ETH moves is the one choice in the design, and we set it so that, while the pool has headroom, the dollar value of a 2x token follows the *square* of the ETH price. A 3x token follows the cube, and a 1.5x token the price to the power 1.5. No trade happens to make this so, and no target is maintained. The split is a fixed function of the price: once you know the price, you know the split.

For the curious

Write p for the price of ETH in dollars. A token with leverage ℓ has a dollar value proportional to p^ℓ . In ETH terms that is $p^{\ell-1}$, so a 2x token’s share of the reserve grows in proportion to the price and a 3x token’s in proportion to the price squared. The full paper writes the exponent as $q = \ell - 1$ and proves two things, both while the market is inside its headroom. The rule delivers exactly this payoff. And the return of the token between any two moments depends on the two prices alone, and neither on the path nor on the trades of anyone else.




Figure 2: The two seats on one reserve. Nothing is borrowed and nothing is held anywhere else. Each seat is worth its share of the ETH at the current price.

What would it take to wipe a leverage token out? ETH would have to go to zero. That is why there is no liquidation: the collateral is already in the reserve, and the token's claim on it shrinks smoothly instead of being closed. A 2026 proposal on ethresear.ch to build synthetic assets on fully collateralised option pairs rather than debt makes the same point (Buterin, 2026): two claims that always add back to the collateral leave nothing to liquidate. Our reserve is a perpetual form of that split, with a fee to open and one to close in place of an option's premium.

Table 1: What a move in ETH does to each leverage token at full leverage, that is, inside the token's headroom. Each entry is the token's move in percent. Small moves look like the plain multiple. Large moves compound, so the token gains more than the plain multiple on the way up and loses less than it on the way down.

ETH move	1.25x token (%)	1.5x token (%)	2x token (%)	3x token (%)
−50%	−58.0	−64.6	−75.0	−87.5
−20%	−24.3	−28.4	−36.0	−48.8
−10%	−12.3	−14.6	−19.0	−27.1
+10%	+12.7	+15.4	+21.0	+33.1
+20%	+25.6	+31.5	+44.0	+72.8
+50%	+66.0	+83.7	+125.0	+237.5
+100%	+137.8	+182.8	+300.0	+700.0

Where does the ETH a rising token gains come from? Figure 3 shows the split for one 2x market as the price moves, in a simplified example with a single token, which the next section also uses. At the entry price the tokens own a quarter of the reserve and the pool owns three quarters. As ETH rises, the pool hands ETH to the trader seat. As ETH falls, the trader seat shrinks toward zero and the pool ends up holding the reserve.

So the pool is the counterparty of all the leverage tokens, and the LPs are paid fees for taking that side (Section 8).

But the pool cannot hand over more ETH than it has. The rule that makes the token a 2x works up to a price and no further, the price where the pool has given all it can at that rate. The next section is about that price and what happens past it.




Figure 3: Who owns the reserve's ETH as the price moves, for one 2x market in which the tokens own a quarter of the reserve at the entry price. When the price rises, ETH flows from the pool to the token holders. When it falls, the ETH flows back and the pool ends up with almost everything. The dashed line is the headroom limit, +100% in this example. Past it, the pool has handed over all it can at the 2x rate (Section 4).

What this means for a holder

Your token is a claim on ETH that is already in the reserve. There is no rebalancing to bleed, no borrowing to pay funding on and no debt to call, which is why the token has no decay, no funding and no liquidation. Between the moment you open and the moment you close, your position's dollar value is multiplied by ETH's price ratio (the closing price over the opening price) raised to your leverage, less two fees, whatever route the price takes in between, as long as the market has headroom at every trade along the way.

4 How far can the price run?

Go back to Figure 3. The pool starts with three quarters of the reserve and the tokens with a quarter, and the rule moves ETH to the tokens as the price rises. At some price the pool has moved all it can move at the 2x rate. In that example it sits at +100%, where ETH has doubled.

We call that price the *headroom limit*, and the distance to it, +100%, the *headroom* of the market. Up to the limit the token is a full 2x.

Why does the limit sit where it does? The full paper proves that it is the price at which the pool's dollar value stops rising. Below the limit, that value grows with the price, and at the limit it peaks. Past it, the pool is pegged at that peak dollar value, and any further dollar of gain goes to the token holders.

Past the limit, the token keeps all the ETH it holds and keeps gaining as ETH rises, because with the dollar value of the pool pegged, the token side receives all further gains. But the gains no longer double the move.

