

Perpetual Futures and the Architecture of Always-On Markets

Ryan Mahaffy | *Applied Derivatives* | April 17th, 2026

What Is a Perpetual Future?

A perpetual futures contract provides leveraged exposure to an underlying asset with no expiration date. A trader can hold a position indefinitely, paying or receiving a periodic funding rate as the cost of that privilege. That single structural feature, the absence of a settlement date, separates the perpetual from every other futures contract in the history of derivatives markets. It is also what made the instrument both legally ambiguous in the United States and dominant everywhere else.

The concept was not born in crypto. In 1992, Robert Shiller proposed perpetual futures as a mechanism for creating derivative markets on illiquid assets, specifically residential real estate, where transaction frequency was too low to support dated futures contracts. He argued that a single perpetual contract could serve as a durable standard instrument for a given asset class without the liquidity fragmentation caused by rolling across contract months (Shiller, 1993). The mathematics were sound, but the market never materialized. Thirty years later, Bitcoin gave the design what real estate never could.

Mechanics

A perpetual is cash-settled and permits leveraged long or short positions against a notional value determined by an underlying asset price. Unlike a dated futures contract, there is no convergence date forcing the contract price toward spot. Without expiration, the arbitrage mechanism that tethers futures to spot disappears.

The replacement is the funding rate. At regular intervals (typically every eight hours on centralized exchanges, continuously on Hyperliquid), a payment is exchanged between long and short holders. The direction and magnitude are determined by the gap between the contract's mark price and the underlying index price. When the perpetual trades at a premium to spot, longs pay shorts. When it trades at a discount, shorts pay longs. The funding rate prices that imbalance. It incentivizes trades that narrow the

gap and penalizes the crowded side. Over time, this mechanism anchors the perpetual to its underlying through continuous cost-of-carry adjustment rather than hard convergence at expiry.

Mark price is distinct from last traded price. Most exchanges calculate mark as a function of the index price plus a decaying funding basis. Margin calculations, liquidation triggers, and unrealized P&L are all based on mark, not the last print on the order book. A trader can be liquidated by mark price movement even if no trade has executed at that level.

Contract Types

Linear contracts, sometimes called USDC-margined or stablecoin-margined, denominate margin, P&L, and settlement in a stable asset. A 1% move in the underlying produces a 1% move in contract value. Linear contracts are the dominant structure.

Inverse contracts, developed by Alexey Bragin for the ICBIT exchange in 2011 and popularized by BitMEX, denominate everything in the underlying asset (He, Manela, Ross, & von Wachter, 2024). A trader posting Bitcoin as margin to trade a Bitcoin inverse perpetual is simultaneously long the contract and short Bitcoin as collateral. The nonlinear payoff introduces convexity absent in linear contracts. The structure made sense before stablecoins were widely available and has since given way to linear products.

Quanto contracts fix the exchange rate between the underlying and the settlement currency. They exist, but volumes are negligible, and the structure introduces additional basis risk. Not material to this paper beyond acknowledgment.

Risk Mechanisms

Two mechanisms are unique to perpetual futures markets and absent from cleared derivatives: the insurance fund and auto-deleveraging (ADL). The insurance fund is a pool of assets, accumulated primarily from liquidation fees, that absorbs the difference when a liquidated position closes at a price worse than its bankruptcy price. Under normal conditions the fund grows, under stress it draws down.

If the insurance fund is insufficient, ADL activates. ADL forces the most profitable, highest-leverage traders on the opposite side of the market to have their positions reduced involuntarily at the bankrupt trader's bankruptcy price. The ADL counterparty does not choose to exit. They are exited at a price that may be substantially worse than the current mark. This separates perpetual futures from CME-cleared futures, where the clearinghouse guarantee, default fund, and member liability structures prevent one participant's insolvency from touching another's P&L. In perpetual markets, that protection is conditional on insurance fund balance. When the fund runs dry, the loss socializes.

