

External Data for the Next 15 Minutes

What the internet evidence says can — and cannot — improve direction calls on Kalshi 15-minute markets: BTC · Gold · ETH · Copper. Evidence review read 2026-10-21, cross-checked against our own tests on 6,482 BTC / 4,616 Gold / 3,656 Copper settled windows and 35k+ blind synthetic windows.

BOTTOM LINE — Almost nothing published clears the bar as a standalone next-15-minutes direction signal. One family does: live order-book / trade-flow imbalance, used in the last 1-5 minutes of a window, on the venues that set the settlement price — worth ~1-4 probability points, so it pays only when the Kalshi quote has not caught up. Everything else credible is a volatility / window-selection input, a same-second repricing input, or an effect measured at 72-hour to daily horizons that does not transfer. The good news: the one real signal is FREE to pull.

The pull list — all of it, in order

#	Feed	Cost	Assets	What it is for
1	Multi-venue crypto order-book + trades websockets: Binance spot & perp, Bybit, Coinbase, Kraken, Bitstamp, Gemini	FREE, no key	BTC, ETH	Top-5 order-book imbalance (OBI) and perp-vs-spot trade flow (CVD) over rolling 1-3 min — the only measured signal at this horizon (rank-IC +0.113 BTC / +0.082 ETH at 1 min).
2	CME Globex via Databento: GC, HG, CME BTC/ETH, ES, NQ, ZN, 6E in one feed	Paid, ~\$199/mo	Gold, Copper, macro	GC is where gold price discovery happens (peer-reviewed); HG for copper. No free real-time CME source exists. Also carries the coincident macro complex (yields, dollar proxy, risk).
3	Settlement + cross-venue bundle: Kalshi API, Pyth Hermes (metals), Polymarket CLOB websocket	FREE	All four	Price the last 5 minutes off the settlement series itself (60-second RTI average for crypto; Pyth candles for metals). Polymarket is a free second book — act only on fresh-both-sides gaps over ~4¢.
4	Exchange-native derivatives flow: Binance/Bybit liquidation stream (forceOrder), mark price/funding, open-interest polling	FREE native	BTC, ETH	Cascade state gate on signal #1: ride a running liquidation cascade, skip fragile crowded windows. Never a standalone trigger.
5	Deribit public API: DVOL, book summaries (open interest, mark IV), expiry calendar	FREE	BTC, ETH	Volatility calibration (fair-probability check on distance-to-target) and avoid-flags: 08:00 UTC expiries, macro-release windows. Explicitly NOT a direction input.

Sections 1-8 give the evidence per candidate, §9 the ML literature, §10 the cross-check against our own data (the two agree on every point), §11 the build-ready pull list with endpoints, §12 the useless list, §13 what nobody could verify, §14 sources.

1 · The hurdle any signal must clear

Kalshi taker fee = $\text{roundup}(0.07 \times C \times P \times (1-P))$. At $P = 0.50$ that is \$1.75 per 100 contracts = 1.75¢ per contract (3.5% of the \$50 cost). Add a typical 1-2¢ half-spread on 15-minute markets and a taker needs the true probability to be $\approx 3-4$ points above the quoted price just to break even. (The "7%" people quote is the fee coefficient, not the cost at mid.)

The market to beat is already calibrated: a practitioner audit of 2,764 Kalshi BTC 15-minute windows (29 Aug – 27 Sep 2026) found every 10¢ price band won within 1 point of its price, and exactly half of 3,320 windows settled up. Price-alone information — including the move so far in the window — is already in the quote. Any external signal must add on top of that.

Settlement defines what "external" can even mean. Crypto settles on the CF Benchmarks Real-Time Index: the average of 60 one-second prices in the final minute vs the same average from the minute before the window opened (tie → YES) — so in the last minute each second is only 1/60 of the settlement average, and signals must predict the running average, not the last print. Gold and copper settle on Pyth 1-minute candle closes; the public Pyth Metal.XAU/USD feed reproduced the settled verdict on 99.5% of 200 measured gold windows. Pyth aggregates first-party quotes from the same firms that quote COMEX GC/HG — so for metals, futures flow is settlement-adjacent in a way it is not for crypto.

