Measuring Executable Net Carry in Cross-Venue Perpetual Funding-Rate Arbitrage

humarb Research · November 2026 · humarb.ai

Abstract

A quoted funding-rate differential is not, by itself, a measure of the cash flow available from a cross-venue position. This paper studies executable net carry as a measurement problem: expected funding must be placed on a robust multi-day basis, evaluated at a stated position size, and reduced by fees and depth-dependent entry and exit costs before routes can be compared. Using only the methodology and reported measurements in the humarb Methodology Pack, hereafter the humarb framework, the paper formalizes a route-level framework that combines contract-identity controls, seven-day medians of paired daily funding differences, separate order-book walks for entry and exit, book-specific freshness tests, and explicit treatment of unmeasured inputs. In the full-list snapshot published at 2026-11-01 19:19 UTC, 509 of 3,164 routes had a reference estimate; among those 509 routes, 462 (90.8%) had negative estimated net income after all measured costs at USD 5,000 per side over a seven-day hold. In the separate 2026-11-01 12:00–13:05 UTC window of 343 consecutive publishes, median net-positive route counts on the reference side were 71, 69, and 106 for the reported default, USD 5K, and USD 2K categories, respectively, whereas the conservative side reported zero for all tiers. These are model-based snapshot measurements, not realized trading outcomes. The evidence shows how cost and data-quality controls alter the economic interpretation of apparent funding spreads.

Keywords: perpetual futures; funding rates; net carry; transaction costs; order-book depth; market microstructure

1 Introduction

1.1 From quoted funding spreads to executable net carry

Perpetual futures replace terminal convergence at expiry with an ongoing funding mechanism that links the derivative to its reference market. The original analysis of Shiller (1993) motivates perpetual contracts as instruments whose cash settlement depends on an index, while Ackerer, Hugonnier, and Jermann (2026) provide a modern treatment of perpetual-futures pricing under alternative funding specifications. In a cross-venue funding position, the relevant object is not the perpetual contract in isolation but a route consisting of two legs, normally a long and a short exposure to the same underlying asset. The sign and size of funding cash flows depend on each venue's rate, accrual convention, and settlement schedule. A comparison of raw prints can therefore combine quantities that are not economically commensurate. The problem becomes more acute when a route is intended to be held across several settlements: the latest observation may be transient, while the position's funding outcome is accumulated over time.

Executable net carry adds a second distinction. Even a correctly normalized funding differential is a gross quantity until the position's entry and exit are financed. Both legs incur

fees, and both must be executed against available liquidity rather than against an abstract mid-price. The humarb framework defines net income as expected funding cash flow over the holding period minus fees, entry cost, and exit cost, with a seven-day default holding period and a stated per-side notional (humarb, 2026, §3.1). This definition changes the unit of analysis from an observable quote disparity to a modelled cash-flow residual at a specified size. It also creates a conservative interpretation of missing information: if depth or another required input cannot be measured, the framework reports that limitation instead of assigning a zero cost or manufacturing a synthetic estimate. The approach is thus concerned with what can be measured consistently, not with labelling a position as certain or continuously attainable.

1.2 Research questions and contribution

This paper addresses three questions. First, after funding expectations are robustified and all measured costs are deducted, how many reference-estimated routes remain net-positive in the reported snapshots? Second, how does the classification change when each leg is evaluated on the conservative basis specified by the humarb framework, including adverse-rate selection and peer-median replacement? Third, how do depth availability, book freshness, and contract-identity requirements affect estimate coverage? These questions are descriptive rather than predictive. They concern the internal economic meaning of a screener's output at documented times and position sizes. They also separate two notions of coverage: the share of published routes for which an estimate exists, and the share of those estimates whose net income is negative. Confusing these denominators would materially change the interpretation of the reported evidence.

The contribution is a disciplined reconstruction of a measurement framework and its snapshot results, situated in the literature on perpetual contracts, carry, execution, and crypto-derivative microstructure. The paper does not introduce a new trading strategy or an independent return sample. Instead, it makes explicit why seven-day paired medians, walked depth, fresh-book constraints, and named exclusions belong in the definition of net carry. The empirical section preserves the humarb framework's separate measurement windows: a single full-list publish, a median over six default-list publishes, a one-hour engine window, and an earlier prototype audit. No route-level observations, parameter estimates, or cost values beyond those disclosed in the humarb framework are added. Where a quantity required for a broader statistical analysis is absent—such as route-level gross-to-net bridges or conservative-estimate denominators—it is reported as **not disclosed**.

2 Related Research

2.1 Perpetual futures and funding-rate mechanisms

The theoretical lineage of perpetual futures emphasizes the relationship between an index, derivative valuation, and cash settlement. Shiller (1993) examines repeated-measures indices and perpetual futures as mechanisms for markets in assets whose values are difficult to observe continuously. Ackerer et al. (2026) subsequently develop pricing results for perpetual futures under linear and inverse contract designs and funding rules. This literature matters for cross-venue carry because funding is not merely an auxiliary fee appended to a futures price; it is part of the contract's economic design. Consequently, venue-specific accrual and settlement conventions must be normalized before two rates are differenced. The humarb framework's distinction between discrete settlement histories and continuously accruing venues follows the

same conceptual requirement, although its exact seven-day and hourly-snapshot procedures are source-specific methodological choices rather than results imported from the pricing literature.

The time-series dimension is equally important. Inan (2026) studies the predictability of perpetual-futures funding rates, providing a direct motivation for treating funding as a changing empirical object rather than a fixed yield. A single current print may differ from the rate that governs a multi-day holding period, and an estimate based on recent history should therefore be explicit about its temporal construction. The humarb framework answers this problem with daily pairing and a median over seven days, supplemented by shorter-history labels and snapshot procedures for particular venue types (humarb, 2026, §3.3). The relevant contribution here is not a claim that a median forecasts future funding with a quantified accuracy; the humarb framework supplies no such accuracy estimate. Rather, the median limits the influence of a single spike day and aligns the two legs by calendar day before their difference is summarized.

2.2 Basis, carry, and arbitrage frictions

Carry is commonly defined through the return component that would be earned if market conditions remained unchanged, and Koijen, Moskowitz, Pedersen, and Vrugt (2018) document its relevance across asset classes. In crypto markets, Schmeling, Schrimpf, and Todorov (2026) analyze futures carry and its relation to leveraged positioning and arbitrage capital. These perspectives help distinguish a quoted basis or funding differential from the return available after financing, positioning, and liquidation constraints. A cross-venue perpetual route is related to, but not identical with, conventional cash-and-carry: it may involve two perpetuals rather than a spot asset and a dated future, and its cash flows depend on two funding schedules. The appropriate comparison is therefore conceptual. General carry research supplies the distinction between gross pricing relationships and implementable positions; it does not supply the numerical cost parameters used in this paper.