Instead, the multiple the token delivers slides down from 2 toward 1 as the price runs further (Figure 4), and we call it the *effective leverage*. It never reaches 1 at any finite price, and it never falls below it. No ETH is sold, and no position is force-closed or liquidated.

Table 2 puts numbers on it for \$1,000 opened at the example's entry price. Its last column counts the same position in ETH. Inside the headroom the token's ETH balance grows in proportion

to the price, which means that if you count in ETH, you are gaining ETH on the way up as well. Past the limit the balance keeps growing, more slowly.




Figure 4: Effective leverage of the 2x token of Figure 3. It is exactly 2 up to the headroom limit. Beyond the limit it slides toward 1, approaching a plain ETH position without ever becoming one. The curve has no jump at the limit, so crossing it changes nothing abruptly.

Table 2: \$1,000 in the 2x token of Figure 3, opened at the entry price with +100% of headroom, before fees. Up to the limit the token is worth exactly what a full 2x pays. Past it, the token keeps gaining at the effective leverage shown, and stays below what a full 2x would have paid. The last column counts the same position in ETH, relative to the ETH it held at entry.

ETH move	Token value (\$)	Gain	Effective leverage	Full 2x (\$)	ETH held (entry = 1)
-40%	360	-64%	2.00	360	0.60
-20%	640	-36%	2.00	640	0.80
+25%	1563	+56%	2.00	1563	1.25
+50%	2250	+125%	2.00	2250	1.50
+100%	4000	+300%	2.00	4000	2.00
+200%	8000	+700%	1.50	9000	2.67
+400%	16000	+1500%	1.25	25000	3.20

Crossing the limit is also reversible. The split of Figure 3 is a fixed function of the price. The rule sells no ETH on the way up and has no ETH to buy back on the way down, and that leaves no ratchet.

Suppose ETH runs past the limit and then comes back inside it, and no one trades in the market while ETH is past the limit. Once ETH is back, the token is exactly where a full 2x would be at that price, as if the excursion had never happened. A trade made past the limit restarts the token from its compressed value and changes what it is worth when ETH comes back. Figure 5 puts the whole story on one line.

Against a rebalancing product the trade-off looks like this. Figure 6 compares the Decayless 2x token with a rebalancing 2x token at a range of end prices, for two levels of choppiness along the way. Inside the headroom the Decayless token pays the full square of the price ratio, the end price over the entry price. The rebalancing token pays less, and the gap grows with the choppiness.

But past the limit the Decayless token gives up part of the square. That is the price of not rebalancing, and you pay it on the way past the limit and nowhere else. The route itself adds no cost. We accepted that cost on purpose, because it depends on how far the price runs and not on the route it took.

How do you know how much headroom there is? The app shows each market's headroom next to the trade, as a dot, a word ("Full 2x" or "Compressing") and a number. The number is how




Figure 5: The two states of a single market. Crossing the limit is smooth in either direction, and reversible as long as no one trades in the market while ETH is past the limit.




Figure 6: Value at the end price, per dollar entered, for the 2x market of Figure 3. The Decayless token follows the square of the price up to the limit and stays above the spot line beyond it. The two rebalancing curves show a continuously rebalanced 2x product reaching the same end price by a calm route and by a choppy one. Their gap to the square is the bleed.

much further ETH can rise before the token you are about to hold starts compressing, and it includes the size you are entering.

Headroom grows when LPs deposit and shrinks when traders open more leverage against the same pool. It concerns big gains rather than losses, since a falling price never takes a token past its limit.

What this means for a holder

Full leverage while the pool has headroom. If price runs past the pool's headroom, your leverage compresses; it never liquidates. Compressing means you keep gaining as ETH rises, at a multiple that slides from 2 toward 1. You're back at a full 2x the moment ETH is back inside the limit. Check the gauge before you open. Its number is how far ETH can rise before your position stops being a full 2x.

5 The four-token menu

So far we have used one market, with one pool and one 2x token. The design puts four leverage tokens on the same pool of a pair: 1.25x, 1.5x, 2x and 3x. Each keeps its own rule. The 1.25x token's dollar value follows the price to the power 1.25, and the 3x token's follows the price cubed. The pool is what is left of the reserve after the tokens have been given their shares.