2 · Order-book imbalance / trade flow — the one real signal

VERDICT — USE, late in the window only. Snapshot imbalance predicting the whole next 15 minutes is worthless (~50.8% accuracy). Fresh imbalance (last 1 minute) predicting the next 1-5 minutes is ~53.6% BTC / 52.6% ETH — at or just above the fee breakeven, and only against a Kalshi quote that is stale by those 2-4 points at that moment.

Best evidence — Imbalance Labs study (Jul 2026, vendor/practitioner, grade C+, ~1.5M one-minute bars per asset, fully reproducible): top-5 notional order-book imbalance vs forward mid-price returns on Hyperliquid L2 data:

Asset	Rank-IC @1 min	@5 min	@15 min	@60 min
BTC	+0.1129	+0.0508	+0.0244	+0.0105
ETH	+0.0824	+0.0359	+0.0163	+0.0012

BTC deciles are cleanly monotonic: mean next-1-minute return runs from −1.13 bps (most ask-heavy decile) to +1.12 bps (most bid-heavy) — a 2.25 bps top-to-bottom spread per minute. The vendor itself states the boundary honestly: "the raw signal sits inside the taker fee; harvesting it requires real execution." Conversion used throughout this report: accuracy $\approx 0.5 + \arcsin(\text{IC})/\pi$.

Do not confuse this with the famous order-flow result: Cont-Kukanov-Stoikov (2014) show order-flow imbalance explains price changes with $R^2 \approx 0.65$ contemporaneously (same 10-second window) — that is description, not prediction. Genuinely forward regressions collapse fast: one BTC draft reports forward R^2 of 0.094 at 1 second and 0.019 at 10 seconds. Signed trade flow (CVD) likewise explains ~60–85% of returns contemporaneously, which is exactly why lagged CVD adds little — it is already in the price. The incremental use is divergence: perp CVD accelerating while spot CVD is flat in the final minutes = forced/liquidation flow (\$\$).

The BRTI-constituent question — no published study exists. Nobody has measured whether imbalance on CF Benchmarks constituent exchanges (Coinbase, Kraken, Bitstamp, Gemini, LMAX Digital, Bullish, Crypto.com — official list) predicts the index better than imbalance elsewhere. What is documented: the original BRTI is built from constituents' consolidated order books, so constituent depth mechanically enters the settlement index contemporaneously, while Binance flow reaches it only via arbitrage (sub-seconds to seconds). Inference, not evidence: constituent books are the settlement-relevant books; the Binance perp is the information-leading book. A live system wants both — they cost the same (free). This is the single most valuable original backtest to run (\$13).

Access — free. Exchange-native public websockets, no API key: Binance spot & USD-M futures (depth20@100ms, aggTrade), Bybit v5 (orderbook.200, publicTrade), Coinbase Advanced Trade WS, Kraken v2 WS, Bitstamp, Gemini. Aggregators (Kaiko ~\$1,500/mo, Tardis) are not needed for a first

build.

3 · Cross-exchange / cross-venue lead-lag

VERDICT — Leads are real at 100 milliseconds to ~1 second, not minutes. No credible evidence of a seconds-to-minutes lead usable at the 15-minute scale — with one structural exception: for gold, COMEX futures are where price discovery happens, so GC, not spot XAU, is the series to model.