Origin and Adoption Timeline

The instrument that reached crypto descended from Shiller's paper and was built into a working product by Bragin. Bragin's 2011 inverse perpetual at ICBIT referenced Shiller explicitly and established the funding rate as a practical mechanism for keeping a no-expiry contract anchored to spot. BitMEX refined and popularized the structure beginning in 2016, launching the first widely traded Bitcoin perpetual swap (He et al., 2024). Volume migrated rapidly from dated futures to perps. By the time Binance, OKX, and Bybit had built out their perp offerings in 2018 and 2019, dated crypto futures had become a secondary venue.

Why Perps Took Over

Perpetuals now account for the majority of crypto futures activity globally, on the order of 90% by most industry estimates, exceeding spot trading volume across major venues (Coinbase Institutional, 2025; CFTC, 2025). Three factors explain the concentration: structural advantages over dated futures, network effects as liquidity consolidated, and a regulatory environment that pushed the entire market offshore.

Eliminating the roll: In dated futures, continuous exposure requires closing the expiring contract and reopening the next. The roll carries bid-ask twice, slippage in back-months, and basis risk on execution. For high-frequency strategies or large positions, the cost is material. The perpetual removes it. One

contract, one position, held indefinitely. The funding rate is the cost of carry made explicit and continuous rather than embedded in the basis and realized at roll.

Leverage accessibility explains the rest. Dated futures on regulated U.S. venues operate under CFTC margin requirements. CME Bitcoin futures carry initial margin requirements in the range of 40 to 50% of notional, implying maximum effective leverage of roughly 2 to 3x for a standard account (CME Group, 2025). Offshore perpetuals launched with leverage caps of 20x, 50x, and on some altcoin markets up to 1000x before position limits were introduced. The risk profile is different, and so is the appeal to a certain class of speculator.

Funding as Price Discovery

The mechanical framing of funding as expiry replacement is accurate but incomplete. Funding is also the most information-dense signal in crypto derivatives. Reading it is a direct edge for anyone operating in the space.

Positive funding means longs pay shorts. Modest positive funding reflects normal conditions: the convenience of perpetuals creates a small structural premium shorts collect. *Persistently elevated positive funding reflects crowding, and the conditions for a liquidation cascade build with it.* Negative funding is the mirror image. Deeply negative funding in a declining market indicates the short side is overcrowded and sets up the conditions for a squeeze on any recovery.

The FTX collapse in November 2022 produced the cleanest single example. Funding on solvent exchanges turned sharply negative as traders rushed to hedge exposure, while FTX itself showed anomalous positive funding as trapped longs could not exit. For participants watching cross-venue basis in real time, that divergence was a solvency signal before any public announcement.

Cross-venue basis is always the granular version of the same signal. Persistent divergence between Hyperliquid's BTC funding rate and Binance's BTC funding rate can reflect venue-specific positioning, oracle differences between reference prices, or friction that prevents arbitrageurs from closing the spread.

He et al. (2024) document that deviations of crypto perpetual futures from no-arbitrage prices are substantially larger than those observed in traditional currency futures, and that these deviations have declined over time. The residual is gas cost, bridge latency for capital moving between chains, and the fact that some CEX-native market makers do not yet operate on-chain at comparable scale. As those frictions compress, the basis compresses further. Monitoring that basis provides a real-time read on market efficiency that did not exist before perpetuals created a continuous, publicly visible cost-of-carry market for crypto.

The Regulatory Cage and the Onshoring Wave

For roughly a decade, U.S. regulators addressed perpetual futures through enforcement rather than rulemaking. The CFTC charged BitMEX in 2020 for operating an unregistered futures exchange and violating anti-money-laundering requirements (Pillsbury, 2025a). That case established that perp exchanges serving U.S. customers were operating illegally. It did not resolve the underlying classification question: futures, requiring DCM registration, or swaps, falling under the Dodd-Frank swap dealer framework. Exchanges responded by geofencing U.S. users. Volume continued to grow offshore.

On April 21, 2025, the CFTC issued Release 9069-25, a formal Request for Comment on the trading and clearing of perpetual-style derivatives. The RFC asked whether perpetuals should be classified as futures or swaps, whether existing margin and disclosure requirements were sufficient, and what manipulation risks were specific to the no-expiry structure (CFTC, 2025). This was the first formal acknowledgment by U.S. regulators that perpetuals warranted a dedicated framework rather than enforcement-by-analogy.