The full paper proves that four tokens keep what one had. The payoff of each token is exact, the pool has one well-defined headroom limit, and checking whether the reserve is inside it takes a single comparison.

We chose one pool per pair because we ran the alternative. Our predecessor, Sir Trading on Ethereum mainnet, ran one pool per market, that is, per pair and leverage. A pair with several leverages split its liquidity providers across several pools. With one pool per pair, capacity that the 3x token is not using at the moment is available to the 1.25x token, and one deposit backs the whole menu.

You might worry that pooling lets the gentle tokens ride free on the aggressive ones. It does not, because the demand a token makes on the pool grows with its leverage. To keep a 3x token at full leverage, for example, the pool has to release ETH 8 times as fast as for a 1.25x token holding the same amount of ETH.

The full paper calls this demand the capacity of a token and adds it up across tokens. It shows that the pool reaches its limit at the price where the total capacity of the tokens equals the ETH the pool has left. Under the fee schedule of Section 7, the 3x token also pays 8 times the fee, and per unit of pool capacity the four tokens pay the same.

The one-token example of Sections 3 and 4 was a simplification, and from here on we replace it with a reserve that carries all four tokens, the running example of the rest of the paper. Today the four tokens hold a tenth of the ETH in the reserve apiece, and the pool holds the remaining six tenths. Figure 7 shows who owns the ETH as the price moves. With four tokens drawing on one pool, the pool reaches its limit at +23%.

What this means for a liquidity provider

You deposit once and back the whole menu at the same time, and you can't pick a token. Capacity that one token isn't using is available to another. Section 8 shows what your position is worth as the price moves.

6 The aggressive token steps aside first

Once the pool reaches its limit, it can no longer keep all four tokens at full leverage. As the price keeps rising the pool can still release some ETH, and the design has to decide which tokens get it.




Figure 7: Who owns the reserve's ETH as the price moves, in the running example (each token holds a tenth of the reserve today, the pool six tenths). When the price falls, the tokens shrink, the 3x fastest, and the pool ends up with the reserve. When it rises, the pool hands ETH to the tokens until it reaches its limit, the first dashed line. The other dashed lines mark the prices at which the 3x, 2x and 1.5x tokens, in turn, go flat under the rule of Section 6.

The full paper studies all admissible ways of dividing that ETH, and it proves two facts that hold whichever rule is chosen. First, the value of the pool does not depend on the rule, which leaves LPs indifferent to the choice. Second, the total ETH the pool can ever release past its limit is finite.

That second fact rules out the naive fix of giving the gentle tokens priority forever. Someone has to absorb the shortfall.

Suppose ETH rallies past the limit of the pool. Under the rule we adopted, the 3x token is the first to compress. The three gentler tokens keep their full leverage, and the 3x token absorbs the whole shortfall, its effective leverage sliding from 3 toward 1.

The 3x token compresses until the price where it would otherwise begin to lose ETH. There it is frozen at the ETH it holds, which makes it a plain ETH position, flat at 1x, and from there the 2x token takes over the shortfall, then the 1.5x. The 1.25x token comes last and never goes flat, compressing toward 1x as the single market of Section 4 does.

We call this queue, ordered by leverage, the cascade. For a holder, the full paper proves two things about it. It is smooth, with no jump in any value or in any effective leverage at any price. And it takes ETH from no token, which means that as the price rises, no token loses ETH.

The choice of this rule over the others rests on a third result. Among all admissible rules, the cascade gives the gentlest token the most ETH at any price, then the next gentlest, and so on.

So instead of one shared limit for the reserve, each token has a limit of its own. Table 3 and Figure 8 give the four limits for the running example.

The 3x token compresses from +23%, the price at which the pool itself reaches its limit. The 2x token keeps its full leverage until ETH has risen +105%, the 1.5x until +343%, and the 1.25x until +1216%. Figure 9 shows the effective leverage of all four tokens along the way.

Seen from the token you hold, the path has three stages. Up to its limit, you get full leverage. Past it, your leverage glides toward 1 while your token is the one absorbing the shortfall. Then, unless you hold the 1.25x, you have a plain ETH position.

Table 3: Each token’s headroom in the running example: how far ETH can rise from today before the token stops delivering its full leverage.

Token	ETH can rise by	ETH price then (today = 1)
3x	+23%	1.232
2x	+105%	2.047
1.5x	+343%	4.429
1.25x	+1216%	13.161




Figure 8: The same headroom as a ladder: each bar runs from today’s price to the price at which that token starts compressing, on a logarithmic price axis. The gentler the token, the further the price can run.