Research on crypto derivatives further shows why arbitrage frictions belong inside the measurement. Alexander, Chen, Deng, and Wang (2024) examine arbitrage opportunities and efficiency tests in crypto derivatives, while De Blasis and Webb (2022) connect Bitcoin-futures arbitrage to contract design and market structure. Lee, El Meslmani, and Switzer (2020) study pricing efficiency and arbitrage between Bitcoin spot and futures markets, and Augustin, Rubtsov, and Shin (2023) analyze how futures-market introduction affects spot markets. Together, these studies establish that price alignment, market structure, and arbitrage activity are interdependent. The humarb framework's quote-currency matching and depth-walked costs can be read as route-level controls within this broader setting. Its earlier prototype audit also demonstrates that an apparent spread may be an identity error rather than an economic discrepancy, a point that cannot be addressed by a funding-rate model alone.

2.3 Transaction costs and order-book depth

Execution research provides the microeconomic basis for rejecting top-of-book pricing as a universal execution assumption. Almgren and Chriss (2000) formalize the trade-off between execution cost and volatility risk in portfolio transactions. Bouchaud, Mézard, and Potters (2002) document statistical properties of order books, and Cont, Kukanov, and Stoikov (2014) relate price impact to order-book events. Fruth, Schöneborn, and Urusov (2014) and Ackermann, Kruse, and Urusov (2021) further model execution when liquidity varies over time. The common implication is that available quantity and price are joint attributes of an execution. A small trade at the best quote and a larger trade consuming several levels need not have the same average

price. The humarb framework's walked entry and exit costs operationalize this principle at a stated per-side amount without importing any external liquidity parameter.

Crypto-market evidence reinforces the need to measure costs directly. Dyhrberg, Foley, and Svec (2018) analyze Bitcoin liquidity and transaction costs, making clear that market access and investability depend on more than an observed last price. For a two-leg funding route, this issue is compounded: entry and exit each require execution on both sides, and insufficient depth on either leg can prevent a consistent estimate. The humarb framework therefore takes the minimum of entry and exit fill quantities, refuses extrapolation when a book cannot fill the required amount, and rejects crossed or empty books (humarb, 2026, §3.4). These choices deliberately trade coverage for coherence. They do not establish that walked prices equal subsequently realized fills; partial fills and inter-leg delays remain separate implementation risks identified in the humarb framework's limitations.

2.4 Cryptocurrency-derivatives market microstructure

Studies of crypto-derivative microstructure provide the final research context. Alexander, Choi, Park, and Sohn (2020b) examine price discovery, informational efficiency, and hedging effectiveness in BitMEX Bitcoin derivatives. Alexander, Choi, Massie, and Sohn (2020a) analyze ether spot and derivative markets, showing the importance of interactions between related trading venues. Almeida and Gonçalves (2024) synthesize a broader cryptocurrency-microstructure literature spanning efficiency, liquidity, and price discovery. This body of work supports treating quote formation and market structure as central to derivative valuation rather than as secondary implementation details. In a cross-venue setting, asynchronous or unconfirmed book states can create apparent discrepancies that do not represent simultaneously available prices.

The gap addressed by the present paper is narrower and primarily measurement-oriented. Existing research supplies theories of perpetual pricing, carry, execution, and price discovery, but the humarb framework combines them into a single route-level reporting rule: a robust paired funding basis, two-sided walked execution costs, separate reference and conservative estimates, book-specific freshness checks, and explicit missing-data labels. The framework also recognizes that coverage is endogenous to data quality. Refusing stale books or incomplete depth may reduce the number of estimates while improving their interpretability. The empirical question is therefore not simply whether apparent spreads exist, but how many routes survive a sequence of identity, liquidity, temporal, and cost controls at a documented instant. This is the sense in which "executable" is used here: measurable under the stated rules, not guaranteed to be realized.

3 Data and Methods

3.1 Data classes and measurement architecture

All methodological inputs and empirical values in this paper come from the humarb framework. The source describes funding-rate records, realized funding-settlement histories, hourly rate snapshots for continuously accruing venues, bids and asks with exchange and receive timestamps, contract information, maintenance-margin tiers, and daily UTC candles (humarb, 2026, §2.2). Each record carries a record class, the exchange's own timestamp where available, a receive timestamp, and quality flags. The exchange timestamp remains null when the venue supplies none; receive time is not substituted for it. Values are stored as decimal strings with explicit

units, while missing values carry reason codes rather than being converted to zero. Coverage differs across venues and data classes. The humarb framework does not disclose a complete route-level dataset, a complete asset inventory, or a fully enumerated rule for membership in the user-facing default list; these items are therefore **not disclosed** here.

The empirical evidence comprises four distinct measurement architectures, summarized in Table 1. Table A is a single full-list publish and supports statements about reference-estimated routes at the default amount. Table B is a median across six consecutive default-list publishes and supports statements about estimate coverage and missing depth. Table C contains 343 consecutive publishes without gaps during a stated one-hour window and supports median route counts under reference and conservative bases. Table D is an earlier prototype audit concerned with apparent positive-spread value, not net income. These observations are neither pooled nor treated as a common sample. In particular, the full-list denominator cannot be attached to the one-hour net-positive counts, and the prototype's quote-mixing share cannot be compared directly with the full list's negative-income share.

Table 1. Measurement windows and estimands reported in the humarb framework.

Source table	Measurement time and aggregation	Position or output basis	Quantity supported
A	Single publish, 2026-11-01 19:19 UTC	Reference estimate; USD 5,000 per side; seven-day hold	Full-list coverage and sign of net income after measured costs
B	Median of six consecutive publishes, 2026-11-01 14:39 UTC	Default list shown to users	Routes with estimates and principal disclosed missing-data reason
C	343 consecutive publishes, no gaps, 2026-11-01 12:00–13:05 UTC	Median route counts; reference and conservative sides	Net-positive route counts by reported category
D	Earlier prototype audit, 2026-10-14	Apparent positive-spread value	Share associated with USDC–USDT quote-currency mixing

Source: (humarb, 2026, §4A–D). Each row retains its own time and measurement definition.

3.2 Contract identity and normalization

A route is valid only if its legs refer to compatible economic instruments. The humarb framework uses contract identifiers, multipliers, quote currencies, and listing information for identity matching, while a contract-identity check can veto a route (humarb, 2026, §§2.2, 3.6). This control addresses superficially similar names that conceal different multiplier prefixes or separators. It also prevents USDC-quoted and USDT-quoted contracts from being mixed as though their prices and funding rates were directly interchangeable. Funding intervals are normalized before comparison because the source identifies 1-, 4-, and 8-hour intervals among the covered contracts. These are measurement prerequisites: an apparent rate differential that fails identity or unit consistency is not interpreted as carry.

