Framing Net Carry for Cross-Venue Perpetual Funding Positions: A Settled-History, Depth-Walked, and Liquidation-Aware Method

Hang Liu
humarb Research, London, United Kingdom
ORCID: https://orcid.org/0009-0008-1059-3926
admin@humarb.ai · humarb.ai
November 2026

Document status. Methodological preprint prepared for repository submission. DOI: https://doi.org/10.5281/zenodo.23897014. This version is not peer reviewed.

Abstract

Cross-venue perpetual funding positions are often summarized by annualizing the current funding-rate spread, a practice that conflates transient quotes with holding-period cash and omits execution and margin risks. This paper develops an accounting framework based on settled funding history, asynchronous settlement grids, depth-walked entry and exit costs, and tier-adjusted liquidation distances. Net income is reported at reference and conservative ends, with unavailable inputs left explicitly unavailable. A nominal cost-coverage horizon is distinguished from recovery under time-varying funding. A limited replay motivates the historical basis but does not validate the estimator; out-of-sample performance is not reported.

Keywords: perpetual futures; funding rates; net carry; execution costs; liquidation risk; forecast evaluation
JEL classification: G12; G13; G14

Conflict-of-Interest Disclosure

Hang Liu. This research was conducted at humarb Research, London, United Kingdom. The author is affiliated with humarb Research, whose products may present quantities related to those discussed in this paper. No external sponsor influenced the reported framework.

中文摘要

跨交易场所永续合约资金费套利常以“当前费率差”年化展示，但该指标把瞬时报价误作持有期现金流，并忽略不同结算时刻、交易规模对应的盘口成本、出场成本及单腿强平风险。本文提出一个可审计的净持有收益会计框架：以已结算资金费历史的稳健汇总形成参考基准，在单次结算层面应用费率上下限，并在混合结算网格上计算两腿现金流；以四次吃单成交、逐档盘口成交成本、已测量的其他成本和明示的价差回归项估计净收入。结果分为参考端与保守端，后者是压力情形而非置信区间。关键输入缺失时，估计不可用而非以默认值补齐。本文还区分名义成本覆盖时长与时变资金费路径下的真实覆盖时间，并以维持保证金档位和近期方向性单日波动构造强平缓冲指标。有限的离线 replay 仅用于说明当前费率外推的偏差，不构成对生产估计器的样本外验证；本文未报告该估计器的样本外结果。框架旨在明确测量、压力测试与概率预测的边界，不构成投资建议。

1. Introduction

A perpetual future has no expiry, so its anchoring mechanism is not convergence at maturity but a sequence of transfers between long and short positions. When a matched route holds a long leg on one venue and a short leg on another, the difference between their funding rates can become a meaningful source of cash income. The simplest screen therefore takes the two currently displayed rates, subtracts them, and annualizes the difference. That operation is attractive because it is immediate and comparable, but it embeds three strong assumptions: the observed rates remain representative over the intended hold, both legs pay on a compatible schedule, and the route can be entered and exited at prices that do not depend on trade size. None is generally harmless. A method paper must consequently begin not with an annualized percentage, but with the accounting identity that connects funding cash, execution costs, other measured frictions, and price-spread changes over a stated holding period.

The distinction between a displayed rate and settled history is central. A current rate describes a particular interval or a venue’s current calculation state, whereas a multi-day position receives a sequence of realized transfers. A transient premium can revert while the position remains open, and a cap or dead zone can censor the published rate without removing the underlying price pressure. The appropriate object for a holding-period estimate is therefore not the latest print held constant, but a robust summary of what has actually settled, accompanied by an explicit stress case. He et al. (2022) and Dai et al. (2025) provide the broader treatment of perpetual-futures fundamentals and arbitrage; the present paper addresses a narrower measurement problem. It asks how a net-carry number should be framed so that its inputs, timing, costs, and omissions remain visible to a professional reader rather than being compressed into a superficially precise rate spread.

Execution creates a second distinction between quotation and measurement. A mid-price is not an executable fill, and the best bid or ask describes only the first displayed level. Once a route is evaluated for a stated per-side notional, entry and exit can consume several levels of the book. The cost must consequently be measured by walking the available depth, with the same base quantity used for entry and exit on both legs. This approach is related to implementation shortfall because cost is measured against a decision-time reference price (Perold, 1988), and to optimal-execution research because spread and book impact are

economically distinct (Almgren & Chriss, 2001). It deliberately stops short of fitting a parametric impact model: the measured ladder, rather than an assumed price-impact function, determines the cost whenever the displayed book can fill the amount.

Risk creates a third distinction. Market-neutral price exposure does not imply neutral margin exposure when the legs are isolated on separate venues. A favorable move on one leg cannot automatically rescue the other, so route-level risk depends on the weaker leg's distance to liquidation relative to the asset's own recent adverse movement. Public maintenance-margin tiers and deductions matter because a nominal tier rate alone can misstate the effective requirement. The method therefore combines a tier-aware isolated-margin approximation with a directional daily range statistic, while stating explicitly what it excludes: cross-margin offsets, accrued funding, liquidation fees, mark-price details, and auto-deleveraging. The result is a cushion ratio, not a probability of survival and not a route-level guarantee when either leg cannot be measured.

This paper makes four contributions. It defines reference and conservative net-income estimates within one accounting identity and treats unavailable inputs as unavailable results. It gives an exact settlement-block construction for mixed funding intervals and applies caps at the settlement scale. It defines common-quantity, depth-walked entry and exit costs together with a nominal cost-coverage horizon whose extrapolation semantics are not confused with an integral equation. Finally, it separates implemented measurements from planned probabilistic evaluation, including proper scoring rules, interval tests, and backtest-overfitting diagnostics. A companion snapshot study, *Measuring Executable Net Carry in Cross-Venue Perpetual Funding-Rate Arbitrage* (humarb Research, 2026a), concerns measured route outcomes; this paper instead documents the conceptual model and supplies no new performance claim.

The scope of evidence is intentionally limited. The only realized-versus-extrapolated figures are from a short offline replay of a current-rate predictor, reported to motivate the settled-history basis. They are not an evaluation of the estimator specified here. No out-of-sample result for reference-end net income, no coverage result for a probabilistic interval, and no PBO or CSCV statistic is claimed. This separation matters because a conservative end is an auditable stress case rather than a quantile, and a walked exit cost is a measurement on a current book rather than a forecast of a future fill. The paper's contribution is thus a disciplined measurement framework and an evaluation design, not a demonstration of profitability.

2. Related Literature

The economic role of funding follows from the contract’s lack of maturity. Periodic transfers discourage a persistent disconnect between the perpetual price and its underlying index, and they turn a matched long-short position into a carry exposure rather than a simple directional hedge. He et al. (2022) analyze the fundamentals of perpetual futures, while Dai et al. (2025) study arbitrage in perpetual contracts and the implications of funding-rule structure. This literature motivates treating funding as a cash-flow mechanism with explicit rule constraints, not as an exogenous yield quote. The present framework adopts that perspective at the level of accounting: signs are assigned by leg direction, rates are normalized only after their settlement scale is known, and caps are applied to the amount paid at each settlement. It does not derive equilibrium prices or no-arbitrage bounds, because its object is the ex ante framing of net carry under observed public inputs.

Cross-venue price differences also connect the method to cryptocurrency arbitrage and relative-value trading. Makarov and Schoar (2020) document trading and arbitrage in cryptocurrency markets, including the economic significance of segmentation across venues. Gatev et al. (2006) provide a foundational relative-value framework in which convergence between related prices can generate returns. These studies support the inclusion of a spread-reversion term when a defensible historical estimate exists, but they do not justify crediting convergence mechanically. In the accounting identity developed below, reversion is a separate signed term with reference and adverse treatments. If it cannot be formed, it is omitted and named rather than silently set to a favorable value. This separation prevents an assumed convergence benefit from being mixed with settled funding or measured execution cash.