- Crypto venues: a 2026 practitioner study over 29 assets finds Binance leads Hyperliquid on all 29, peak correlation at a ~600–800 ms lag (BTC); Binance vs Lighter peaks at ~100 ms. These lags tell you which feed to treat as fair value — not a 15-minute signal.
- CME vs spot (BTC/ETH): the peer-reviewed literature splits by sampling frequency, which is itself the answer — Corbet et al. (2018, 1-min) and Baur & Dimpfl (2019, 5-min) find spot leads; Kapar & Olmo (2019, daily) find futures dominate. No study shows a minutes-long exploitable CME lead in modern data. CME crypto futures also do not trade Saturday, while Kalshi trades 24/7.
- BTC → ETH: transmission at 100 ms – 5 s. BTC's move is a legitimate input to ETH windows (high contemporaneous beta) but the Kalshi ETH price sees it within seconds too. ETH's own-book imbalance is ~0.73× BTC's strength.
- GOLD (peer-reviewed, grade A–): Hauptfleisch, Putniņš & Lucey, Journal of Futures Markets — COMEX GC plays a larger role in gold price discovery than the London spot market despite trading under a tenth of its volume; nobody finds spot leading futures. Discovery share is not a minutes-ahead forecast — it means GC flow/imbalance is the gold analogue of \$2, measurable on CME book data.
- COPPER: no leader. Information shares ≈ COMEX 54.8% / LME ~15% / SHFE ~15% in one study (with bounds so wide the ranking is fragile), LME 47–49% in another. Discovery is shared and session-dependent. Usable fact: SHFE sessions (01:00–07:00 and 13:00–17:00 UTC) are live marginal-venue hours; outside them HG is the market.

Value of a leader feed = repricing speed (knowing fair value before the Kalshi book updates), worth the 1–2¢ of staleness that occasionally appears — not a directional edge. Access: crypto leaders free (\$2). CME: no free real-time source; Databento CME Globex (Standard ≈ \$199/mo) is the self-serve route. CME's own websocket is licensed. 10-minute-delayed CME quotes are useless here by construction.

4 · Polymarket vs Kalshi

VERDICT — No published evidence either venue's 15-minute price leads or systematically corrects the other. Useful as a free second book / stale-quote detector, not a signal.

A matched-pair audit of 186 Polymarket-Kalshi outcome pairs with both quotes fresh found a median gap of 0.78¢; only 1.5% of pairs differed by ≥5¢, and none of the 12 largest headline gaps survived a staleness check (one side simply had not traded recently). Over six months: 385 sustained gaps ≥2¢ — "the honest size of the opportunity before fees." Median gap 0.78¢ is far under the 1.75¢ one-leg fee; only maker-rested, fresh-both-sides gaps over ~4¢ are candidates, and they are rare. Two structural warnings: (i) the venues settle differently (Kalshi crypto: 60-second RTI average; Polymarket: Chainlink stream), so a late-window "gap" is often settlement basis, not edge; (ii) a Stanford/SMU working paper (Jul 2026) documented settlement manipulation in Polymarket 5-minute BTC (Binance net flow in the final 10 seconds up ~50%, ~3.9× in near-even cycles, ~\$8.2M extracted) — 15-minute contracts showed far less of this, which is reassuring for Kalshi 15m. Access: free — Polymarket CLOB websocket (no auth) + Gamma API; Kalshi REST is public for market data, its websocket needs an API-key handshake.

5 · Options: DVOL, skew, expiries, max pain

VERDICT — Useless for 15-minute direction. Genuinely useful for two things: converting distance-to-target into a fair probability (volatility calibration) and marking windows to avoid.

- DVOL is direction-free by construction (a variance index). No study shows DVOL changes predicting sub-hour direction. Correct use: σ per 15 min $\approx$ DVOL $\times$ 0.53% (DVOL 40 $\rightarrow$ $\pm 0.21\%$ one-sigma) to sanity-check whether Kalshi's price for the current distance is rich or cheap — the market mostly embeds this already.
- 25-delta skew / risk reversal: a days-to-weeks regime gauge. Intraday skew moves are responses to spot (hedging demand during the move) — coincident, not leading. No published sub-hour effect size exists.
- Expiries / pinning / max pain: Deribit expiries are 08:00 UTC (no hourly expiry of size), delivery price a 30-minute TWAP (07:30–08:00 UTC) — those Kalshi windows are structurally noisy; avoid or size down. On pinning, Deribit itself reports "limited evidence of a consistent pinning effect" at recent quarterly expiries. Max pain has no forcing mechanism (public open interest does not reveal who is long or short, so dealer-hedge direction is unknowable) — treat "price gravitates to max pain" as unsupported. COMEX gold options / GVZ: same structure, weaker — vol-only.

Access: Deribit public API/websocket free (volatility index data, book summaries with OI and mark IV — skew is computed from these). CME options: Databento, paid.