Normalization does not eliminate all contractual heterogeneity. Discrete-settlement venues provide realized per-settlement payments, continuously accruing venues require hourly snapshots, and venues without settlement history initially require a current-rate approximation. The humarb framework preserves these bases as labels rather than concealing them behind a common numerical field. Unknown funding-accrual modes produce no number, and pre-market or session constraints that make a leg untradable during the holding window are flagged. This

treatment is important for interpretation: a route with no estimate is not evidence of zero carry, just as an excluded cost is not a cost measured as zero. The output distinguishes an unfavourable measured value from an unavailable measurement.

3.3 Headline net-carry measure

The humarb framework's headline quantity is estimated net income over a stated holding period, accompanied by net APR on a per-side-notional basis. The default amount is USD 5,000 per side, or USD 10,000 across the two legs, and a USD 2,000-per-side tier is also computed. The default holding period is seven days, equivalent to 168 hours. Routes are ranked by estimated net income, but a route for which no estimate can be computed remains in the list, ranked last and labelled accordingly; it is not assigned a simulated value (humarb, 2026, §3.1). The resulting measure is therefore conditional on amount, holding period, rate basis, and available execution information. Net APR annualizes the same estimate on per-side notional; it does not transform the snapshot into a realized annual return series.

$$\begin{aligned}\text{Net income} = {} & \text{expected funding cash flow over the holding period} \\ & - \text{fees} - \text{entry cost} - \text{exit cost}.\end{aligned} \tag{1}$$

The cost definition includes taker fees from published venue fee tiers and separate depth-walked entry and exit costs. Transfer and withdrawal costs between venues are excluded and named as excluded rather than priced as zero. The humarb framework also identifies borrow cost and certain currency-conversion costs as unpriced in its limitations. The phrase "after all measured costs" is consequently more precise than "after all costs": unmeasured or excluded items can still affect realized cash flow. This distinction also prevents an unsupported reconciliation between gross and net values. The humarb framework does not disclose route-level expected funding, fee, entry-cost, and exit-cost components, so a numerical decomposition of the 90.8% negative-income result into individual cost contributions is **not disclosed**.

3.4 Reference estimate: a seven-day robust rate basis

For venues with settlement history, the humarb framework buckets settlements by UTC day and forms a reference hourly rate from recent daily rates. At route level, the legs are paired by day, and the median is taken over the daily short-minus-long differences during the last seven days. This is not the difference between two independently calculated leg medians. If fewer than four same-day pairs exist under the current setting, the estimate is refused. A leg with less than 24 hours of settled history is not estimated; a leg with one to six days is estimated on a shorter basis and labelled new. The humarb framework also defines a regime-shift flag when the mean of recent paired days has the opposite sign to, or differs by more than half from, the seven-day median. The flag does not alter the estimate, and the number of recent paired days used in that mean is **not disclosed** (humarb, 2026, §3.3).

The rationale for the median is robustness to a single spike day, not a demonstrated forecasting optimum. Daily pairing is equally consequential because it preserves the route's daily short-minus-long cash-flow geometry before temporal aggregation. A median of paired differences can differ from a difference of marginal medians when daily observations are asynchronous or differently distributed. The seven-day horizon also matches the default holding-period scale, although the humarb framework does not claim equality between historical and future funding. Accordingly, this paper interprets the reference estimate as a robust historical basis for a

screening number. It does not assign a confidence interval, hit rate, or out-of-sample predictive coefficient, none of which is disclosed in the source.

Continuously accruing venues require a different input but retain the same robust-summary principle. The humarb framework uses the median of daily means of hourly rate snapshots over six previous complete UTC days plus the current day when that day contains at least 20 hourly observations; a day with fewer than 20 observations is excluded. For venues publishing no settlement history, the initial estimate uses the current hourly rate, is labelled current-rate and approximate, and ranks after seven-day-settled routes. In parallel, the humarb framework records observed rates and switches to the median of seven daily means after seven complete days exist, each with at least 20 hourly observations. These procedures do not make snapshots equivalent to realized payments. They document precisely which observable basis supports each estimate and when a weaker basis remains in use.

3.5 Conservative estimate and peer-median replacement

The conservative estimate evaluates each leg on the less favourable of several candidate rates. One candidate is the current rate, but a favourable outlier is replaced by the peer-venue median for the same asset when it deviates from that median by at least both of the humarb framework's stated thresholds: 0.01% per hour or 50% of the median magnitude, whichever is larger, and three times the median absolute deviation. If no peers exist, the conservative side is unavailable; it never falls back to the raw current rate. The humarb framework does not enumerate the complete candidate set, so that set is **not disclosed**. When a paying leg touches a venue funding cap and the venue may shorten its interval, the conservative side counts the additional settlements (humarb, 2026, §3.2). The estimate is thus a scenario construction with explicit replacement rules, not a statistical lower confidence bound.

The two estimate bases also differ in exit slippage. The humarb framework assigns an empirical median of historical exit slippage to the reference side and a high percentile to the conservative side; the exact high percentile is **not disclosed**. Reporting both ends separates a recent-history calculation from a less favourable rate-and-exit construction without merging them into a single point estimate. The difference cannot be attributed solely to funding-rate replacement, because exit-slippage assumptions change as well. A paired route-level decomposition is not provided. The empirical comparison in this paper therefore concerns counts under the complete reference and conservative definitions, not an isolated causal estimate of any one conservative component.

3.6 Depth-walked entry and exit costs

Entry cost is calculated by walking the book to fill the stated per-side notional: the long leg consumes asks and the short leg consumes bids. Exit is walked separately for the same quantity rather than inferred from the entry walk. Taker fees are then obtained from the venue's published fee tier. The fill quantity is the minimum of entry and exit fill. If the book cannot fill the required amount, cost is marked unmeasured rather than extrapolated, and a smaller default walk is not substituted for a missing tier at a larger amount. Crossed or empty books are rejected (humarb, 2026, §3.4). These rules make the estimate size-specific and prevent an apparently liquid top quote from standing in for unavailable depth.

Walking successive levels captures the difference between a marginal quote and an average execution price for the specified amount. It also explains why a flat fee rate or quoted top-of-

book spread is insufficient as a complete execution model. The humarb framework separately defines an entry spread from mid-prices, $b = (\text{short mid} - \text{long mid})/\text{mean mid}$, but uses it only when both books are no older than 60 seconds and each top-of-book width is below 1%. A wide or stale book would contradict the walked entry associated with the same row. The mid-price spread is therefore a constrained diagnostic, not a replacement for the walked cost. Realized slippage may still differ because the humarb framework does not model every aspect of order handling, partial execution, or inter-leg timing.