Execution-cost measurement draws on two complementary ideas. Perold’s (1988) implementation shortfall compares the outcome of implementation with a paper benchmark established at the decision point. Almgren and Chriss (2001) formalize execution trade-offs and market impact in an optimal-execution setting. The method here uses the decision-time mid as a measurement reference and decomposes crossing into spread and depth effects, but it does not estimate an optimal trading trajectory or a parametric impact law. That choice is deliberate for a route estimate made at a stated amount: a book walk responds directly to the displayed liquidity that the amount would consume. The trade-off is that the measurement is conditional on a fresh snapshot and says little by itself about the future exit book, a limitation addressed through a separate conservative exit-cost input and explicit evaluation design.

The statistical choices are informed by robust estimation and forecast evaluation rather than by a single return distribution. Hampel (1974) provides the influence-curve foundation for robust estimation, and Leys et al. (2013) explain why median absolute deviation can be preferable to mean-and-standard-deviation screening in the presence of outliers. Hyndman and Fan (1996) document alternative sample-quantile conventions, which matters when evaluation statistics are compared across implementations. Parkinson (1980) relates extreme prices to variance estimation, but the daily high-low information here is used only as a range statistic, not as a fitted volatility model. For planned evaluation, proper scoring rules (Brier, 1950; Gneiting & Raftery, 2007), interval tests (Kupiec, 1995; Christoffersen, 1998), forecast comparisons (Harvey et al., 1997), selection adjustment (Smith & Winkler, 2006), and backtest-overfitting diagnostics (Bailey & López de Prado, 2014; Bailey et al., 2017) serve distinct purposes and are not interchangeable.

3. Funding Mechanism and Settlement Timing

3.1 Cash signs and the premium family

Consider one settlement on leg $j \in \{\mathrm{L}, \mathrm{S}\}$, where L and S denote the long and short legs of a matched route. Let A be per-side notional in USD and let r be the settlement rate as a fraction. A positive rate is paid by longs and received by shorts. The cash sign is therefore $s_{\mathrm{S}} = +1$ and $s_{\mathrm{L}} = -1$, so a position of notional V has funding cash $s_j V r$ for that settlement. This convention makes the accounting identity independent of whether a displayed venue number is described as a payment, a receipt, or a funding percentage. It also makes clear that a favorable negative rate on the long leg and a favorable positive rate on the short leg can both contribute positively. A route is not evaluated by the absolute size of either rate, but by the signed sum over actual payment events on both legs.

A common funding family begins with a premium index comparing impact bid p^{bid} and impact ask p^{ask} with index price p^{idx}. Define the instantaneous premium P by

$$P = \frac{\max(0, p^{\mathrm{bid}} - p^{\mathrm{idx}}) - \max(0, p^{\mathrm{idx}} - p^{\mathrm{ask}})}{p^{\mathrm{idx}}}. \tag{1}$$

If the index lies inside the quoted spread, the two positive parts vanish and P has a point mass at zero. Let $\bar{P}$ be the interval average, I the interest component, and δ the dead-zone

half-width. For an interval of τ hours expressed on an eight-hour convention, with lower and upper per-settlement caps c^{lo} and c^{hi}, the settlement rate is

$$r = \mathrm{clip}_{[c^{\mathrm{lo}},\,c^{\mathrm{hi}}]}\left(\frac{\tau}{8}\left[\bar{P} + \mathrm{clamp}(I - \bar{P}, -\delta, +\delta)\right]\right). \tag{2}$$

Inside the dead zone, the bracket equals I; outside it, the premium continues with a dead-zone adjustment until the caps bind.

The dead zone and caps matter statistically as well as institutionally. They can create repeated published values and compress the variance of the observed rate relative to the latent premium, so an estimator should not interpret low variance in the published series as low uncertainty in the underlying price pressure. They also create a regime distinction between a leg at a cap and a leg away from it. Published venue parameters illustrate the scale: an interest component of 0.01% per eight hours and a dead-zone half-width of 0.05% are common, but they are not universal constants. The formula family must therefore be checked before two venues are compared. Some families omit the $\tau/8$ scaling, some scale I directly to the interval, and others clamp the premium itself rather than $I - \bar{P}$, producing different behavior outside the dead zone.

Decentralized and alternative venues widen this institutional problem. A venue may pay a fraction of an eight-hour-scale rate every hour, subject to a protocol cap, or may determine funding from long-short open-interest imbalance and borrowing rather than from a premium index. Such differences prevent a universal formula from being applied merely because two products are both called perpetual futures. The methodological response is to isolate the common accounting layer—sign, notional, payment instant, and settlement amount—from the venue-specific rate-generation layer. The net-carry framework requires the former and treats the latter as a public rule input. It does not infer an undisclosed rule or force unlike venues into a single premium equation.

3.2 Interval normalization and the settlement-block grid

Published funding intervals include 1, 2, 4, and 8 hours. For leg j, let τ_j be its interval in hours and let $y_j = r_j/\tau_j$ be the rate normalized to a fraction per hour. Normalization permits comparison of rates, but it does not authorize the assumption that cash arrives continuously. A position earns only at settlements that occur while it is open and for which it is eligible. A venue can also shorten a paying leg’s interval after a cap is touched, which changes the

number of cash events in a stress case. Continuously accruing venues are the exception in which the holding-period expression $s_j A H y_j$ is appropriate. The naive alternative—multiplying an hourly gap by elapsed hours for every venue—substitutes a convenient clock for the actual contractual cash calendar.

For discrete settlements, suppose leg j settles at $t_{j,n} = t_j^{\mathrm{next}} + n\tau_j$ and pays $\gamma_j = s_j A r_j$ at each event. When the legs have different intervals, their merged schedule repeats over the settlement-block period $\Lambda = \mathrm{lcm}(\tau_{\mathrm{L}}, \tau_{\mathrm{S}})$. Net cash per block is

$$d = \sum_j \gamma_j \frac{\Lambda}{\tau_j}, \tag{3}$$

and cumulative cash $S(t)$ satisfies $S(t + n\Lambda) = S(t) + nd$ for integers $n \geq 0$. This grid preserves both payment phase and lumpiness. It also supplies a useful early-stopping property: if $d \leq 0$, no later block can lift the running maximum above its value within the first block. A search for the first time cumulative funding covers a fixed cost can therefore stop after one block in that case, without approximating the payment path as continuous.

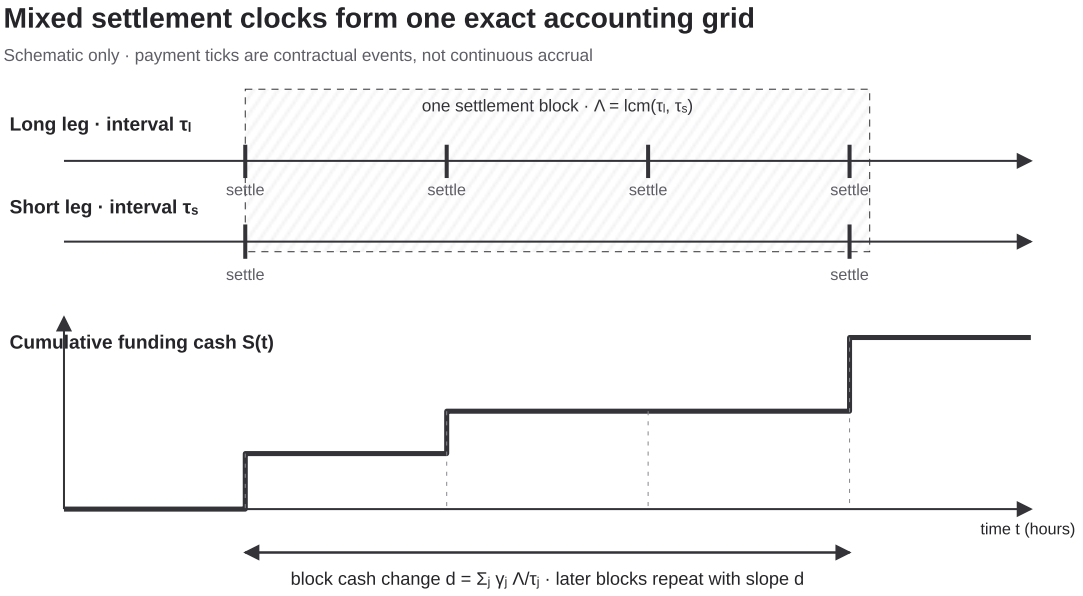


Figure 1. Settlement-block grid. Two legs settle on different clocks; their merged block has length Λ, and cumulative funding follows a staircase whose block-to-block change is d. Schematic only.