Figure 9: Effective leverage of the four tokens in the running example. Each holds its promised leverage up to its own limit, slides to 1 while it is the token absorbing the shortfall, and stays at 1 once flat. The 1.25x token has no flat stage. No curve jumps anywhere.

And the path is reversible. When ETH comes back inside your limit you are at full leverage again, with no ratchet and no rebalancing. You are also back at the value full leverage would have given you, as long as no one trades in the pair while your token is past its limit, because the split is then a fixed function of the price. A trade made while your token is compressing or flat restarts your token from its compressed value and changes what it will be worth when ETH comes back.

Between opening and closing, the dollar value of your position is multiplied by ETH's price ratio (the closing price over the opening price) raised to your leverage, less two fees, as long as *your* token is inside its own limit at every trade in between. It does not matter whether the pool as a whole, or a more aggressive token, was past its limit at the time.

Whose trades can move your headroom? The 3x token's limit is the pool's limit. The 2x token's limit depends on the pool's peak dollar value and on the sizes of the 2x, 1.5x and 1.25x tokens, and on no other quantity. How much 3x has been opened enters through that peak value and no other way.

The full paper turns this into a precise statement. Suppose someone opens or closes any amount of a token while the pool is past its limit. Before fees, that leaves the limit of every gentler token where it was. With fees, it moves those limits slightly *up*, because the half of the fee that goes to the pool raises its peak, and with it the limit of each gentler token still at full leverage.

For example, take a 3x position worth a tenth of the reserve, opened while the pool is past its limit. It moves the 2x token's limit from a price of 2.047 to 2.069 (today = 1), and that shift is the fee alone. Without the fee, the move is zero to machine precision.

Now open the same position while the pool still has headroom. This time it does move the gentle limits, though not by much: the 1.25x limit moves by -2.1% , while the 3x token's own limit collapses to $+4.7\%$.

Compare equal sharing, the rule we started from, in which all four tokens are compressed by the same fraction and share one limit. Under equal sharing, that position would drag the headroom of all four tokens down to $+4.7\%$. Figure 10 sweeps the size of the position.




Figure 10: How far a 3x position moves the 1.25x token's headroom in the running example, against the size of the position (fees included). Under the cascade a position opened while the pool is past its limit leaves the 1.25x limit unchanged or slightly higher; one opened while the pool has headroom lowers it by a few percent. Under equal sharing the same position drags the single shared limit down by an order of magnitude more.

The cascade protects the gentle tokens from the aggressive ones, and from two other things it does not protect them. When LPs withdraw, the limits of all four tokens fall. When the gentle tokens themselves grow, their own limits fall.

That protection has a payer. The 3x token pays for the headroom of the gentle tokens by taking the first loss of capacity instead of a proportional share, and Table 4 measures the cost against equal sharing.

In a rally past the limit of the pool, the 3x token holds less ETH than it would under equal sharing, and the 1.25x token holds more. The percentages are the same in dollars, since both are valued at the same price. The gap converges, because the ETH the pool can release is finite.

So the cascade re-slices the same ETH among the tokens, and the value of the pool is the same under both rules. Section 7 explains how the fee schedule takes the cost into account.

Table 4: Cascade against equal sharing in the running example. Columns: the ETH held by the 3x token under the cascade relative to equal sharing; the same for the 1.25x token; the 3x token’s shortfall as a share of the reserve. The last row is the limit of a runaway rally.

ETH price (today = 1)	3x token (%)	1.25x token (%)	3x shortfall (% of the reserve)
1.5	−4.4	1.0	0.9
2.047	−19.8	5.1	5.6
3	−35.3	12.5	12.4
5	−44.8	25.3	18.4
∞	−55.0	—	27.7

What this means for a holder

Your token has its own headroom, and the gauge shows it. Up to your limit you get the full leverage you bought. Past it you keep gaining as ETH rises, at a shrinking multiple, until (for any token but the 1.25x) your position is a plain ETH position. You’re never liquidated, and you’re back at full leverage as soon as ETH is back inside your limit. The more aggressive tokens in the pool compress and go flat before yours does, and what anyone opens in them while the pool is stretched does not move your limit down.

