How Do Stablecoins Stay Pegged?

Nobody sets the price — so what holds it? The two-tier market behind every peg, the different toolkits each design uses, and the anatomy of a break.

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

A stablecoin trades on the open market like any other token: its price is whatever buyers and sellers agree on at that moment. No exchange enforces a dollar, no regulator sets it, and the issuer cannot dictate it. Yet the price sits stubbornly at a dollar, day after day, across billions in daily turnover. The mechanism responsible is worth understanding in detail, because it is also the mechanism that fails — and knowing how a peg is held is the only way to judge whether a particular one will hold tomorrow. This report goes deeper than the overview in what is a stablecoin, covering the market structure behind the peg and the different tools each design uses to defend it.

AT A GLANCE

NOBODY SETS IT **The price is a market price**	HELD BY **Arbitrage, not enforcement**	TWO TIERS **Primary issuance, secondary trading**
FIAT-BACKED TOOL **Redemption at par**	CRYPTO-BACKED TOOLS **Liquidations, fees, swap modules**	BREAKS WHEN **Arbitrageurs stop believing**

01

The peg is a consequence, not a rule

Start by discarding the intuition that something holds the price at a dollar the way a central bank might defend a currency. Nothing does. The price of a stablecoin on any exchange is set by the orders in that market, exactly as it is for any other asset — the mechanics are those described in how order books work, with no special case for stablecoins.

What produces the peg is that deviations create profit opportunities, and people take them. If the token can be converted into a dollar somewhere else, then any price below a dollar is free money to whoever can perform that conversion, and their buying pushes the price back. The peg is the residue of that activity — an equilibrium, continuously re-established by self-interested traders. This framing explains everything that follows, including the failures: a peg does not break because someone stopped enforcing it. It breaks because the arbitrage stopped being profitable, stopped being possible, or stopped being believed.

02

The two-tier market

Here is the structural fact that most explanations omit, and it does most of the work. There are two separate markets in every stablecoin, and they operate under different rules.

The *primary market* is the issuer's own window: creating new tokens by delivering dollars, and redeeming tokens for dollars. This is where the peg is anchored, because it is where a token can be exchanged for its backing at face value. But access is typically restricted — verified institutional counterparties, minimum sizes that run to substantial sums, and processing on business days. The *secondary market* is where everybody else trades: exchanges, on-chain pools, brokers. Here the price floats freely, and here is where you and almost every retail holder actually transact.

The peg therefore depends on a small number of large participants bridging the two. When the secondary price dips, an authorised firm buys tokens cheaply there and redeems them at par in the primary market, pocketing the difference — and their buying is what lifts the secondary price. Notice the dependency this creates: your ability to sell at a dollar rests entirely on *someone else's* ability to redeem at a dollar. If that channel narrows — slower processing, higher minimums, a banking holiday, a paused redemption window — the connection weakens and the secondary price can drift regardless of how sound the reserves are.

03

The fiat-backed toolkit

For reserve-backed tokens the defence is straightforward in principle: hold assets that can be turned into cash quickly, and honour redemptions promptly. The quality of the peg is largely a function of the quality of that promise, which is why reserve composition matters

so much — short-dated government paper can be liquidated at scale in a day, while less liquid assets cannot, and a redemption queue that slows under pressure is precisely what breaks the arbitrage loop.

Two frictions deserve naming because they are invisible until they matter. The first is *timing*: traditional banking settles on business days, so a weekend crisis cannot be arbitraged away until banks reopen, which is why depegs frequently occur and persist across weekends. The second is *capacity*: the arbitrage only works to the extent that firms have capital available to deploy into it, and in a genuine panic those firms may be unwilling to take the other side at any price they can justify — the same withdrawal behaviour described in what is a market maker. The reserves can be entirely real and the peg can still slip, simply because the bridge between the two markets is closed for the weekend or nobody wants to cross it.

04

The crypto-backed toolkit

Over-collateralised systems have no company to redeem from, so they defend the peg with a different and considerably more active set of instruments. *Liquidations* are the first line: if the collateral behind issued tokens falls toward the danger threshold, positions are automatically sold to restore backing, which keeps the system solvent and the token credible.

Beyond that, these systems adjust *incentives*. Fees charged for issuing tokens can be raised or lowered to make creating supply more or less attractive, and rates paid for locking tokens away can be adjusted to change demand — a monetary policy in miniature, usually decided by governance. Many such systems also operate a direct swap facility holding reserve assets that can be exchanged for the stablecoin at a fixed rate, which is effectively a redemption window rebuilt on-chain — and it is often the most effective single tool they have, because it restores the primary-market anchor that a purely collateral-based design lacks. The vulnerability is different in kind from the fiat-backed case: the danger is not a slow bank but a fast market, in which collateral falls so quickly that liquidations cannot clear at viable prices, particularly when network congestion delays them.