The grid’s advantage is most apparent over short holding periods or when a larger payment occurs late in a block. An hourly normalization can attribute income to hours that contain no settlement, while the grid leaves cash at zero until the contractual instant. This is not a claim that the phase of every future rate is known; rather, it is a requirement that known settlement mechanics not be discarded when forming the accounting path. The block construction also distinguishes the timing problem from the forecasting problem. Timing determines when a candidate rate would pay if it were realized; historical estimation determines what candidate rate is defensible. Combining the two into one annualized current gap obscures both and makes short-horizon coverage calculations particularly misleading.

4. Headline Estimator: Net Income with Two Auditable Ends

4.1 Accounting identity and annualization

The headline quantity is estimated net income in USD over holding period H hours for stated per-side notional A, together with a simple net annual percentage rate on the same basis. Let G denote funding cash, F the taker fees on entry and exit for both legs, K^{ent} and K^{ext} the walked entry and exit crossing costs, K^{oth} other costs when measured, and B the spread-reversion term. The reference end is

$$\Pi^{\mathrm{ref}} = G^{\mathrm{ref}} - F - K^{\mathrm{ent}} - K^{\mathrm{ext,ref}} - K^{\mathrm{oth}} + B^{\mathrm{ref}}. \tag{4}$$

The conservative end is defined by

$$\begin{aligned} \Pi^{\mathrm{con}} = \min\big(\Pi^{\mathrm{ref}},\ & \\ & G^{\mathrm{con}} - F - K^{\mathrm{ent}} - K^{\mathrm{ext,con}} \\ & \quad - K^{\mathrm{oth}} + B^{\mathrm{con}}\big). \end{aligned} \tag{5}$$

For either end Π, the corresponding annualized figure is

$$\mathrm{APR} = \frac{\Pi}{A}\frac{365 \cdot 24}{H} \times 100\%, \tag{6}$$

which is simple annualization, not compounding and not a realized return forecast.

The outer minimum in equation (5) ensures that the conservative end never exceeds the reference end even when separate input adjustments interact unexpectedly. The two ends are

not endpoints of a confidence interval. The reference end uses a robust settled-history basis and current measured exit cost; the conservative end combines less favorable funding evidence, a higher recent exit-cost candidate, and adverse spread treatment where those quantities are supported. No probability is attached to the interval between them. This is preferable to assigning a nominal coverage level because the available rate history does not, by itself, establish a distribution for future net income. A reader can audit each changed input instead of interpreting a statistically precise but unsupported band.

For a discretely settling leg with m_j eligible settlements inside the hold, funding cash is

$$G = \sum_j s_j A m_j \, \mathrm{clip}_{[c_j^{\mathrm{lo}}, c_j^{\mathrm{hi}}]}(\tau_j y_j). \tag{7}$$

The cap is applied to $\tau_j y_j$, the per-settlement rate, rather than to the hourly normalized rate. A spot leg contributes zero funding cash, and a continuously accruing leg contributes $s_j A H y_j$. This separation prevents a long-interval venue from being clipped at the wrong scale. It also makes m_j a function of entry time and eligibility, not simply H/τ_j. The naive current-gap screen misses both features: it ignores how many payments can still be earned and can compare rates whose caps belong to different settlement intervals.

Other costs enter only when measured. Depending on the route, they can include borrow, currency conversion, transfer, settlement-currency basis, or an overnight gap for a session-bound market. An unmeasured cost is neither zero nor hidden: it is omitted from the arithmetic and identified in an excluded-costs list attached to the estimate. If an essential input is missing, the estimate itself is unavailable. This rule differs from completing a screen with default fees, default depth, or a favorable cached book. Defaults create the appearance of full coverage while shifting unmeasured risk into an invisible assumption. The framework instead sacrifices the number of available routes to preserve the meaning of each published number.

4.2 Spread reversion and unavailable routes

Let $\beta = (p_{\mathrm{S}}^{\mathrm{mid}} - p_{\mathrm{L}}^{\mathrm{mid}}) / [\tfrac{1}{2}(p_{\mathrm{S}}^{\mathrm{mid}} + p_{\mathrm{L}}^{\mathrm{mid}})]$ be the entry spread between leg mid-prices as a fraction of their mean. Over the hold, convergence contributes $A(\beta_{\mathrm{entry}} - \beta_{\mathrm{exit}})$. The reference end credits a typical reversion and the conservative end uses an adverse one, so $B^{\mathrm{ref}} \geq B^{\mathrm{con}}$. The economic motivation follows relative-value convergence and cross-venue price persistence (Gatev et al., 2006; Makarov & Schoar, 2020), but the method does not infer

that every spread must converge within H. The way typical and adverse levels are formed, and the circumstances under which the term is omitted, are not disclosed. When the term cannot be formed, it is omitted and named rather than replaced by zero without comment.

This treatment separates three questions that a raw funding gap combines: how much funding may be paid, what trading costs must be crossed, and how the relative prices of the legs may change. The spread term can have either sign, and its exclusion is relevant to interpretation even though it is not literally a fee. Likewise, the conservative end is not obtained by applying a uniform haircut to the reference result. Funding, exit, and spread inputs can move by different amounts and in route-specific directions. A uniform haircut would be easier to display but would erase the accounting sources of stress. Figure 2 therefore presents the ends as an accounting decomposition rather than as a statistical density or a pair of forecast quantiles.

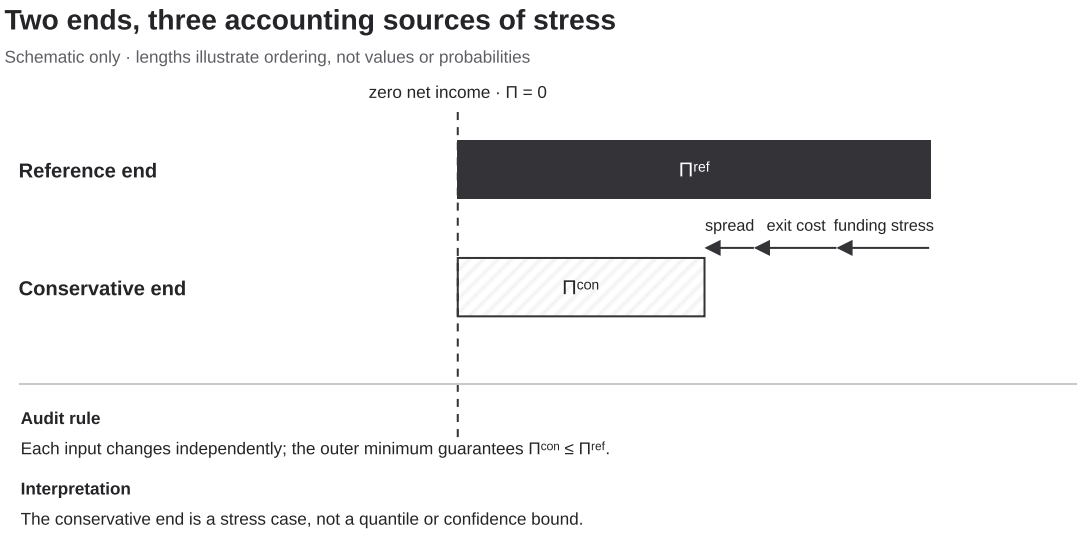


Figure 2. Reference and conservative ends. The conservative case combines less favorable funding, higher exit cost, and adverse spread treatment; the outer minimum preserves ordering. The zero line is accounting breakeven, not a probability threshold. Schematic only.