7 Fees

You pay two fees, and the app puts the whole schedule in one line. One fee when you open, one when you close, nothing while you hold. Table 5 is the launch schedule. The base fee is 10%, the fee of the 2x token, and a token pays the base fee times its leverage minus one. The 1.25x token pays a quarter of the base, and the 3x token pays double.

The fee is quoted on the amount that goes into your position, which makes it a little lower as a share of what you actually pay. Of a dollar you pay to open a 2x position, 9.09% is the fee and the rest goes into the position.

Closing charges the same quoted fee on the amount you receive. Of what leaves the position when you close a 2x, 9.09% is the fee, and the rest is paid to you.

Half of the fee is added to the pool as ETH at the moment it is paid. The other half leaves the reserve for the protocol’s auctions, which are specified elsewhere.

Between opening and closing you pay no funding, no management fee and no rebalancing cost. You may pay or take dollars instead of ETH. If you do, the swap through the price pool adds its own price impact on top, and the trade details show it.

Why should the fee grow with leverage? The demand a token makes on the pool grows with its leverage minus one (Section 5). Consider two positions of the same size. A 3x position ties up 8 times the pool capacity of a 1.25x position, and under this schedule it pays 8 times the fee. The pool is paid the same for a unit of capacity it commits, whichever token consumes it, and in that sense the schedule creates no subsidy between tokens.

You might ask for another kind of fairness. For the same ratio of pool to the capacity in use, a gentle token has more headroom than an aggressive one. In the full paper’s example, when the

Table 5: Launch fee schedule. The fee is charged on the way in, on the amount that goes into the position, and on the way out, on the amount you receive; the third column is the same fee as a share of what you pay in.

Token	Fee to open or close (%)	Share of what you pay (%)
1.25x	2.50	2.44
1.5x	5.00	4.76
2x	10.00	9.09
3x	20.00	16.67

pool holds twice the capacity in use, the limit of the 1.25x token is +107% away and that of the 3x token is +29% away. That is a property of the payoff rather than of the fee, and it is one reason the menu stops at 3x.

We considered a steeper schedule during the design, one charging the 3x token three times the base fee instead of twice. The argument for it was that the exposure of the pool to a token grows faster than its leverage.

We did not adopt it because of the cascade. Under the cascade the 3x token already pays for its aggression by taking the first loss of headroom (Table 4), and charging it the steepest fee on top would price the same thing twice.

Relative to the steeper schedule, the launch schedule lowers the 3x token’s fee by 6.4 percentage points of what you pay. That is of the same order as the ETH the 3x token forgoes in a moderate rally under the cascade.

8 Providing liquidity

When you deposit ETH into a pair’s pool, you get LP tokens, a share of the pool seat of Figure 2. You back all the leverage tokens of that pair at once, and you cannot choose one. Your share of the reserve is whatever is left after the tokens have been given theirs.

So your position is the other side of a leveraged bet, and counted in ETH, its balance moves with the price in the opposite direction to the balances of the tokens. As ETH rises, the pool hands ETH to the tokens. As ETH falls, it takes ETH back.

You earn half of each open and close fee, added to the pool as ETH at the moment it is paid. Your LP token’s claim on the reserve grows with each trade. This is income proportional to trading in the pool, paid in the collateral you already hold, and it is neither fixed nor guaranteed: a quiet pool earns no fees, and a busy one earns in proportion to what traders open and close.

The other half of each fee feeds the protocol’s auctions. Rewards paid from those auctions in the protocol’s own token are specified separately, and we do not count them here.

What do you risk, and against what? Your position is a blend of the pair’s two assets, and how it looks depends on which of them you would otherwise have held. Figure 11 and Table 6 show an LP position in the running example against both, counted in dollars and in ETH.

Counted in dollars, the position rises with ETH up to the limit of the pool and is then pegged flat. A rally past the limit adds no further dollars, because all further gain goes to the tokens, but it does not take any away either. In a crash the position falls with ETH, as holding ETH would and holding dollars would not.

Counted in ETH, that same peg means your balance shrinks as ETH rises past the limit, in proportion to one over the price. Suppose ETH climbs to twice the limit price. You then hold half the ETH you held at the limit. Relative to having held ETH, a large enough rally can cost you most of your ETH, and the further the price runs, the larger that share. Counted in dollars against the same benchmark, the shortfall grows without limit as the price rises.

A crash runs the other way. The tokens shrink first, the 3x fastest, the pool ends up holding the reserve, and counted in ETH you gain.