> *Every peg is a bridge between two markets. The reserves are one bank of the river and the exchange price is the other. What matters is not how much sits on the far bank, but whether anyone can cross — how fast, at what size, and on what day of the week.*

05

The algorithmic toolkit, and why it inverts

Designs without meaningful collateral attempt the same arbitrage logic without anything to redeem against. In place of reserves, the protocol lets the stablecoin be exchanged for a fixed dollar's worth of a second, volatile token it issues. Below the peg, arbitrageurs buy the cheap stablecoin, swap it for a dollar of the companion token, and sell — removing stablecoin supply and pushing the price up.

The logic is identical to the redemption arbitrage above, with one substitution: the thing you receive is not a dollar but a claim on a token whose price floats. That substitution is the whole problem. In calm conditions the mechanism works and the peg holds. Under stress, the arbitrage itself creates selling pressure in the companion token, which drives its price down, which means more of it must be issued to cover each stablecoin, which drives it down further. The stabiliser becomes an amplifier. This is why the category has produced sudden, near-total collapses rather than gentle drifts, and why the distinction between redeeming for an *asset* and redeeming for a *claim on confidence* is the single most important thing to check in any peg design.

06

Anatomy of a break

Depegs follow a recognisable sequence. It begins with a *doubt* — a revelation about reserves, a failed counterparty, a collateral crash, or simply a rumour. Holders begin selling on the secondary market, and the price ticks below the peg. So far this is ordinary and self-correcting.

The second stage decides everything: do arbitrageurs step in? If the redemption route is open and credible, they buy the discount aggressively and the price snaps back within hours — the mechanism working as designed. If instead they doubt they will be paid, or cannot access redemption, or fear the discount will widen further before they can complete the round trip, they stand aside. Now the discount deepens with no counter-pressure, which is itself read as confirmation that something is wrong, prompting more selling. The final stage is either *restoration*, when the doubt is resolved and arbitrage resumes with force, or *collapse*, when the backing proves inadequate and the price finds a new level reflecting what the token is actually worth. The variable throughout is not the reserves. It is whether the people who could profit from the gap believe they will collect.

07

Why some recover and some do not

Comparing outcomes across the record, the recoveries share a feature: the backing was real and the interruption was to the *channel* rather than to the collateral. Where a token fell because reserves were temporarily inaccessible — trapped in a failed institution, or unreachable over a weekend — resolution restored the arbitrage and the price returned quickly, because the underlying claim had never actually been impaired.

The permanent failures share the opposite feature: the backing was insufficient, or existed only as confidence in a related asset. No amount of time or intervention restores a peg whose collateral has evaporated, because there is nothing for the arbitrage to converge toward. This gives a practical test for reading any depeg in progress. Ask what broke — the reserves, or the route to them? A channel problem is usually survivable and often a buying opportunity for those who can verify it. A backing problem is usually terminal, and the discount is the market pricing that in correctly rather than irrationally.

08

What to watch

Four things tell you how robust a peg is before you need to know. First, who can redeem, at what minimum size, and how quickly — because that defines the strength of the bridge. Second, what the reserves consist of, since liquidity under stress matters far more than headline value. Third, the depth of secondary markets, since a thin market lets modest selling move the price further than the underlying justifies, as our note on market liquidity describes. And fourth, the record — how the token behaved during previous stress, which is the only real evidence available.

Then keep the correct expectation. Small deviations are the mechanism breathing, not failing; a token trading a hair off a dollar simply means the arbitrage has not yet been worth someone's time. What matters is the size and persistence of a gap, and above all whether the redemption channel remains open. A peg is not a promise that a price will never move. It is a bet that when it does, somebody will find it profitable to move it back — and the whole art of judging a stablecoin is estimating how confident that somebody will be on the worst day.

“Let us hold fast the profession of our faith without wavering; for he is faithful that promised.”

HEBREWS 10:23

Methodology & Sources

This report explains peg mechanics across stablecoin designs; it contains no coin names, issuer names, reserve figures, deviation sizes, fee levels, or collateral ratios, all of which vary by issuer and change continuously — consult current issuer documentation and market data directly. The account of peg maintenance is mechanical and conventional: a stablecoin's market price is set by supply and demand on secondary venues, and proximity to par is sustained by arbitrage between those venues and a primary issuance-and-redemption channel typically restricted to verified counterparties subject to minimum sizes and business-day processing. The description of crypto-collateralised defences — automated liquidation of deteriorating collateral, adjustable issuance fees and savings rates set by governance, and direct swap facilities holding reserve assets exchangeable at fixed rates — reflects mechanisms in general use; specific implementations differ substantially between protocols. The reflexive failure mode attributed to designs redeeming into a volatile companion token rather than an external asset reflects the documented record of that category. The depeg sequence in section 06 and the channel-versus-backing distinction in section 07 are analytical framings offered as diagnostic tools, not predictions; no specific incident, issuer, or protocol is named. Nothing here is a recommendation regarding any token, issuer, or trade, and a peg holding historically is not evidence it will hold in future.

Alain AI Lab — independent crypto & AI research.
intelligencecrypto.org · Research is educational and analytical, not financial advice.