A route without an estimate receives a reason-specific unavailable state. Missing settlement history, insufficient depth, stale books, absent public margin tiers, and unsupported conservative evidence are different failures and should not be collapsed into a generic “no

data” label. Nor should a route be assigned a simulated number merely to maintain coverage. This discipline is part of the estimator rather than an interface detail: an accounting quantity exists only when its necessary inputs exist. It also prevents favorable selection through imputation, because a venue with sparse public information cannot be made to look attractive by filling gaps with optimistic assumptions. The result is a smaller but more interpretable set of estimates.

5. Reference Rate Basis

5.1 Why settled history replaces the latest print

A multi-day hold is paid over many settlements, while the latest displayed rate can reflect a transient premium. The reference basis therefore uses a robust summary of settled history for each leg and for the route differential. Its exact construction, statistical windows, thresholds, weighting, trimming, and operating rules are not disclosed. At the conceptual level, robustness limits the influence of a single dislocated print on a multi-day expectation. Median-centered deviations measured in median-absolute-deviation units are preferable to mean-and-standard-deviation screening when extreme observations are the object of concern (Hampel, 1974; Leys et al., 2013). This choice does not assert that the median is always the best location estimate; it asserts that a single spike should not automatically dominate a holding-period basis.

The naive alternative is to hold the current gap flat over the entire holding period. When premiums are mean-reverting toward a long-run level rather than necessarily toward zero, that alternative overstates the magnitude of both favorable income and adverse cost. A current extreme can therefore make a route appear unusually attractive or unusually expensive precisely when reversion is most consequential. Settled history does not remove regime risk, and a robust summary is not a forecast of the next premium. It provides a less reactive reference quantity whose connection to actual past cash is direct. A next-interval premium forecast exists only as a research component and does not enter the headline estimate described here.

A limited offline replay illustrates why this distinction matters. The replay used official settlements over 14 days in October 2026, per-side notional of USD 10,000, non-overlapping trades, and entry when predicted funding minus cost was positive. It extrapolated the then-

current rate gap. The realized-to-extrapolated funding ratio ϱ was 0.333 for a single settlement, 0.099 for 8 hours, 0.145 for 24 hours, and 0.128 for 72 hours. The replay relied mostly on top-of-book cost proxies rather than walked depth. It tested a current-rate extrapolation, not the estimator in this paper, and it covers one window and one regime. These figures motivate the historical basis; they are not performance results for the reference or conservative ends.

The asymmetry among the replay horizons is a warning against converting a short-window observation into a long-window expectation by multiplication alone. Realized funding over the replay was roughly 10–33% of the extrapolation, but that summary does not license a universal correction factor. Different entry selections, assets, rate regimes, and holding periods can produce different decay. A robust historical basis addresses the reference input directly rather than repairing a current-rate forecast with an estimated multiplier. It also preserves the distinction between measured replay evidence and undisclosed production construction. No bias magnitude is claimed beyond the reported replay, and no general law is inferred from four ratios.

5.2 Venues without public settlement history

Some venues publish no usable settlement history. For these venues, an observed rate series can be recorded from public displays. Until enough history has accumulated, the funding cash for such a leg uses a flat projection of the current rate, explicitly labeled as a current-rate basis and a product assumption rather than a statistical estimate. The conservative end gives no credit to a favorable current-rate leg and charges an adverse leg in full. Once the observed series is long enough, the reference switches to a robust summary of that series. The sufficiency rule and other implementation details are not disclosed; the important public distinction is between a displayed current-rate assumption and an observed-history estimate.

This two-stage treatment resolves two undesirable alternatives. Waiting for a long observed series would leave every such venue blank, while presenting the initial current-rate projection as equivalent to settled history would conceal a known bias direction. Under mean reversion, a flat current rate overstates the magnitude of both income and cost, although the size of that bias is not estimated. The conservative rule is intentionally asymmetric because favorable evidence is weaker than adverse evidence in the first stage. This is not a confidence statement and does not make the initial projection safe. It makes the epistemic status visible and

prevents an unseasoned series from acquiring the same interpretation as a history-backed reference.

6. Execution and Exit Cost

6.1 Depth-walked VWAP at a stated amount

Entry and exit costs are measured by walking a recent order book rather than by using the mid-price or the top-of-book spread alone. For an ask-side or bid-side ladder (p_i, q_i) ordered from the best level, a budget walk spends the stated amount by consuming levels in sequence until the budget is used. VWAP is money spent divided by quantity bought, or proceeds divided by quantity sold. The long leg buys from asks and the short leg sells into bids at entry; exit reverses both sides. This makes cost an explicit function of per-side notional. The same route can therefore have different net economics at different amounts even when funding rates and published fee schedules are unchanged.

For leg j with mid-price p^{mid} and entry VWAP $\bar{p}^{\mathrm{ent}}$, single-side entry slippage in basis points is $\sigma_j = 10^4 |\bar{p}^{\mathrm{ent}} - p^{\mathrm{mid}}| / p^{\mathrm{mid}}$. Route entry crossing cost is

$$K^{\mathrm{ent}} = A \frac{\sigma_{\mathrm{L}} + \sigma_{\mathrm{S}}}{10^4}. \tag{8}$$

The mid serves as a decision-time benchmark, consistent with the implementation-shortfall perspective (Perold, 1988), while depth consumption captures spread and book impact without fitting a parametric impact model. A mid-price alternative would assume execution between bid and ask and ignore depth altogether. A touch-only alternative recognizes the spread but assumes that the entire amount fills at the best level. Both understate cost precisely when trade size is large relative to displayed liquidity.

Exit must price the same position that was entered, so a common base quantity q is used across all four fills. The long leg pays $N_{\mathrm{L}}^{\mathrm{en}}$ to buy q and receives $N_{\mathrm{L}}^{\mathrm{ex}}$ when selling it; the short leg receives $N_{\mathrm{S}}^{\mathrm{en}}$ when selling q and pays $N_{\mathrm{S}}^{\mathrm{ex}}$ to buy it back. The exit halves measured against each leg’s mid are

$$\begin{aligned} K_{\mathrm{L}}^{\mathrm{ext}} &= q p_{\mathrm{L}}^{\mathrm{mid}} - N_{\mathrm{L}}^{\mathrm{ex}}, \quad K_{\mathrm{S}}^{\mathrm{ext}} = N_{\mathrm{S}}^{\mathrm{ex}} - q p_{\mathrm{S}}^{\mathrm{mid}}, \\ K^{\mathrm{ext}} &= K_{\mathrm{L}}^{\mathrm{ext}} + K_{\mathrm{S}}^{\mathrm{ext}}. \end{aligned} \tag{9}$$

The mids cancel in total crossing cost:

$$K^{\text{ent}} + K^{\text{ext}} = (N_{\text{L}}^{\text{en}} - N_{\text{L}}^{\text{ex}}) + (N_{\text{S}}^{\text{ex}} - N_{\text{S}}^{\text{en}}). \tag{10}$$

If a book cannot fill the required amount, cost is unmeasured rather than extrapolated beyond visible depth.