3.7 Freshness and reused-book guards

A rate older than 60 seconds is marked stale for estimation. For books, freshness is specific to the individual book: the top is fresh only when the latest message for that book is at most 10 seconds old, using the earlier of exchange time and receive time. A live connection or heartbeat does not refresh the book's clock. The reason is evidential rather than infrastructural: connection health does not establish that the venue has confirmed an unchanged book. The humarb framework explicitly prefers publishing fewer estimates to publishing an unconfirmed one (humarb, 2026, §§2.3, 3.5). This choice can reduce coverage in thin markets where valid quotes update infrequently, but it prevents an old price from improving a route's estimated economics without current confirmation.

A reused book must also remain consistent with the current market. When two walks conducted at different times appear to cross, the humarb framework checks the pair against the legs' current best bids and asks and drops the pair if the top-of-book spread and walk spread do not agree within a tolerance; the numerical tolerance is **not disclosed**. The motivating observation was an apparent 40+ bp opportunity produced by walks taken 37 seconds apart in a moving market while the same venues quoted within 1 bp. This example is a methodological diagnostic, not a separate performance observation. Together, the two guards ensure that a cached book that looks better than the current top cannot improve an estimate merely because it remains available in memory.

3.8 Liquidation geometry and unmeasured risk inputs

The humarb framework reports liquidation distance separately from net income. For an isolated, single-position leg with leverage L and maintenance-margin rate r in the tier containing the per-side amount, the fractional price distance is $(1/L - r)/(1 - r)$ for a long and $(1/L - r)/(1 + r)$ for a short. A spot leg has no liquidation in this model. The volatility comparison uses the largest single-day range over the last 30 closed UTC days, $\max((\text{high} - \text{low})/\text{open})$, with directional versions for upward and downward moves. The nearer leg's distance is compared with that range. If a single-day move exceeds 100% of price, the display states that any leverage can be hit. Missing tier or candle data produces an unmeasured label rather than a safety conclusion (humarb, 2026, §3.7).

The model's scope is deliberately limited. A maximum one-day range is not a cumulative seven-day drawdown, and isolated-margin geometry does not describe cross-margin liquidation. The humarb framework bounds the displayed worst-case loss on one leg by the margin posted on that leg under the isolated-margin convention; it does not aggregate account-level loss across venues. These risk measures are not converted into an additional numerical deduction in the net-income equation. The empirical results should therefore be read as carry estimates accompanied by separately measured risk indicators, not as a unified expected utility or risk-adjusted return. A complete liquidation conclusion was available for only about 10–14% of routes at the time of

writing, a limitation addressed again below.

4 Empirical Results

4.1 Full-list net income after measured costs

Figure 1 reports the full published list at a single publish, 2026-11-01 19:19 UTC. Of 3,164 published routes, 509 had a reference estimate, corresponding to the 16% coverage reported by the humarb framework. Among those 509 reference-estimated routes, 462 had negative net income after all measured costs, or 90.8%. The estimate basis was USD 5,000 per side with a seven-day holding period. In the subset for which both legs' current rates were available, 282 of 440 routes—64.1%—had at least one leg currently paying funding, where paying means that the long leg's current rate was positive or the short leg's current rate was negative (humarb, 2026, §4A). These measurements describe distinct subsets and should not be combined into a single conditional probability.

The 90.8% result is the central descriptive finding: within the reference-estimated subset, net income after fees and depth-walked entry and exit was overwhelmingly negative at the stated snapshot. Its denominator is 509, not all 3,164 published routes and not a disclosed set of routes with positive gross funding spreads. The humarb framework also does not disclose a separate count of exactly zero net income. Accordingly, Figure 1 plots only the disclosed negative class against the full reference-estimated track; it does not label every remaining route as positive. A positive/negative pie chart would therefore introduce an undisclosed classification.

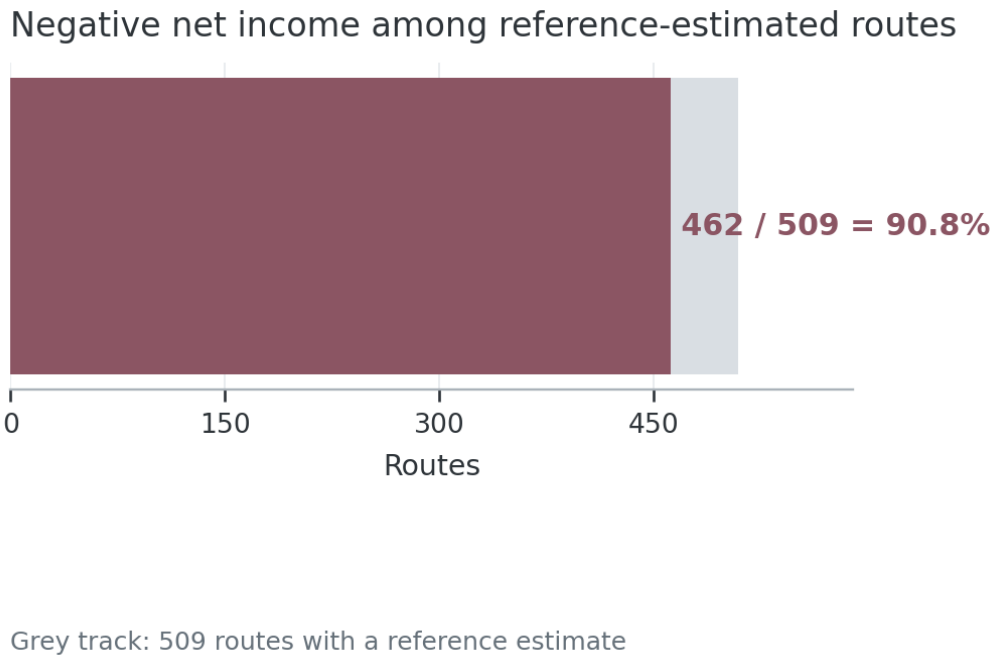


Figure 1. Negative net income among routes with a reference estimate. Source: (humarb, 2026, §4A); single publish at 2026-11-01 19:19 UTC; USD 5,000 per side, seven-day hold. Only the disclosed negative class is coloured; zero-valued outcomes are not separately disclosed.