On July 21, 2025, Coinbase Derivatives, a registered CFTC Designated Contract Market, self-certified perpetual-style Bitcoin and Ethereum futures contracts. These carry a five-year expiration and economic mechanics approximating a true perpetual, a structure that fits inside the existing DCM framework without requiring a new product classification (Pillsbury, 2025a). The CFTC raised no objection within the 30-day review window. The SEC did not challenge the ETH listing, an absence that

reinforced the market's working assumption that Ethereum is treated as a commodity under current regulatory posture.

On September 5, 2025, the SEC and CFTC chairs issued a joint statement announcing regulatory coordination under the label "Project Crypto" and identifying perpetual contracts as a priority for concurrent action. The statement proposed innovation exemptions using existing exemptive authority to bring perps onto U.S.-regulated platforms (SEC & CFTC, 2025; Pillsbury, 2025b). CFTC Chair Michael Selig stated in March 2026 that a U.S. perp framework would be released within weeks and framed prior U.S. posture as the direct cause of offshore liquidity migration (CoinDesk, 2026a). The Digital Asset Market Structure Clarity Act of 2025 passed the House and awaits Senate action. Its passage would provide statutory definitions distinguishing digital commodities from digital securities. This “innovation exemption” framework provides the workable interim path.

The competitive implication is asymmetric. CFTC-regulated perps on BTC and ETH do not directly compete with a permissionless ecosystem that lists perpetuals on crude oil, single equities, and custom indices (more on this in the next section). The onshoring wave addresses U.S. institutional access to the two largest crypto markets. The offshore innovation addresses everything else.

Case Study: Hyperliquid and Programmatic Transparency

Hyperliquid is the most analytically useful venue to examine because every order, fill, and liquidation is public in real time. Funding, open interest, vault balances, liquidations, and oracle updates are on-chain events rather than periodic disclosures. The distinguishing feature is transparency; vertical integration alone does not capture it. CME is also vertically integrated across matching, clearing, and risk management. Binance operates spot, margin, and perpetuals on a single platform. What Hyperliquid adds is that every state transition is auditable by any observer without exchange cooperation.

The infrastructure is a purpose-built Layer 1 blockchain. HyperCore, the execution layer, runs an on-chain central limit order book capable of processing roughly 200,000 orders per second with sub-second

finality (Hyperliquid Foundation, 2025). There is no off-chain matching engine publishing snapshots. The book is the chain. HyperEVM is the Ethereum Virtual Machine-compatible (EVM) layer that shares state with HyperCore, allowing vault contracts to place orders directly into the book without bridging or latency overhead. HYPE is the native token: gas, validator stake, governance, and the target of a rare fee-buyback mechanism that routes a share of protocol trading fees toward token demand.

Until October 2025, every perp market on Hyperliquid was operated by the protocol's validator set. Listing required validator consensus. HIP-3, activated October 13, 2025, changed that (CoinDesk, 2025). Any party meeting the staking requirement can deploy a perpetual exchange on HyperCore using the same order book, margining, and API as validator markets. A deployer stakes 500,000 HYPE, currently roughly \$20 to 25 million depending on HYPE price, and gains the right to list markets. The first three markets in any deployment skip auction. Additional markets are acquired through a Dutch auction shared across deployers. Deployers set their own leverage limits within protocol bounds, publish their own oracles, and receive 50% of fees generated by their markets. Users pay 2x the fee of validator-operated markets.

The staking requirement is bonded accountability rather than a permission gate. A deployer whose oracle fails, whose market is manipulated, or whose actions harm the network is subject to slashing, which can eliminate the staked position. Slashed HYPE is burned. A seven-day unstaking queue prevents exit before consequences are applied. The structure is a novel solution to the oracle trust problem in permissionless derivatives. Traditional permissionless DEXs either rely on governance votes to vet oracles or accept any input without recourse. HIP-3 replaces both with economic skin in the game: the deployer who runs the oracle has 500,000 HYPE at risk if it misbehaves.