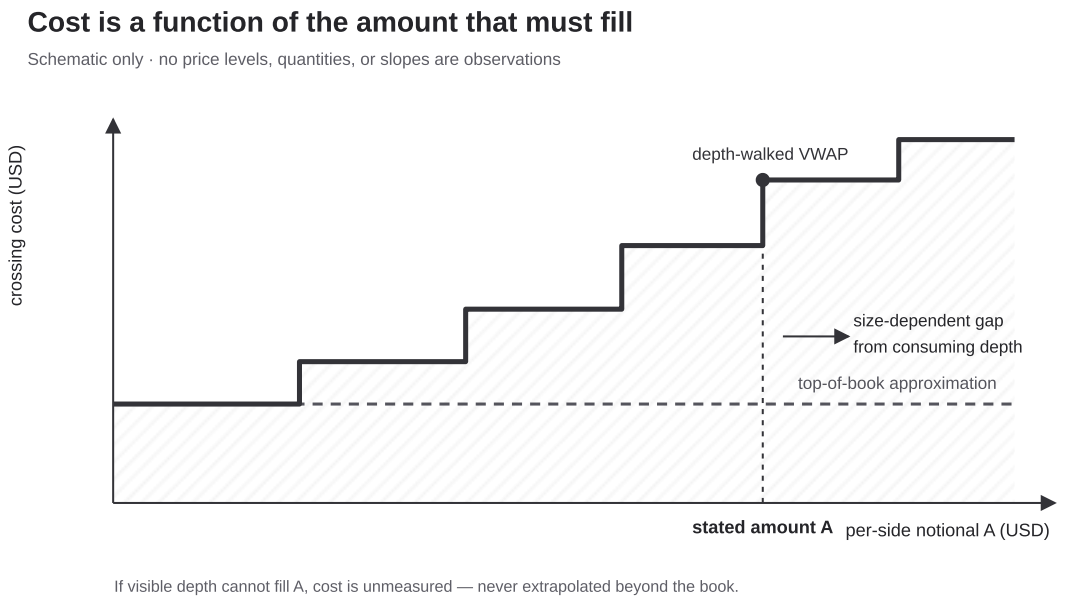


Figure 3. Depth-walked cost. The stepped VWAP path reflects consumption of successive book levels; the gap from a top-of-book approximation is size dependent. Schematic only.

Fee cash F is the sum over the four fills of the venue's published regular-user taker rate multiplied by actual fill notional; a negative rate is a rebate. Taker fees correspond to the spread-crossing assumption. A maker-fee alternative would require fill probability, queue position, cancellation, and adverse-selection assumptions that are not implemented. Using maker fees merely because they are lower would therefore contradict the execution model. Actual fill notional also matters when a walk crosses levels: applying a fee to an assumed mid-price amount would introduce a small but avoidable accounting mismatch between crossing cost and fee cash.

6.2 Freshness and conservative exit evidence

A live connection alone does not establish that an order book is current. The estimator refuses books whose recency cannot be confirmed and follows the principle that a cached book never improves an estimate. This rule trades fewer available estimates for protection against a favorable but stale liquidity picture. Exact freshness rules are not disclosed. The conceptual comparison is with a screen that reuses the last available book whenever new data are delayed: that alternative preserves coverage but can silently reduce measured cost or make an amount appear fillable. Freshness is thus part of measurement validity, not merely an operational safeguard.

The reference exit half is a walk on the book available at estimation time. It prices leaving under current displayed liquidity, not under a forecast book at the end of the hold. The conservative end uses a higher exit-cost candidate reflecting how much worse exit has recently been on the same route, with $K^{\mathrm{ext,con}} \geq K^{\mathrm{ext,ref}}$; its exact construction is not disclosed. Route-specific recent evidence is preferable to assuming that today's depth persists, because book depth can vary over minutes and one snapshot can be unusually deep or thin. It is also preferable to claiming a probabilistic exit-cost band without a validated book-evolution model. The residual risk remains explicit: neither end is a guaranteed future fill cost.

7. Cost-Coverage Horizon

The cost-coverage horizon summarizes the time an observed funding gap would need to offset measured round-trip cost if that gap remained constant. Let $C = F + K$ be one-off round-trip cost, with $K = K^{\mathrm{ent}} + K^{\mathrm{ext}}$, and let g be the current funding gap, short minus long, in fraction per hour. If $g > 0$, the nominal horizon is

$$T_{\mathrm{cc}} = \frac{F + K}{A g} \quad \text{hours.} \tag{11}$$

If $g \leq 0$, no horizon exists: the current gap does not cover cost. Reporting "infinite" would be less informative than stating that the sign prevents coverage. The numerator contains the four taker fees and walked crossing cost, so the horizon is amount specific. It is not a property of the funding rates alone. A route with the same gap can have different horizons at different notionals because depth consumption changes K.

The semantics of T_{cc} are deliberately narrow. It is not the solution of $\inf\{H : \int_0^H A g(t)\, dt \geq C\}$ for a time-varying true gap $g(t)$. The two coincide only when $g(t) = g$

is constant. The nominal horizon also excludes cost of capital and terminal value, so it must not be called a capital-recovery period. Its value is that it answers one transparent arithmetic question: how long would today's gap take to pay today's measured costs. The naive alternative of comparing one settlement's cash with the entire round trip makes nearly every route appear uncovered and ignores the intended holding period; flat extrapolation answers a more useful question but must carry its assumption in the name.

Replay evidence can be displayed beside, but not substituted into, this measured quantity. In the limited replay described above, realized funding was roughly 10–33% of the current-rate extrapolation, corresponding to $\varrho \in [0.099, 0.333]$. A mechanically adjusted reference $T^{*} = T_{\mathrm{cc}}/\varrho$ can illustrate sensitivity to that historical decay, but it never replaces T_{cc}. Replacing the nominal measure would rewrite a transparent current-gap calculation with a forecast-derived multiplier and would blur the line between measurement and prediction. It would also suggest unwarranted stability in ϱ across regimes. The two quantities can therefore be compared side by side only with distinct labels and an explicit statement that the replay did not evaluate the production estimator.

8. Liquidation-Risk Model

8.1 Tier-adjusted isolated-margin distance

Each leg is treated as a separately margined isolated position. A price move can push one leg to liquidation while the other is profitable, and profits on one venue are not assumed to offset margin deficiency on another. For one position with entry price E, quantity Q, notional $V = QE$, leverage L, tier maintenance-margin rate μ, and tier maintenance deduction D in USD, isolated margin is QE/L and the maintenance requirement at price p is $Qp\mu - D$. Setting margin plus profit-and-loss equal to the requirement gives

$$\begin{aligned} p^{\mathrm{liq}}_{\mathrm{long}} &= \frac{E(1 - 1/L) - D/Q}{1 - \mu}, \\ p^{\mathrm{liq}}_{\mathrm{short}} &= \frac{E(1 + 1/L) + D/Q}{1 + \mu}. \end{aligned} \tag{12}$$

Expressed as a fraction of entry price, the distances are

$$
\begin{aligned}
d_{\text{long}} &= \frac{1/L - \mu + D/(QE)}{1-\mu}, \\
d_{\text{short}} &= \frac{1/L - \mu + D/(QE)}{1+\mu}.
\end{aligned}
\tag{13}
$$

A spot leg has no liquidation distance.

The effective maintenance-margin rate is $\mu_{\text{eff}} = \mu - D/V$, the requirement divided by notional. Within a tier that includes a deduction, it is lower than nominal μ, and the deduction makes the requirement continuous across tier boundaries. Since the distance expressions increase with D, setting $D = 0$ understates distance and is conservative, but it is not the same calculation as the venue's own figure. The source illustration uses $\mu = 2.5\%$, $D = \text{USD } 1{,}875{,}000$, notional USD 200 million, and $L = 20$: distance is 3.526% with the deduction and 2.564% without it, a difference of about 0.96 percentage points. The example is an accounting illustration, not a recommended leverage or a route observation.