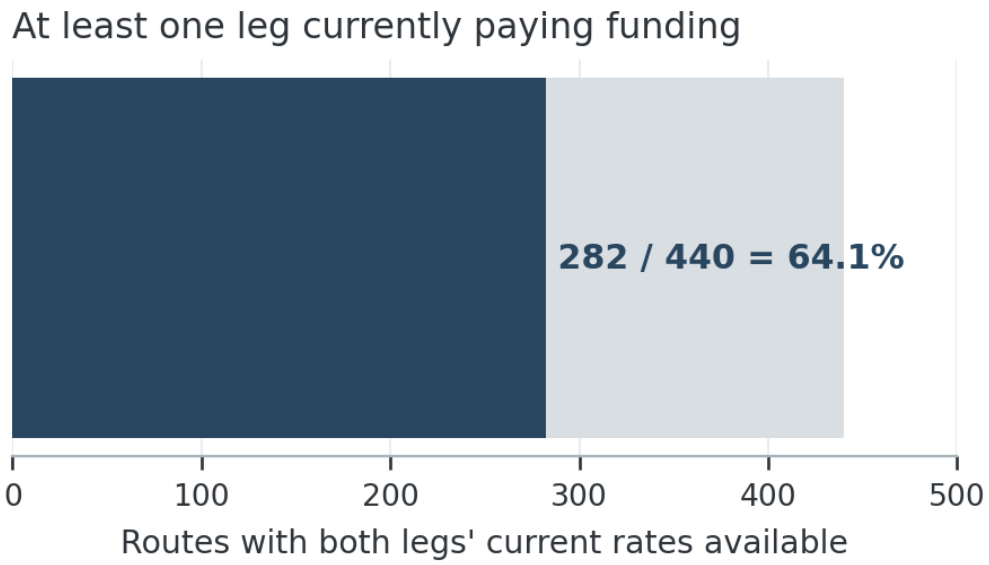


Figure 2. Share of routes with at least one leg currently paying funding, among routes with both legs' current rates available. Source: (humarb, 2026, §4A); single publish at 2026-11-01 19:19 UTC. The denominator is 440, distinct from the 509 reference-estimated routes in Figure 1.

4.2 Reference and conservative estimates

The one-hour window provides a different estimand: net-positive route counts across repeated publishes rather than the sign distribution within one full-list publish. Source Table C covers 2026-11-01 12:00–13:05 UTC and contains 343 consecutive publishes with no gaps. On the reference side, the median net-positive counts were 71, 69, and 106 for the categories labelled default, USD 5K tier, and USD 2K tier, respectively. On the conservative side, the humarb framework reports zero net-positive routes for all tiers (humarb, 2026, §4C). The category labels are reproduced exactly; the source does not provide a further reconciliation of the default label with the separately listed USD 5K tier.

Figure 3 makes the contrast visible without converting the counts to percentages. The relevant denominators are **not disclosed**, and the window's population of eligible routes cannot be inferred from Table A's 509 reference estimates because the measurement windows and list definitions differ. The contrast is consistent with the combined effect of less favourable candidate rates, peer-median replacement, and more conservative exit slippage, but the humarb framework supplies no component-wise bridge. The zero conservative counts should therefore be interpreted strictly: no net-positive routes were reported on that basis during this window. They are not evidence that positive net income is impossible at other times or under other disclosed inputs.

4.3 Estimate coverage and missing depth

The default-list measurement concerns whether an estimate exists, not whether it is positive. At 2026-11-01 14:39 UTC, using the median of six consecutive publishes, the default list contained 418 routes, of which 234 had an estimate, or 56%. The principal disclosed reason that a route lacked an estimate was order-book depth not yet measured, accounting for 165 routes (humarb, 2026, §4B). The humarb framework's limitations add that these gaps were concentrated in several named venues and that work to close them was scheduled but incomplete. This paper does not infer that the missing routes would have been profitable or unprofitable had depth

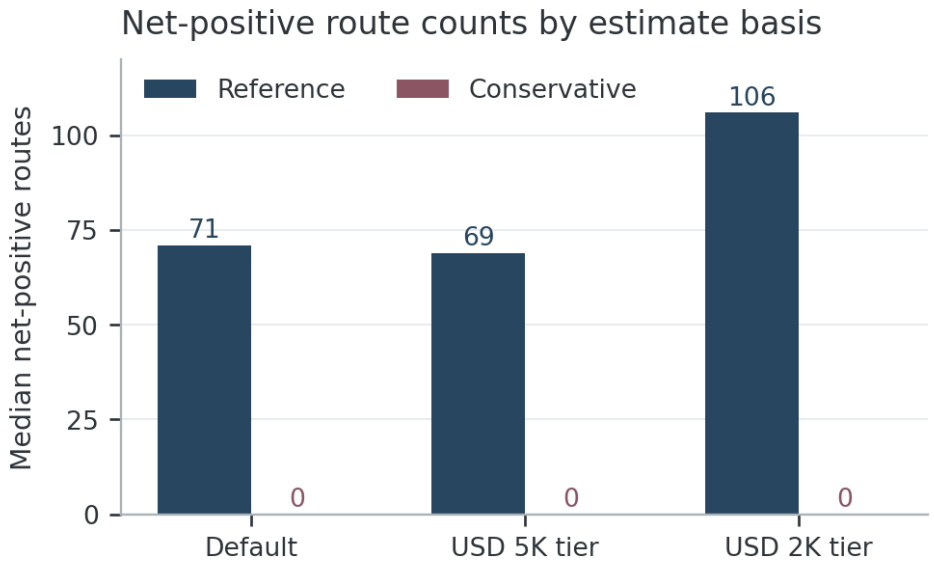


Figure 3. Net-positive route counts under reference and conservative estimates. Source: (humarb, 2026, §4C); 2026-11-01 12:00–13:05 UTC, 343 consecutive publishes with no gaps. Reference values are medians across publishes; conservative values are zero for all reported tiers. No percentages are calculated because denominators are not disclosed.

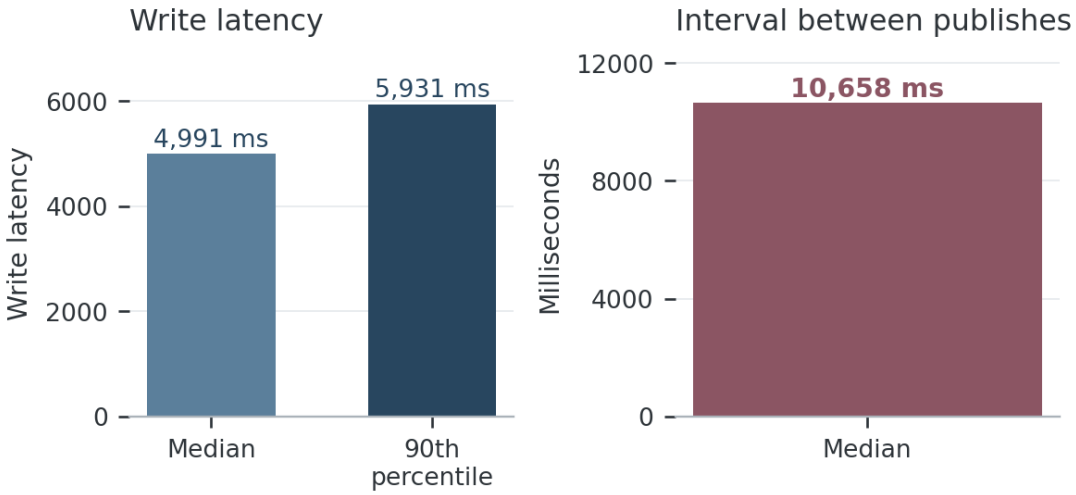


Figure 4. Engine cadence during the one-hour window. Source: (humarb, 2026, §4C); 2026-11-01 12:00–13:05 UTC, 343 consecutive publishes with no gaps. The left panel reports median and 90th-percentile write latency; the right panel reports the median interval between publishes.

been available; their net carry is simply unmeasured under the framework.