Trade.xyz, the first HIP-3 deployer, launched 24/7 perpetual markets on U.S. equities including Microsoft, Apple, Nvidia, and Amazon, along with a synthetic Nasdaq index. By early April 2026, real-world-asset open interest on HIP-3 reached roughly \$2.3 billion, and HIP-3 markets represented a material share of total Hyperliquid trading volume (The Block, 2026). Of the top 30 HIP-3 markets by

open interest, only seven were crypto pairs. The remaining 23 were tokenized equity futures, S&P 500 and Nasdaq index contracts, gold, silver, and crude oil (Ripple, 2026). In March 2026, S&P Dow Jones Indices licensed the S&P 500 to trade.xyz, producing the first officially licensed on-chain S&P 500 perpetual (CoinDesk, 2026b). Institutional recognition of permissionless on-chain venues moved from lawsuit to license within eighteen months.

The Weekend Market

Traditional market hours are a constraint finance has treated as immutable for generations. NYSE and NASDAQ operate 9:30 a.m. to 4:00 p.m. Eastern, five days a week. COMEX commodity futures close Friday afternoon and reopen Sunday evening. Geopolitical events, OPEC decisions, and macroeconomic shocks do *not* respect that schedule. The gap between when information arrives and when it can be traded has always been a structural feature of traditional markets.

On the weekend of February 28, 2026, coordinated U.S. and Israeli strikes hit Iranian nuclear facilities on a Saturday. COMEX was closed. CME crude oil futures were dark. Traditional commodity traders had no venue to act on the information until Sunday evening at the earliest. Hyperliquid was open. Over that weekend, oil perptuals on Hyperliquid rose over 5%, gold rose 1.3%, and silver gained 2%. The platform processed over 175,000 trades totaling \$257 million in commodity notional while legacy venues were unavailable (Hiperwire, 2026). Open interest exceeded \$1.1 billion. Bloomberg cited Hyperliquid's on-chain oil prices as a reference in Iran risk coverage ahead of CME and NYMEX reopening (Yahoo Finance, 2026).

When underlying traditional markets are open, HIP-3 deployers publish oracle prices derived from live external data feeds, updated roughly every three seconds. Mark price tracks the reference asset closely. When the underlying closes, deployers transition to an EMA-anchored internal price derived from order-book impact rather than live external feeds. For equity perptuals, the weekend oracle remains largely anchored to Friday's close. The 1% per oracle tick clamp that governs mark price updates prevents runaway moves but also means the mark can lag a sharp gap by a meaningful margin before catching up.

The practical implication: gap risk on Hyperliquid is distributed across oracle update cycles during closure, with funding rates spiking as one-sided demand builds against an anchored mark. On CME the Sunday evening open is the gap. On Hyperliquid the gap unfolds in slow motion. Position sizing should reflect it.

Liquidity, Liquidation, Liability

The plumbing of a perpetual futures exchange determines what happens when markets move fast. Centralized exchanges obscure the plumbing. Hyperliquid puts it on-chain. The same mechanics operate, less visibly, on every perp venue in the market.

Liquidity on HyperCore comes from three sources: professional market makers connected via API, vault-based strategies open to any depositor, and retail limit orders. The HLP, Hyperliquid's native liquidity provider vault, is the largest single source of resting liquidity across the platform. It accepts USDC deposits, runs algorithmic market-making across listed perps, and distributes returns based on vault performance. The HLP's strategy, P&L, positions, and balances are publicly visible on-chain (Hyperliquid Foundation, 2025). CEX market maker arrangements are not. HLP returns come from realized spread capture, funding rate income when the vault holds net positions on the favored side, and a share of liquidation proceeds when the vault absorbs liquidated positions. The HLP is the first-line absorber when no other counterparty takes the other side at the liquidation price. Weekend and after-hours exposure amplifies the risk: when the oracle is EMA-anchored and a real gap is building in the underlying, the HLP can accumulate inventory that looks profitable against a stale mark and is severely underwater at the next traditional open.

During a liquidation cascade on October 10, 2025, approximately \$10 billion in liquidations were processed on Hyperliquid out of roughly \$19 billion globally (CoinDesk, 2025). The insurance fund absorbed a significant drawdown. ADL triggered in some markets. The fund held. The event tested the architecture under the conditions most stress tests do not anticipate: a sequential cascade where each wave

of liquidations depressed prices enough to trigger the next wave. The event revealed the boundary conditions of the current capitalization model.