The applicable tier is the one containing per-side notional, and maintenance parameters come from public tier tables. Some venues publish a rate directly; some decentralized venues define maintenance as half the initial-margin rate at the tier's maximum leverage, $\mu = 0.5/L_{\max}$. No distance is defined when leverage exceeds the tier maximum, the amount exceeds the top tier, or $1/L \le \mu$. The calculation remains approximate: liquidation fees, accrued funding, mark-versus-last-price distinctions, and auto-deleveraging are excluded, and the exchange's own figure governs. A venue without usable public tiers is reported as unmeasured with a reason, never as safe and never with an inferred tier.

8.2 Directional move and route cushion

Distance alone is not comparable across assets. A fixed percentage cushion that appears large for a stable pair may be small for a newly listed token. The method therefore scales each leg's distance by that asset's recent adverse move. From a recent window of closed UTC daily candles with open o, high h, and low l, define the largest upward and downward single-day directional moves by

$$
w^{\uparrow} = \max_{u} \frac{h_u - o_u}{o_u}, \qquad w^{\downarrow} = \max_{u} \frac{o_u - l_u}{o_u}. \tag{14}
$$

A short leg is hurt by rises, so $w_j = w^{\uparrow}$; a long leg is hurt by falls, so $w_j = w^{\downarrow}$. If directional values are missing, the two-sided range is a conservative substitute. A missing day produces unknown rather than silently substituting an older day.

These are single-day range statistics, not cumulative drawdown and not fitted volatility forecasts. Although extreme prices have a long history in variance estimation (Parkinson, 1980), the model does not convert the range into a variance parameter. A multi-day hold can drift much farther than the largest recent one-day move. The cushion is consequently a short-horizon stress comparison, not a probability that liquidation will or will not occur. When both distance and swing are known for each measured leg, the route's cushion ratio is

$$k = \min_j \frac{d_j}{w_j}, \tag{15}$$

using the nearer leg against its own recent adverse move. A ratio below one means the largest recent move would have reached the approximate liquidation line; a larger ratio indicates more room under that specific comparison.

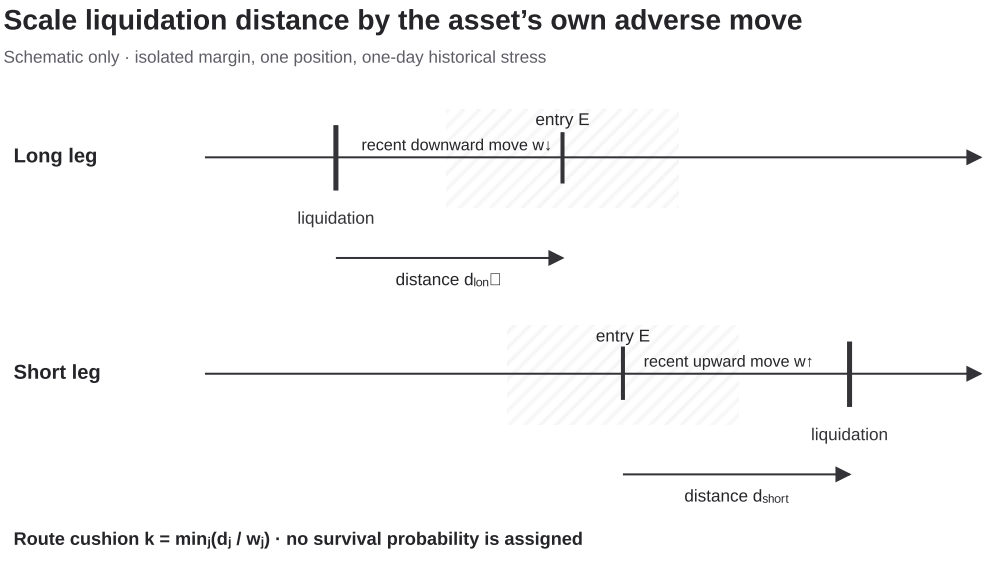


Figure 4. Liquidation cushion. The recent directional move is compared with each leg's own approximate distance; the nearer comparison governs the route ratio. Schematic only.

The largest loss on a failed leg is bounded by its assigned margin, expressed as $\min(Aw, A/L)$ under the model. This does not bound the operational loss of an unmanaged two-leg route after one leg fails, because the remaining leg can retain directional exposure. Data validation is also directional: a fall cannot exceed 100%, so a two-sided figure above that level is treated as bad data, whereas a rise above 100% in one day can be a real event for a short position in a thinly traded new listing. If one leg is measured and the other is not, only the measured leg's distance is meaningful and no route-level verdict is drawn. This prevents partial public information from becoming an unjustified safety label.

9. Data and Evaluation Framework

9.1 Input classes, provenance, and replay scope

The conceptual estimator uses public input classes: officially settled funding prints where venues publish them, current funding-rule parameters, recent order-book snapshots, regular-user taker schedules, maintenance-margin tiers, and closed daily candles. Current rate observations can form a self-recorded public-display series when official settlement history is absent. Public data provenance does not imply that every input is equally complete or historically reconstructable. Order-book depth is ephemeral, venue rules change, and missing values must remain missing. The framework therefore attaches freshness, coverage, and excluded-cost information to estimates rather than treating a successful API response as proof that an accounting quantity is well defined.

No new independent empirical sample is introduced for this method paper. The limited replay summarized in Section 5 used official settlements over its stated 14-day October 2026 window and the stated entry rule, with mostly top-of-book cost proxies. Its four realized-to-extrapolated ratios are reproduced from the source materials supplied for this work and were not independently re-estimated here. They should be read as motivating evidence about a different, current-rate predictor. They are not a backtest of reference-end net income, conservative-end coverage, or realized net APR. In particular, the replay's cost proxies are not substituted for the depth-walked methodology specified in this paper.

The production estimator's exact statistical construction, thresholds, weights, candidate sets, and operating rules are not disclosed. This restriction limits full independent reproduction of the headline estimate and is itself a methodological limitation. The public framework instead

makes the accounting identities, units, failure states, and evaluation logic inspectable. It also avoids a common form of accidental disclosure through empirical detail: no parameter values or route-selection rules are reverse engineered from examples. The four figures are schematic and contain no observations. The only numerical replay results and margin illustration are those explicitly present in the source materials supplied for this work.

9.2 Implemented measurements without reported outcomes

Evaluation begins by freezing the relevant quantity at decision time and comparing it with what can be observed later, but different quantities require different meanings of calibration. A measured cost can be checked for stability after a delay; a predictive distribution can be checked for coverage; a net-income estimate can be compared with realized settled cash net of measured costs. These are not interchangeable. Neither a walked cost nor a conservative end is a probabilistic forecast, so neither should be scored as though it claimed a coverage level. Keeping these objects separate prevents an apparently successful stability check from being mislabeled as predictive calibration.

A research forecast of the next-interval premium is stored with a mean and spread and scored after the interval completes against naive persistence, signed bias, the share inside its nominal band, and standardised-error quantiles relative to Gaussian reference values. Results are not reported. Its spread uses a widening-only, in-sample, one-sided calibration. It is not a conformal predictor and has no finite-sample coverage guarantee; any statement of 95% coverage is withdrawn. This example shows why a displayed interval must carry its construction semantics. A band can be useful for research diagnostics while still being unsuitable for a formal coverage claim.

Exit-cost stability is evaluated by re-walking the same route amount after a delay. Unfillable re-walks and expiries are counted rather than filled in, and the unresolved share accompanies any stability figure because dropping unresolved cases biases measurement toward cheaper outcomes. The permitted interpretation is stability of the taker exit-cost estimate on published books, not realized fill cost and not forecast error. Decision-outcome tracking separately freezes route, amount, horizon, and estimate at publication, then compares them with cash from official settled prints rather than sampled running rates. Band coverage is at most a point estimate at present, with no formal coverage test. Results for both components are not reported.