Figure 5 separates the coverage measure from the disclosed missing-depth count. It does not present a complete reason-composition pie chart. The humarb framework identifies the main reason but does not disclose a full mutually exclusive classification of every other no-estimate route, so the remaining reason counts are **not disclosed**. The distinction is substantive: missing depth is a measurement constraint generated by the refusal to extrapolate, not an economic zero. A screening framework that silently replaced missing books with top-of-book or smaller-size estimates could publish more routes, but those additional numbers would not share the same cost definition.

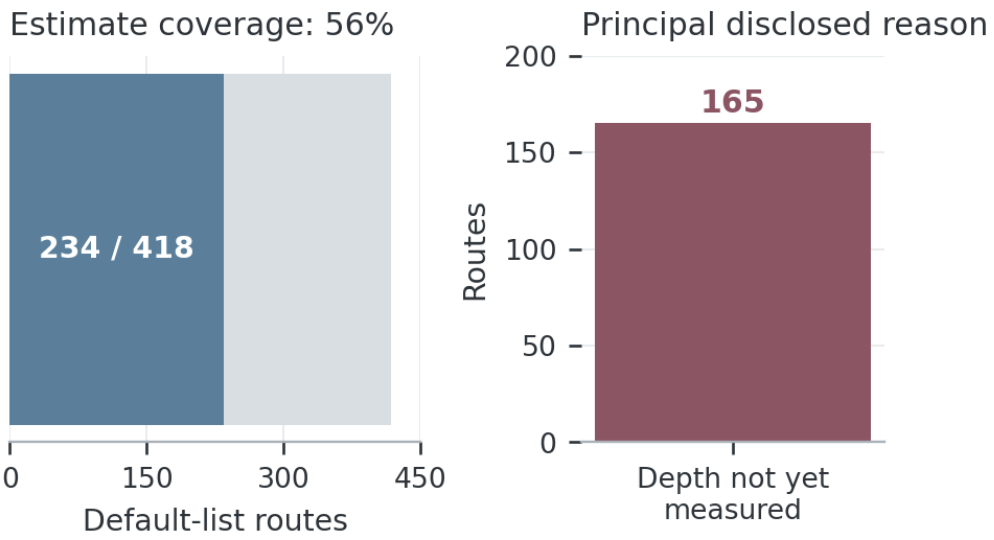


Figure 5. Estimate coverage and the principal disclosed reason for no estimate. Source: (humarb, 2026, §4B); median of six consecutive default-list publishes at 2026-11-01 14:39 UTC. Other no-estimate reasons and their individual counts are not disclosed.

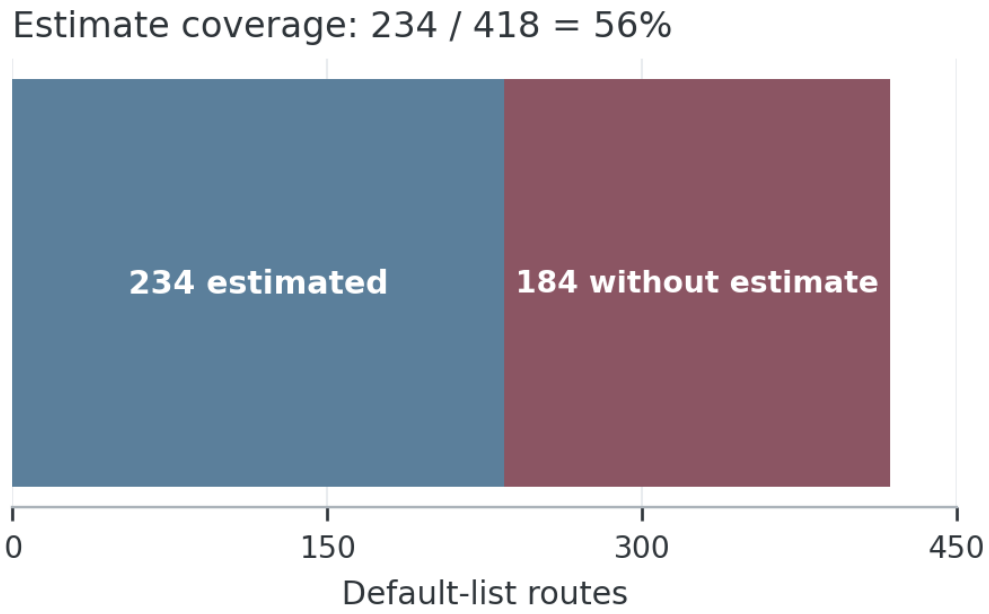


Figure 6. Default-list routes with and without an estimate. Source: (humarb, 2026, §4B); median of six consecutive publishes at 2026-11-01 14:39 UTC. The chart distinguishes coverage from net-income sign; it does not classify unestimated routes as positive or negative.

4.4 Quote-currency mismatch in the prototype audit

The earlier prototype audit provides a different form of evidence about measurement validity. On 2026-10-14, the humarb framework found that 73% of apparent positive-spread value came from mixing USDC-quoted and USDT-quoted contracts, with the source reporting "618 occurrences" in parentheses for this audit quantity (humarb, 2026, §4D). This finding motivated strict identity and quote-currency matching. The statistic concerns apparent positive-spread value, not the proportion of routes, the proportion of observations, or net income after costs. The occurrence

unit is not further defined in the humarb framework and is therefore not reinterpreted here as a route count.

Figure 7 consequently displays only the reported 73% share of apparent positive-spread value. It does not compare that share numerically with the 90.8% negative-income share in Table A. The two statistics have different denominators, dates, system versions, and economic meanings. The audit demonstrates that apparent spreads can be manufactured by identity mismatch before execution costs are even considered; Table A demonstrates that, in a later full-list snapshot with the stated controls, most reference-estimated routes were negative after measured costs. Keeping these stages separate prevents a false causal claim that quote mixing alone explains the subsequent negative estimates.