Between those cases, inside the headroom, the position moves with the price at an effective leverage between 0 and 1. It behaves like holding ETH when the pool is lightly used, and like holding dollars when it is fully used.

So at any price the position sits between the two benchmarks, neither better than both nor worse than both, and the fees are what you are paid on top. You can deposit and withdraw in either asset, which lets you choose the benchmark that is yours.




Figure 11: An LP position in the running example against holding ETH and against holding dollars, all equal to 1 today. The dashed line is the pool's limit. Left, in dollars: the position rises with the price up to the limit and is pegged flat past it, which is where holding dollars is; in a crash it falls with ETH. Right, in ETH: it grows in a crash, when the pool ends up holding the reserve, and shrinks in a rally past the limit, in proportion to one over the price. In both panels, at every price, the position lies between the two benchmarks.

Table 6: The same LP position at a range of ETH moves, relative to its value today, counted in dollars and in ETH, next to what holding either asset would be worth (all 1 today). The pool's limit is at +23%; past it the dollar value is pegged.

ETH move	LP, in dollars	holding ETH, in dollars	LP, in ETH	holding dollars, in ETH
-50%	0.642	0.500	1.284	2.000
-20%	0.896	0.800	1.120	1.250
+20%	1.044	1.200	0.870	0.833
+50%	1.044	1.500	0.696	0.667
+100%	1.044	2.000	0.522	0.500
+200%	1.044	3.000	0.348	0.333
+400%	1.044	5.000	0.209	0.200

What this means for a liquidity provider

Your balance shifts with price like an LP position; backing higher leverage means bigger swings. This is not fixed yield. You're paid fees, in proportion to trading, for taking the other side of all the leverage tokens in the pair. Counted in ETH, a large enough rally can leave you with much less ETH than you deposited, and the further the price runs, the less you hold. Counted in dollars, a rally never takes value away, but a crash does. At any price your position sits between holding ETH and holding dollars, and you can enter and leave in either. No position on either side is liquidated.

9 Safety, in one page

Start with trades at a single price. Suppose you open and close leverage tokens at one price, in any tokens and in any order. When you stop, you have paid the fees and no more, even if you crossed a headroom limit partway through. The full paper proves this and confirms it on 2400 random sequences, including ones that straddle limits.

Real trades do not all see one price, though. Instead of reading ETH's price once, the protocol has two readings in each block, and both come from its oracle on Uniswap v3 ([Adams et al., 2021](#)). One is an average over the last half hour, with a cap on how fast it can move. The other is the price the Uniswap pool showed at the start of the current second, which is fixed before anything in that second happens.

A trade fills at whichever reading is worse for the trader: you buy a token at the higher reading and sell it at the lower one. A pool deposit is counted at the lower reading, and a withdrawal at the higher.

We chose the worse of the two because a reading can lag the market. Before this rule, a trader who knew the market had moved could trade against the stale average. Now that trader gets no better price from the pool than the slow average would give, and usually a worse one. The gap between the two readings becomes a cost the trader pays, and it goes to the people already in the pool.

The full paper proves that a trader can't profit from any sequence of their own trades at the two prices, in any state of the pool, as long as fees are left aside. With fees counted, the proof covers leverage-token trades alone. The one case it checks numerically but doesn't prove is a trader who mixes in pool deposits while paying fees. The paper confirms all of it on 7500 random sequences.

You pay for the rule on both legs of a position, since you open at the higher reading and close at the lower one. What that costs depends on the gap and on the leverage of the token. At a gap of 0.5% on both legs, for example, a round trip in the 2x token pays 0.50% extra on top of the fees, and one in the 3x pays 0.99%. When the two readings agree, there's no extra cost.

The oracle comes with defences of its own, a liquidity measure that a brief withdrawal collapses and a cap on how fast the average can move, and we take those as given here.

Someone who pushes a reading off the market never gets a better price than the average, though a push can hold the average back by one small step. A push can also make the other trades in that block worse, which is what the minimum you set on a trade protects you from. And when the market moves and both readings miss it, for as long as the Uniswap pool hasn't traded or after a market halt, the people already in the pool take the loss, as they would with any price feed.

A manipulated reading moves all the positions in the pair at once. That is a reason to prefer deep pairs, and it isn't a new attack.

Can a trade by someone else change what your tokens are worth? At the price that trade uses, it cannot. A trade in any token leaves the value of all the other tokens where it was at the price that trade used. A leverage-token trade can add to the value of the pool and cannot subtract from it, and what it adds is half its fee.