9.3 Planned probabilistic and overfitting diagnostics

Proper scoring rules are planned for genuinely probabilistic forecasts. For predictive CDF $\mathcal{F}$ and outcome x, the Brier score for event $\mathcal{E}$ is $(\chi_{\mathcal{E}} - 1[x \in \mathcal{E}])^2$, where $\chi_{\mathcal{E}}$ is the forecast event probability. Pinball loss for quantile ℓ_α is $\rho_\alpha(x - \ell_\alpha)$ with $\rho_\alpha(v) = v(\alpha - 1[v < 0])$, and

$$\mathrm{CRPS}(\mathcal{F}, x) = \int_{-\infty}^{\infty} \left(\mathcal{F}(z) - 1[x \le z]\right)^2 dz. \tag{16}$$

A strictly proper score is minimized in expectation only by the true distribution (Brier, 1950; Gneiting & Raftery, 2007). These scores answer different questions from exit-cost stability and should not be applied to a stress case that has no probability statement. A Gaussian closed form can be used where a forecast is explicitly Gaussian, but that convenience does not establish that the underlying premium is Gaussian.

Planned interval evaluation includes Kupiec's proportion-of-failures test and Christoffersen's independence and conditional-coverage tests (Kupiec, 1995; Christoffersen, 1998). They require non-overlapping observations or route-clustered treatment because overlapping holding periods create dependent outcomes. Forecast-comparison tests of Diebold–Mariano type, with the Harvey–Leybourne–Newbold small-sample correction, are planned for comparisons against naive forecasts (Harvey et al., 1997). Split-conformal bands with a genuine calibration and hold-out split are intended to replace the current in-sample widening-only scale. None of these planned items has a reported test statistic or coverage result in this paper.

Evaluation of the headline estimate itself will compare frozen Π^{ref} with realized settled funding net of measured costs, stratified by venue pair and amount. Selection must be modeled because routes are displayed precisely when their estimates are high. The optimizer's curse implies that selecting on noisy estimates can make post-selection outcomes disappointing even when individual estimates are unbiased (Smith & Winkler, 2006). A naive average over displayed routes would therefore mix estimator error with selection effects. Results are not reported. This is a substantive reason not to treat the limited current-rate replay as validation of the headline estimator.

Backtest-overfitting diagnostics are also planned rather than implemented as reported evidence. If configurations $c = 1, \ldots, N$ are compared over a performance series divided into n_b equal blocks, CSCV forms all $\binom{n_b}{n_b/2}$ splits into in-sample and out-of-sample halves. For each split, the configuration ranked best in sample is ranked among all N configurations out of sample. PBO is the share of splits in which that selected configuration falls below the out-of-sample median, equivalently the probability that the logit of its relative rank is no greater than zero (Bailey et al., 2017). PBO addresses selection among configurations, not whether the selected configuration earns money. The deflated Sharpe ratio provides a companion correction for multiple trials and non-normality (Bailey & López de Prado, 2014). No PBO, CSCV, or deflated-Sharpe result is reported.

10. Limitations

The estimates are model outputs, not investment advice and not predictions of realized profit. Funding can change, fills can be partial, legs can execute at different times, and venue events can interrupt an otherwise matched route. The reference basis assumes that recent settled history is representative of the coming hold. A regime shift may be flagged without changing the number, so the presence of a flag does not convert the reference into a regime-aware forecast. For venues without settlement history, the first-stage current-rate basis has a known upward bias in magnitude under mean reversion, but its size is not estimated. These limitations affect the reference and conservative ends differently and cannot be summarized by a single confidence coefficient.

Execution costs are measured on current books. No model governs how depth, spread, or fillability evolves over the holding period, and the conservative exit candidate is not a worst-case bound. A fresh book reduces stale-input risk but does not guarantee a future fill. Venues without an order book, including request-for-quote systems with indicative prices only, cannot receive a walked cost estimate. Transfer, withdrawal, borrow, and currency-conversion costs remain excluded where unmeasured, and excluded costs are named. Spread reversion requires history and is omitted when it cannot be formed. Premium research forecasts approximate impact prices by the top of book; that error is not reported, and their bands ignore dead-zone and cap clamping.

Liquidation risk is approximate and short horizon. The model covers isolated margin and a single position, uses single-day rather than cumulative adverse movement, and excludes cross-margin accounts, accrued funding, liquidation fees, mark-price details, and auto-deleveraging rank. Many venues lack usable public tier information, so a full route comparison is available only for some routes. A cushion above one does not imply safety, and a cushion below one is a historical stress observation rather than a liquidation probability. Sample-quantile conventions also vary across statistical packages (Hyndman & Fan, 1996), which can change small-sample evaluation details by an order statistic; the convention used internally is not restated here.

The most important evidentiary limitation is the absence of an out-of-sample evaluation of the headline estimate. The only realized-versus-extrapolated results concern a different current-rate predictor over one short window and one regime, with mostly top-of-book cost proxies. Proper scoring, interval tests, split-conformal calibration, headline outcome evaluation, PBO, and CSCV remain planned where identified as such. No performance generalization is warranted from the replay. Finally, restricting undisclosed production parameters limits independent replication by design; the paper offers an auditable conceptual and accounting framework rather than a complete implementation specification.

11. Conclusion

A defensible net-carry estimate begins with cash accounting rather than annualized rate subtraction. Settled history supplies a robust reference basis; settlement grids determine when cash can arrive; per-settlement caps preserve venue scale; and depth-walked, common-quantity fills connect cost to trade size. Reference and conservative ends make the principal assumptions inspectable without assigning unsupported probabilities. Missing essential inputs lead to unavailable estimates, and unmeasured costs remain visible. The nominal cost-coverage horizon answers a restricted flat-gap arithmetic question, while the liquidation cushion compares tier-adjusted isolated-margin distance with the asset’s own recent adverse daily move.

The framework’s central discipline is semantic separation. A measured current cost is not a forecast fill, a conservative end is not a confidence bound, a nominal horizon is not capital recovery, and a historical cushion is not survival probability. The limited replay demonstrates why current-rate extrapolation deserves caution, but it does not validate the estimator

presented here. Future evidence must come from frozen decision outcomes, official settled prints, unresolved-case accounting, proper probabilistic scores where applicable, and explicit selection and overfitting diagnostics. Until those results exist, the appropriate claim is methodological: the framework states what a net-carry number means, what it includes, and what remains unknown.

Declarations

Funding. Research and writing were supported by humarb internal resources. No external grant funding was received.

Data availability. The paper describes public input classes, including officially settled funding prints, public venue rules, order-book snapshots, fee and margin schedules, and closed daily candles. The limited replay summary covers 14 days in October 2026; its four aggregate ratios are reported in the text and were not independently re-estimated for this paper. Production implementation details, undisclosed statistical parameters, and internal operating rules are not public. No new route-level dataset accompanies this preprint.

Code availability. No production code is disclosed. The paper contains mathematical and schematic material only.

Non-investment-advice statement. This paper is for research and methodological information only. It does not provide investment, legal, accounting, or tax advice; does not recommend any trade or leverage; and does not promise income, execution, or protection from liquidation.

Generative-AI assistance. Generative AI assisted with language drafting, structural editing, bibliography-format checks, and schematic preparation under human direction. The methodological content is constrained to the source materials supplied for this work. Hang Liu is responsible for review, numerical provenance, confidentiality boundaries, and the final submission.

License. The author intends to make the preprint available under CC BY 4.0, subject to the repository’s submission terms.