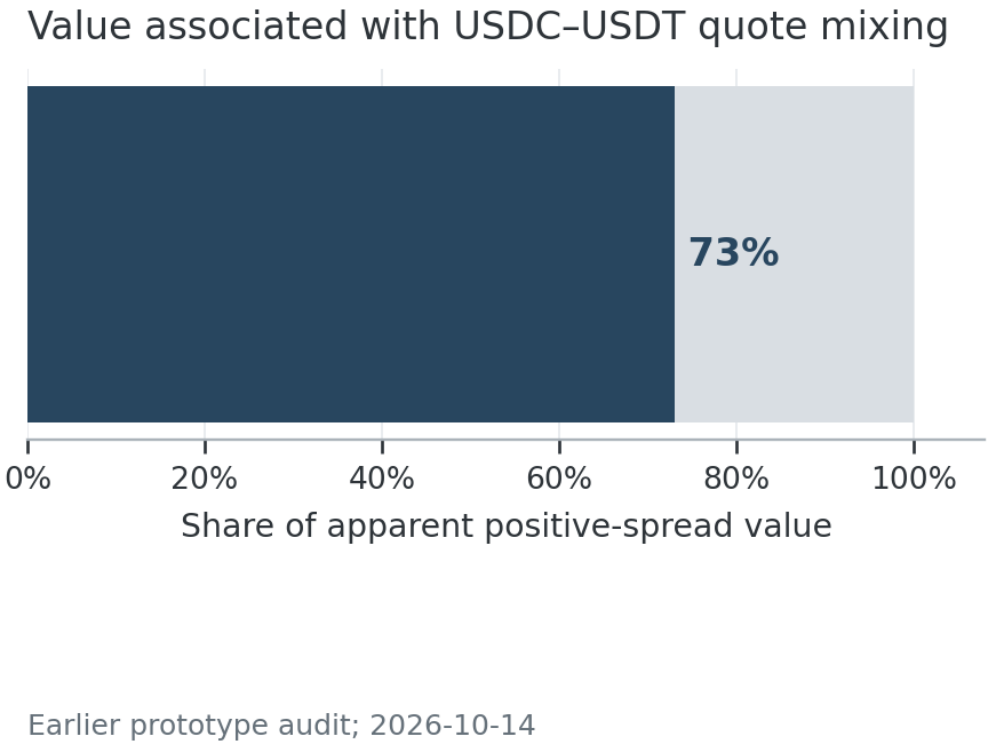


Figure 7. Share of apparent positive-spread value associated with USDC–USDT quote-currency mixing. Source: (humarb, 2026, §4D); earlier prototype audit on 2026-10-14; the humarb framework also reports "618 occurrences" without defining them as routes. This quantity is not comparable with the 90.8% net-income share in Figure 1.

4.5 Interpretation and evidentiary limits

Read together, the snapshots show why a gross-spread ranking and a net-carry estimate answer different questions. The full-list result establishes that most reference-estimated routes had negative net income after all measured costs at one publish. The one-hour result establishes that no net-positive routes were reported on the conservative basis during the specified window. The default-list result shows that a substantial part of the output is constrained by missing estimation inputs, especially unmeasured depth. Finally, the prototype audit shows that identity controls can eliminate artificial spread sources before the funding and cost models are applied. Each result supports a specific stage of the measurement argument; none supplies a long-horizon performance estimate.

The evidence does not identify the independent contribution of fees, entry walking, exit walking, or rate normalization to each negative route. The humarb framework's narrative states that most routes shown by a gross-spread screen do not pay after the measured deductions, but it does not disclose the route-level gross-spread sample or counterfactual net-income values needed to reconstruct that transition. No such decomposition is attempted here. The results also do not establish realized losses, because the estimates are not executed-fill records. The defensible conclusion is narrower and directly supported: under the disclosed reference and conservative constructions, apparently favourable funding differences often fail to survive the full measurement process, and data-quality refusals are an integral part of that process.

5 Risks and Limitations

5.1 Snapshot, coverage, and rate-basis limitations

The empirical results are tied to particular publication windows and cannot be generalized to other dates or market states. Table A is one publish, Table B is a median of six publishes, Table C is a one-hour sequence, and Table D is an earlier prototype audit. Venue coverage also varies by data class. Not every venue publishes margin tiers; two venues publish none, and some others require an authenticated read. At the time of writing, only about 10–14% of routes had a full liquidation conclusion (humarb, 2026, §5). These coverage limits restrict both the representativeness of the estimates and the completeness of accompanying risk measures. Missing values are not neutral observations and should not be interpreted as favourable or unfavourable evidence.

Rate-basis heterogeneity creates a second limitation. Venues without settlement history initially rely on current rates, which the humarb framework characterizes as weaker predictors than seven-day medians. Continuously accruing venues use hourly snapshots rather than realized payments. For one such venue, receiving-leg income depends on open-interest imbalance that the route's own position would dilute; those receiving legs were withheld from estimation at the time of writing. Seven-day medians reduce spike sensitivity but do not eliminate regime change, and the regime-shift flag does not modify the estimate. The reference and conservative values are scenario outputs built on these inputs, not bounds on all possible funding realizations.

5.2 Execution, freshness, and unpriced-cost limitations

Some venues have no order book, and one provides only request-for-quote indicative prices; in that case no depth can be walked and only an indicative spread is shown. Missing depth readers were the main disclosed source of no-estimate rows in the measured default-list window. Book-specific freshness rules can also reduce coverage in thin markets that do not update frequently. These are deliberate consequences of the measurement design, but they mean that the estimated subset may differ systematically from the unestimated subset. The humarb framework does not disclose enough route-level information to test that selection mechanism. Realized execution can further differ through partial fills and delays between legs, neither of which is converted into a numerical adjustment in the reported estimates.

The net-income equation also excludes several costs. Transfer and withdrawal costs, borrow cost, and currency conversion when the rate feed is missing are not priced. Deposit and withdrawal availability can determine whether margin can be topped up, but that channel is outside the numerical model. Consequently, "net" means net of the measured fee and book-

walk components, not net of every possible cash or operational frictional cost. This restriction is especially important for cross-venue positions whose margin is separated: profit on one leg does not automatically make capital available to support the other. The reported net income should not be interpreted as an account-level after-all-expenses return.

5.3 Liquidation and realization limitations

The liquidation model compares isolated-margin distance with a maximum single-day range over 30 closed UTC days. A seven-day holding period can experience a larger cumulative move, and no cumulative drawdown is modelled. Cross-margin accounts liquidate differently, while the displayed isolated-leg loss bound does not describe aggregate account exposure. Venue-specific liquidation events and changes in margin tiers may also alter outcomes. The model's separation of net carry and liquidation distance is useful for transparency, but it prevents the estimates from being read as a complete risk-adjusted valuation. A route with positive estimated carry can still have an unmeasured or inadequate liquidation margin.

Finally, realized funding can differ from either estimate because rates change, settlement rules vary, and venue events intervene. Partial fills, execution delays between legs, and endogenous effects on open-interest-dependent payments add further differences between measurement and realization. The conservative side addresses a defined set of adverse inputs; it is not a worst-case guarantee. These limitations do not negate the descriptive result that cost and freshness controls materially alter route classification. They do, however, confine the result to the status of a screening snapshot. The humarb framework supports no realized-return distribution, confidence interval, or forward performance claim.