But another trade can still change your headroom, and with it your future, in the ways Section 6 works out. A trade made while the pool is past its limit never lowers the limit of a gentler token. A trade made inside the headroom moves all the limits a little, and an LP withdrawal lowers all of them. These moves leave the value of your tokens untouched at the moment they happen.

The reserve also keeps no memory. Two reserves with the same amounts at the same price are the same reserve. After a trade, the reserve is described afresh by what the two seats hold now, and the protocol stores no record of how it got there. The full paper proves that this description always exists and is unique, which leaves the protocol no choice to make. A trade of zero size leaves the reserve as it found it.

The review of the design found one profitable trick, on the LP side. Fees go into the pool the moment they're paid. Suppose an LP deposits just before a large fee and withdraws just after. That LP would take a share of the fee without having backed anything. Our repair, specified in the protocol's tokenomics design, is a one-block quarantine: LP tokens from a fresh deposit earn no share of the fees paid in the block of the deposit, and they can't be moved or closed until the next block. That removes the same-block, risk-free version of the trick. The exact, automatic, constant-time accounting that applies the quarantine is still an open item of that design.

Could an LP try it across two blocks instead? A deposit that stays a block bears a block of price risk, plus the uncertainty of whether the fee is paid at all. A single fee large enough to be worth that risk is still possible, and the interface deals with it by splitting large trades. The protocol itself doesn't.

We prove the results above for exact arithmetic, with one exception. A trader who mixes pool deposits with leverage-token trades at the two prices while paying fees is covered by the numerical checks rather than by a proof. The integer arithmetic has its own obligations (rounding rules, overflow bounds, how a price equal to a limit is decided). The full paper lists them, and we make no claim about them here.

10 Limits of the claims

We state the main claim in the form the app uses. Full leverage while the pool has headroom; past it, leverage compresses toward 1x and never liquidates. That sentence is the claim in its honest form, liquidity providers get a mirror image of it, and the list below sets out where both stop.

- **“No decay” is a claim about headroom.** Inside its headroom, a token's dollar value is multiplied by ETH's price ratio (the closing price over the opening price) raised to its leverage, whatever the route and whatever anyone else does. If the price runs past the limit of a token, the token keeps gaining at a shrinking multiple, and all tokens but the gentlest eventually hold a fixed amount of ETH. They recover their full leverage if the price comes back. The claim never appears without its condition, and the app shows the headroom of each token next to the trade.
- **The pool side can lose most of its ETH relative to holding ETH.** Liquidity providers back all the tokens of a pair and can't choose one. Past the limit of the pool their dollar value is pegged, and counted in ETH their balance shrinks as the price rises, and a large enough rally can take most of it. Counted in dollars against holding ETH, the shortfall grows without limit as the price rises. In a crash they end up holding the reserve, and counted in dollars they lose with the price. At any price the position sits between holding ETH and holding dollars. The fees compensate for that exposure, in proportion to trading, and they are fee income that is neither fixed nor guaranteed.
- **The proofs are in exact arithmetic.** The integer arithmetic has its own list of obligations, which the full paper sets out. No claim here is a claim about the code of any particular deployment.
- **The quarantine still lacks its accounting.** The one-block quarantine of Section 9, which repairs the one trick found there, is specified in the protocol's tokenomics design. Its exact, automatic, constant-time accounting is still open.
- **The protocol token and the auctions are elsewhere.** Half of each fee leaves the reserve for the protocol's auctions. The protocol's own documents specify how that half is used and what rewards liquidity providers receive from it, and we don't count any of it anywhere in these pages.
- **Shorts are not in the launch.** The full paper proves a mirrored construction in which short tokens share a pool with long ones, but it isn't part of the launch design. Mechanically, a short on ETH is a long on the quote asset.
- **The price feed is trusted.** Everything stated here is about the reserve at the prices the feed reports. Both readings of the feed can be stale at once, for as long as the Uniswap pool has not

traded or after a market halt. Filling you at the worse of the two narrows that window but doesn't remove it. The feed's own defences are the protocol's, and we don't re-derive them.