References

Almgren, R., & Chriss, N. (2001). Optimal execution of portfolio transactions. *The Journal of Risk*, *3*(2), 5–39. https://doi.org/10.21314/JOR.2001.041

Bailey, D. H., Borwein, J. M., López de Prado, M., & Zhu, Q. J. (2017). The probability of backtest overfitting. *The Journal of Computational Finance*, *20*(4), 39–69. https://doi.org/10.21314/JCF.2016.322

Bailey, D. H., & López de Prado, M. (2014). The deflated Sharpe ratio: Correcting for selection bias, backtest overfitting, and non-normality. *The Journal of Portfolio Management*, *40*(5), 94–107. https://doi.org/10.3905/jpm.2014.40.5.094

Brier, G. W. (1950). Verification of forecasts expressed in terms of probability. *Monthly Weather Review*, *78*(1), 1–3. https://doi.org/10.1175/1520-0493(1950)078%3C0001:VOFEIT%3E2.0.CO;2

Christoffersen, P. F. (1998). Evaluating interval forecasts. *International Economic Review*, *39*(4), 841–862. https://doi.org/10.2307/2527341

Dai, M., Li, L., & Yang, C. (2025). *Arbitrage in perpetual contracts*. SSRN. https://doi.org/10.2139/ssrn.5262988

Gatev, E., Goetzmann, W. N., & Rouwenhorst, K. G. (2006). Pairs trading: Performance of a relative-value arbitrage rule. *The Review of Financial Studies*, *19*(3), 797–827. https://doi.org/10.1093/rfs/hhj020

Gneiting, T., & Raftery, A. E. (2007). Strictly proper scoring rules, prediction, and estimation. *Journal of the American Statistical Association*, *102*(477), 359–378. https://doi.org/10.1198/016214506000001437

Hampel, F. R. (1974). The influence curve and its role in robust estimation. *Journal of the American Statistical Association*, *69*(346), 383–393. https://doi.org/10.1080/01621459.1974.10482962

Harvey, D., Leybourne, S., & Newbold, P. (1997). Testing the equality of prediction mean squared errors. *International Journal of Forecasting*, *13*(2), 281–291. https://doi.org/10.1016/S0169-2070(96)00719-4

He, S., Manela, A., Ross, O., & von Wachter, V. (2022). *Fundamentals of perpetual futures*. SSRN. https://doi.org/10.2139/ssrn.4301150

humarb Research. (2026a). *Measuring executable net carry in cross-venue perpetual funding-rate arbitrage* [Companion study]. https://humarb.ai/research/executable-net-carry/ https://doi.org/10.5281/zenodo.23868361

Hyndman, R. J., & Fan, Y. (1996). Sample quantiles in statistical packages. *The American Statistician*, *50*(4), 361–365. https://doi.org/10.1080/00031305.1996.10473566

Kupiec, P. H. (1995). Techniques for verifying the accuracy of risk measurement models. *The Journal of Derivatives*, *3*(2), 73–84. https://doi.org/10.3905/jod.1995.407942

Leys, C., Ley, C., Klein, O., Bernard, P., & Licata, L. (2013). Detecting outliers: Do not use standard deviation around the mean, use absolute deviation around the median. *Journal of Experimental Social Psychology*, *49*(4), 764–766. https://doi.org/10.1016/j.jesp.2013.03.013

Makarov, I., & Schoar, A. (2020). Trading and arbitrage in cryptocurrency markets. *Journal of Financial Economics*, *135*(2), 293–319. https://doi.org/10.1016/j.jfineco.2019.07.001

Parkinson, M. (1980). The extreme value method for estimating the variance of the rate of return. *The Journal of Business*, *53*(1), 61–65. https://doi.org/10.1086/296071

Perold, A. F. (1988). The implementation shortfall: Paper versus reality. *The Journal of Portfolio Management*, *14*(3), 4–9. https://doi.org/10.3905/jpm.1988.409150

Smith, J. E., & Winkler, R. L. (2006). The optimizer’s curse: Skepticism and postdecision surprise in decision analysis. *Management Science*, *52*(3), 311–322. https://doi.org/10.1287/mnsc.1050.0451

Appendix A. Symbols

Symbol	Meaning	Unit
$j \in \{\mathrm{L}, \mathrm{S}\}$	Long or short leg of a route	—
s_j	Cash sign: short $+1$, long -1	—
A	Per-side notional	USD
V	Position notional in a venue-rule expression	USD
r, r_j	Per-settlement funding rate	fraction/settlement
H	Holding period	hours
τ_j	Funding interval of leg j	hours
y_j	Funding rate normalized to one hour	fraction/hour
$c_j^{\mathrm{lo}}, c_j^{\mathrm{hi}}$	Per-settlement lower and upper caps	fraction/settlement
m_j	Eligible settlements inside the hold	count
G	Funding cash over the hold	USD
F	Taker fees for four fills	USD

Symbol	Meaning	Unit
$K^{\mathrm{ent}}, K^{\mathrm{ext}}$	Entry and exit crossing costs	USD
K^{oth}	Other costs, only when measured	USD
B	Spread-reversion term	USD
$\Pi^{\mathrm{ref}}, \Pi^{\mathrm{con}}$	Reference and conservative estimated net income	USD
APR	Simple net annualized rate	%/year
$P, \bar{P}$	Instantaneous and interval-average premium index	fraction
I, δ	Interest component and dead-zone half-width	fraction
g	Current funding gap, short minus long	fraction/hour
T_{cc}	Nominal cost-coverage horizon	hours
ϱ	Realized-to-extrapolated funding ratio	—
L	User leverage	×
μ, D	Tier maintenance-margin rate and deduction	fraction; USD

Symbol	Meaning	Unit
d_j	Approximate liquidation distance	fraction of entry price
w_j	Largest recent adverse single-day move	fraction of day-open price
k	Route liquidation-cushion ratio	—
Λ	Settlement-block period	hours
d	Net funding cash per settlement block	USD
$S(t)$	Cumulative funding cash path	USD
γ_j	Signed funding cash at one settlement	USD
β	Relative cross-venue mid-price spread	fraction
q	Common base quantity across four fills	base asset
$\bar{p}^{\mathrm{ent}}$	Entry VWAP	USD/base asset
σ_j	Single-side entry slippage	bp
$N_j^{\mathrm{en}}, N_j^{\mathrm{ex}}$	Entry and exit cash notionals	USD

Symbol	Meaning	Unit
C	One-off round-trip cost	USD
T^*	Replay-sensitive corrected-horizon reference	hours
E, Q	Entry price and position quantity	USD/base asset; base asset
μ_{eff}	Effective maintenance-margin rate	fraction
$L_{\max}$	Maximum leverage of a tier	×
o, h, l	Daily open, high, and low prices	USD/base asset
$\mathcal{F}$	Predictive cumulative distribution function	—
$\chi\varepsilon$	Forecast probability of event $\mathcal{E}$	probability
ℓ_α	Predictive quantile at level α	forecast units
n_b	Number of equal blocks in CSCV	count

Appendix B. Derivation Notes

For the settlement grid, each leg contributes Λ/τ_j payments in one block. Multiplying by signed cash γ_j and summing yields d. Repeating the same candidate settlement amounts gives $S(t+n\Lambda) = S(t) + nd$. If $d \le 0$, every later block is weakly lower than the corresponding point in the preceding block, proving the early-stop statement. This derivation concerns timing under fixed candidate cash amounts; it does not assume that realized future rates are constant or known.

For liquidation, a long has margin plus profit-and-loss $QE/L + Q(p - E)$, while a short has $QE/L + Q(E - p)$. Equating each expression to $Qp\mu - D$ and solving for p gives equation (12). Dividing the absolute entry-to-liquidation price change by E gives equation (13). The maintenance deduction appears with a positive sign in both distances because a larger D lowers the requirement. These are algebraic consequences of the stated isolated-margin approximation, not a reproduction of every venue’s liquidation engine.

For the common-quantity walk, entry and exit notionals are cash amounts for the same q. Adding the two legs' entry-minus-exit cash yields equation (10), so the mid-price benchmark terms cancel from total crossing cost. This cancellation does not remove market risk or price changes during the hold; it concerns only the execution-cost decomposition measured against each leg's own current mid.

The separate **Number and Citation Register** supplied with the submission maps all substantive numerical statements and bibliography entries to their source sections and DOI metadata. Formula numbers, section numbers, author years, and JEL codes are structural metadata rather than empirical observations.