6 Conclusion

Executable net carry requires a measurement definition that is more restrictive than an instantaneous funding-rate differential. The humarb framework combines identity and unit controls, daily paired seven-day medians, size-specific entry and exit walks, published taker fees, book-specific freshness tests, and explicit missing-data labels. Its reference and conservative estimates serve different purposes: the former summarizes recent settled or observed funding on a robust basis, while the latter applies less favourable rate and exit-slippage assumptions. This architecture makes the meaning of a reported number inspectable and prevents missing depth or unverified books from being treated as zero-cost execution. The central methodological lesson is that refusal to estimate can be as informative as a numerical estimate when the required evidence is absent.

In the single full-list publish at 2026-11-01 19:19 UTC, 462 of 509 reference-estimated routes—90.8%—had negative net income after all measured costs at USD 5,000 per side over seven days. In the separate 2026-11-01 12:00–13:05 UTC window, reference-side median net-positive counts coexisted with zero conservative-side net-positive counts across all reported tiers. These findings do not measure realized profitability and do not support extrapolation beyond their snapshots. They demonstrate that the economic interpretation of cross-venue funding arbitrage changes materially when robust funding expectations, executable size, bilateral costs, and current book confirmation are incorporated into one measurement framework.

Disclaimer: This paper describes a screening methodology and a single measurement snapshot; it does not constitute investment advice.

Data Availability and Reproducibility Statement

All methodological parameters and empirical values reported in this article are transcribed from the humarb Methodology Pack (humarb, 2026). The article contains the numerical inputs required to interpret its seven figures and Table 1, but the underlying route-level observations are not distributed with this submission. No additional dataset or implementation code is supplied. Values not reported in the source are identified as **not disclosed** rather than reconstructed. Consequently, the article supports verification of its reported snapshot arithmetic and measurement definitions, but not independent recomputation of route-level estimates.

Use of Language Tools and Author Responsibility

Language-model assistance was used for drafting and copyediting. The named author(s) reviewed the manuscript and assume responsibility for its content, numerical transcriptions, citations, and references. No generative language tool is listed as an author.

References

Damien Ackerer, Julien Hugonnier, and Urban Jermann. Perpetual futures pricing. *Mathematical Finance*, 36(3):481–499, 2026. doi: 10.1111/mafi.70018.

Julia Ackermann, Thomas Kruse, and Mikhail Urusov. Optimal trade execution in an order book model with stochastic liquidity parameters. *SIAM Journal on Financial Mathematics*, 12(2):788–822, 2021. doi: 10.1137/20M135409X.

Carol Alexander, Jaehyuk Choi, Hamish R. A. Massie, and Sungbin Sohn. Price discovery and microstructure in ether spot and derivative markets. *International Review of Financial Analysis*, 71: 101506, 2020a. doi: 10.1016/j.irfa.2020.101506.

Carol Alexander, Jaehyuk Choi, Heungju Park, and Sungbin Sohn. BitMEX Bitcoin derivatives: Price discovery, informational efficiency, and hedging effectiveness. *Journal of Futures Markets*, 40(1):23–43, 2020b. doi: 10.1002/fut.22050.

Carol Alexander, Xi Chen, Jun Deng, and Tianyi Wang. Arbitrage opportunities and efficiency tests in crypto derivatives. *Journal of Financial Markets*, 71:100930, 2024. doi: 10.1016/j.finmar.2024.100930.

José Almeida and Tiago Cruz Gonçalves. Cryptocurrency market microstructure: A systematic literature review. *Annals of Operations Research*, 332(1–3):1035–1068, 2024. doi: 10.1007/s10479-023-05627-5.

Robert Almgren and Neil Chriss. Optimal execution of portfolio transactions. *The Journal of Risk*, 3 (2):5–39, 2000. doi: 10.21314/jor.2001.041. Winter 2000 issue; DOI registered in 2001.

Patrick Augustin, Alexey Rubtsov, and Donghwa Shin. The impact of derivatives on spot markets: Evidence from the introduction of Bitcoin futures contracts. *Management Science*, 69(11):6752–6776, 2023. doi: 10.1287/mnsc.2023.4900.

Jean-Philippe Bouchaud, Marc Mézard, and Marc Potters. Statistical properties of stock order books: Empirical results and models. *Quantitative Finance*, 2(4):251–256, 2002. doi: 10.1088/1469-7688/2/4/301.

Rama Cont, Arseniy Kukanov, and Sasha Stoikov. The price impact of order book events. *Journal of Financial Econometrics*, 12(1):47–88, 2014. doi: 10.1093/jjfinec/nbt003.

Riccardo De Blasis and Alexander Webb. Arbitrage, contract design, and market structure in Bitcoin futures markets. *Journal of Futures Markets*, 42(3):492–524, 2022. doi: 10.1002/fut.22305.

Anne H. Dyhrberg, Sean Foley, and Jiri Svec. How investible is Bitcoin? analyzing the liquidity and transaction costs of Bitcoin markets. *Economics Letters*, 171:140–143, 2018. doi: 10.1016/j.econlet.2018.07.032.

Antje Fruth, Torsten Schöneborn, and Mikhail Urusov. Optimal trade execution and price manipulation in order books with time-varying liquidity. *Mathematical Finance*, 24(4):651–695, 2014. doi: 10.1111/mafi.12022.

humarb. humarb methodology pack: Net-carry estimation for cross-venue perpetual funding arbitrage. Unpublished methodology document, Sections 1–6, 2026.

Emre Inan. Predictability of funding rates. *International Review of Financial Analysis*, 119:105464, 2026. doi: 10.1016/j.irfa.2026.105464.

Ralph S. J. Koijen, Tobias J. Moskowitz, Lasse Heje Pedersen, and Evert B. Vrugt. Carry. *Journal of Financial Economics*, 127(2):197–225, 2018. doi: 10.1016/j.jfineco.2017.11.002.

Seungho Lee, Nabil El Meslmani, and Lorne N. Switzer. Pricing efficiency and arbitrage in the Bitcoin spot and futures markets. *Research in International Business and Finance*, 53:101200, 2020. doi: 10.1016/j.ribaf.2020.101200.

Maik Schmeling, Andreas Schrimpf, and Karamfil Todorov. Crypto carry. *Management Science*, 2026. doi: 10.1287/mnsc.2024.05069. Advance online publication.

Robert J. Shiller. Measuring asset values for cash settlement in derivative markets: Hedonic repeated measures indices and perpetual futures. *The Journal of Finance*, 48(3):911–931, 1993. doi: 10.1111/j.1540-6261.1993.tb04024.x.