- **One knob is still open.** We have not decided whether the slope of the fee schedule ships as a fixed value or as a bounded parameter, though the launch behaviour is the same either way. The menu stops at 3x. What about a more aggressive token later? Under the cascade, adding one would be a decision mostly about that token. The limits of the existing tokens would still move through the peak dollar value of the pool (Section 6), and liquidity providers would carry the exposure of a pool that may be past its limit from the day the token launches.

11 Glossary, and where to go next

Leverage token

A claim on part of a pair's reserve whose dollar value follows a fixed power of the price while it is inside its headroom. Named by its on-chain symbol, which gives the collateral, the quote asset and the multiple, for example WETH-USDC-2X.

Market

One asset at one leverage, for example ETH 2x, backed by the pair's pool.

Reserve

All the collateral held for one pair, split between the leverage tokens and the pool.

Pool

The liquidity providers' side of the reserve, and the counterparty to all of the pair's leverage tokens. It has its own headroom limit, which is the most aggressive token's limit.

LP token

A share of the pool.

Collateral, quote

The asset in the reserve (ETH in our examples) and the asset it is priced in. You can pay and be paid in either, and the protocol swaps inside the transaction, at the price pool's own price impact.

Headroom

How much further the price can rise before a token stops delivering its full leverage. The app shows it live, counting the trade you are entering, and it grows when liquidity providers deposit and shrinks when traders open more leverage.

Headroom limit

The price at which a token's headroom is used up and it starts compressing.

Effective leverage

The multiple a token is actually delivering right now: its stated leverage inside the headroom, sliding toward 1 past the limit.

Compressing, flat

A token past its limit is compressing, which means it is still gaining but at a shrinking multiple. A token other than the gentlest is flat once it has absorbed all it can and holds a fixed amount of ETH.

Fee

Paid once when you open, on the amount that goes into the position, and once when you close, on the amount you receive, with nothing while you hold. Half goes to the pool and half to the protocol.

Where to go next. The full paper gives the mathematics behind this one. Each of its results is tagged as proved, checked numerically by a script, or still open. It works in exact arithmetic and lists what the integer arithmetic still has to establish. Its verification scripts regenerate all the numbers in the three papers, and all the figures and tables, in a few seconds on a stock Node.js. If you disagree with a number, you can change the running example and see what moves. The

ELI5 paper covers the same design from scratch.

Acknowledgments

The design, its proofs and the three papers were developed through an adversarial review process, with large language models as counterparties: Claude Fable 5.1 by Anthropic and Codex by OpenAI. One side proposed and proved, and the other tried to refute. Every claim the results rest on was checked numerically by both. The verification scripts are part of the papers and, like any errors that remain, are the author’s responsibility.

References

- Hayden Adams, Noah Zinsmeister, Moody Salem, River Keefer, and Dan Robinson. Uniswap v3 core. Whitepaper, March 2021. URL <https://uniswap.org/whitepaper-v3.pdf>.
- Binance. About Binance leveraged tokens. Binance Support, 2020. URL <https://binance.zendesk.com/hc/en-us/articles/360042957472-About-Binance-Leveraged-Tokens>. Page dated 2020-05-11. Read from the Wayback Machine capture of 2020-05-29, accessed 2026-10-13.
- Vitalik Buterin. Building index-tracking assets on top of options instead of debt. ethresear.ch, topic 25036, with the discussion that followed, June 2026. URL <https://ethresear.ch/t/building-index-tracking-assets-on-top-of-options-instead-of-debt/25036>. Accessed 2026-10-13.
- FTX. Leveraged token walkthrough. FTX Help Center, 2020. URL <https://help.ftx.com/hc/en-us/articles/360032509552-Leveraged-Token-Walkthrough-READ-THIS>. Page edited 2020-06-15. Read from the Wayback Machine capture of 2020-08-11, accessed 2026-10-13.
- Index Coop. Introducing the flexible leverage index (FLI), by Pulse Inc. Index Coop on Medium, 2021. URL <https://medium.com/indexcoop/introducing-the-flexible-leverage-index-fli-by-pulse-inc-a369bd422ef>. Published 2021-03-17. Read from the Wayback Machine capture of 2021-03-17, accessed 2026-10-13.
- U.S. Securities and Exchange Commission and Financial Industry Regulatory Authority. Leveraged and inverse ETFs: Specialized products with extra risks for buy-and-hold investors. Investor Alert, 2009. URL <https://www.sec.gov/investor/pubs/leveragedetfs-alert.htm>. Read from the Wayback Machine capture of 2009-08-20, accessed 2026-10-13.