ADL in this cascade worked as designed. The mechanism is not Hyperliquid-specific; every major perp venue, CEX and DEX alike, implements some form of the insurance-fund-plus-ADL architecture (He & Manela, 2024). For any trader active in perp markets, the insurance fund balance is a live risk indicator. A fund drawing down under stress is a direct signal that ADL exposure is rising, and on Hyperliquid that data is on-chain and publicly visible in real time. On a CEX, it is not.

Contrast the CME framework. A clearing member at CME bears defined liability for customer account shortfalls under CME Rulebook Chapter 9. The clearinghouse default waterfall begins with the defaulting member's margin and default fund contribution, progresses through the mutualized default fund, and terminates at clearinghouse equity. A futures commission merchant is regulated, audited, and carries customer segregated funds protected under CEA § 4d(a)(2) (CFTC, 2024). A winning position at CME does not become collateral for someone else's bad trade. At Hyperliquid, under extreme conditions, it can. None of the CME legal infrastructure exists on Hyperliquid. On-chain transparency is only a partial substitute for legal accountability. What differs on Hyperliquid relative to a CEX is that fund balances and ADL triggers are visible in real time. A trader who watches the fund can observe drawdown under stress and assess their own ADL exposure before it arrives. Transparency simply converts it from a black-box event into a monitorable one.

Risks

Several of the following risks are idiosyncratic to Hyperliquid. Others are systemic to the perpetual futures architecture regardless of venue.

Leverage and cascade: Hyperliquid's validator-operated markets permit leverage up to 40x on certain assets. HIP-3 deployers set their own leverage limits, bounded only by the threat of slashing. The practical ceiling is therefore defined by deployer risk appetite rather than protocol rule. A deployer who

sets 100x leverage on a low-liquidity equity perp is creating conditions for cascades that can propagate into adjacent markets through shared collateral. The October 2025 cascade demonstrated that. Forced liquidations in a few large BTC positions created selling pressure that reduced mark prices across multiple markets simultaneously. Traders using cross-margin with exposure in unrelated markets found margin ratios deteriorating in positions that were not directly impacted.

Oracle: Every non-crypto perp on Hyperliquid depends on a deployer-published oracle. The oracle is slashable if it misbehaves. A bad oracle price that enters the system can persist for multiple ticks while positions are liquidated against incorrect fair value. The 1% per tick clamp prevents runaway mark moves but also prevents the oracle from correcting quickly after error. Weekend oracle behavior is a specific version of the same risk. Informed traders price it, uninformed traders do not.

Regulatory and access: Hyperliquid prohibits U.S. persons from accessing the platform. The prohibition is IP-based, not KYC-verified across all users. The gap between formal prohibition and actual enforcement creates latent legal risk for both the platform and U.S.-based participants who access it via VPN. The broader question for tokenized equity and commodity perps is unresolved. A perpetual on Nvidia common stock, settling in USDC, operated by a permissionless deployer outside U.S. jurisdiction, occupies ambiguous territory under securities law. CLARITY Act passage would resolve it.

Concentration and deployer. Trade.xyz accounts for the large majority of HIP-3 open interest. That concentration creates a single point of operational dependency for most non-crypto perp volume on the platform. An oracle failure at trade.xyz, a legal challenge to its equity perp operations, or a technical failure in its market-making infrastructure would remove the primary source of commodity and equity perp liquidity simultaneously. Whether 500,000 HYPE in staked capital is sufficient deterrence against involuntary failure modes such as technical error or regulatory seizure remains an open question. Centralized perp venues have failed before. On-chain transparency changes the failure mode, not its possibility.

Closing Observations

Shiller proposed the instrument in 1992 for a market that never adopted it. Bragin built a functional version two decades later for an asset class no one had yet decided how to regulate. What followed was one of the faster instrument adoption cycles in modern derivatives: from niche crypto product to roughly 90% of global crypto futures volume in five years, then to 24/7 commodity and equity price discovery during geopolitical events that traditional markets could not respond to in real time.

This paper is a snapshot of a market in transition. The regulatory framework is being written right now. The infrastructure that made 24/7 oil trading during the Iran strike weekend possible is less than two years old. Liability questions raised by ADL and on-chain liquidation do not have settled legal answers. Competitive dynamics between onshore regulated perp products and offshore permissionless infrastructure have not played out. Perpetual futures are nevertheless the dominant structure for leveraged derivatives in the fastest growing asset class of the last generation. For traders, risk managers, regulators, and market structure researchers, the question has moved past whether to engage with the instrument. The market never closes.

References

- CFTC. (2024). Commodity Exchange Act, 7 U.S.C. § 6d(a)(2): Customer funds segregation. U.S. Commodity Futures Trading Commission.
- CFTC. (2025, April 21). Request for comment on the trading and clearing of perpetual-style derivatives (Release No. 9069-25). U.S. Commodity Futures Trading Commission. <https://comments.cftc.gov>
- CME Group. (2025). Bitcoin futures contract specifications and margin requirements. <https://www.cmegroup.com/markets/cryptocurrencies/bitcoin/bitcoin.contractSpecs.html>
- Coinbase Institutional. (2025). Crypto derivatives market structure overview. Coinbase Institutional Research.
- CoinDesk. (2025, October 13). Hyperliquid's HIP-3 upgrade to unlock permissionless perp market creation. <https://www.coindesk.com/business/2025/10/13/hyperliquid-s-hip-3-upgrade-to-unlock-permissionless-perp-market-creation>
- CoinDesk. (2026a, March 3). CFTC chief Selig to clear path for U.S. perpetual futures in coming weeks. <https://www.coindesk.com/policy/2026/03/03/cftc-chief-selig-to-clear-path-for-u-s-perpetual-futures-in-coming-weeks>
- CoinDesk. (2026b, March 18). Traders can now bet on the S&P 500 around the clock without ever touching a traditional stock exchange. <https://www.coindesk.com/markets/2026/03/18/traders-can-now-bet-on-the-s-and-p-500-around-the-clock-without-ever-touching-a-traditional-stock-exchange>
- He, S., Manela, A., Ross, O., & von Wachter, V. (2024). Fundamentals of perpetual futures (arXiv:2212.06888). <https://arxiv.org/abs/2212.06888>
- Hiperwire. (2026, March 4). Weekend trading and circuit breakers on Hyperliquid. <https://hiperwire.io/explainers/weekend-trading-circuit-breakers-hyperliquid>
- Hyperliquid Foundation. (2025). HIP-3: Builder-deployed perpetuals. Hyperliquid Documentation. <https://hyperliquid.gitbook.io/hyperliquid-docs/hyperliquid-improvement-proposals-hips/hip-3-builder-deployed-perpetuals>
- Pillsbury. (2025a, July 22). CFTC permits listing of perpetual futures on BTC and ETH. <https://www.pillsburylaw.com/en/news-and-insights/cftc-perpetual-futures-btc-eth-crypto-derivatives.html>
- Pillsbury. (2025b, September 8). Prospect of competitive onshoring of perpetual derivatives raised. <https://www.pillsburylaw.com/en/news-and-insights/joint-statement-perpetual-derivatives>
- Ripple. (2026, March 30). Ripple Prime expands Hyperliquid integration to gold, silver, and oil perpetuals. <https://coinpedia.org/news/ripple-prime-expands-hyperliquid-integration-now-trade-gold-silver-and-oil-on-chain/>
- SEC & CFTC. (2025, September 5). Joint statement from the chairman of the SEC and acting chairman of the CFTC. <https://www.sec.gov/newsroom/speeches-statements/joint-statement-atkins-pham-090525>
- Shiller, R. J. (1993). Measuring asset values for cash settlement in derivative markets: Hedonic repeated measures indices and perpetual futures. *Journal of Finance*, 48(3), 911-931.

The Block. (2026, April). Hyperliquid HIP-3 real-world-asset open interest reaches \$2.3 billion.
<https://www.theblock.co/>

Yahoo Finance. (2026, April). Crypto market hedges Iran war risks with 24/7 oil and gold trading.
<https://finance.yahoo.com/news/oil-gold-rally-hyperliquid-24-131937525.html>