6 • Funding, open interest, liquidations

VERDICT — Funding level and OI level: no credible 15-minute direction evidence. Live liquidation FLOW is the one usable piece — as a cascade state gate on \$2, never a standalone trigger. Liquidation heatmaps are modelled guesses that cannot predict.

Funding: the best quantified study (MarketTrace, 789 days, 2024–2026, 6 assets, 3 venues) finds after BTC funding hit its top decile, the median 72-hour forward return was -1.18% with a 40.5% up-rate ($n = 37$) — but on ETH the same condition preceded positive median returns ($+0.7\%$ to $+1.7\%$), and after bottom-decile funding (622 episodes) hit rates were $\approx 50\%$: "carried no measurable edge on direction." Funding is also an 8-hour-interval rate on Binance/Bybit — inside a 15-minute window it is a constant. No 15-minute-horizon funding study exists because the variable barely moves at that horizon.

Open interest: the standard quadrant reading (price up + OI up = new longs, etc.) describes the move already in your baseline. No clean effect size for Δ OI predicting the next 15 minutes was found.

Defensible use: interaction — rising OI + extreme funding + thin book = cascade fragility (size down / skip), not direction.

Liquidations: forced market orders are aggressive flow, so while a cascade is running, continuation over the next 1–5 minutes is the \$2 effect with extra fuel; cascades also exhaust, and the post-cascade reversal is equally documented anecdotally. No published out-of-sample hit rate exists for "liquidation spike $\rightarrow$ next-15-min direction." Heatmaps are weaker still: exchanges do not publish liquidation prices; heatmaps project OI onto modelled levels, and a careful explainer states flatly: "Can it predict where Bitcoin goes next? It cannot." Access: exchange-native and free beats aggregators — Binance futures liquidation stream (note: throttled to the latest liquidation per symbol per 1,000 ms, so it undercounts cascades), mark-price stream (funding + mark), REST open-interest polling; Bybit v5 public websocket carries OI + funding. CoinGlass API v4 (from \$29/mo) aggregates venues at minute granularity — fine for OI/funding, too slow as a primary liquidation feed.

7 • Macro drivers and scheduled releases

VERDICT — DXY / real yields / equity futures are coincident drivers: they move in the same seconds as the asset, so they cannot lead it in quiet times. Their value concentrates in the 1-3 minutes after scheduled releases, and in release-window avoidance.

Announcement science (peer-reviewed, Andersen-Bollerslev-Diebold-Vega; Fed review of the high-frequency literature): macro surprises jump prices "within a few minutes or less," while volatility remains elevated much longer. That asymmetry is the tradable content: in a CPI / payroll / FOMC window, direction is set by the surprise vs consensus — unknowable in advance — while spreads and Kalshi repricing speed both explode. Blackout the 15-minute windows containing 08:30 ET CPI/NFP/PPI and 14:00-14:30 ET FOMC, or trade only post-release continuation with yield confirmation.

Gold vs DXY / real yields: strong at daily frequency, unstable by regime — 52-week daily correlations of gold returns vs 10y TIPS real-yield changes ≈ -0.35 to -0.43 , vs DXY ≈ -0.35 to -0.42 — and 2026 supplied a break (gold spiked to records in March 2026 while yields rose). At minute scale these are betas, not leads. There is no free real-time TIPS yield (FRED is daily); the real-time proxies are CME ZN (10y futures) and 6E (EUR futures $\approx$ inverse DXY — DXY itself is an ICE product) on the same Databento feed. Equity futures: BTC-Nasdaq correlation is regime-switching; no credible study shows ES/NQ leading BTC by minutes; overnight and weekends, crypto routinely moves first.

The release calendar itself is free (economic calendars) for scheduling blackouts. The surprise is only tradable with a low-latency machine-readable news feed (paid: Bloomberg / Reuters / Dow Jones) — without one, you are the exit liquidity in release windows.

8 · Gold-specific and copper-specific

- London fix: LBMA Gold Price auctions at 10:30 and 15:00 London are liquidity/benchmark events, not direction signals.
- Asian-hours drift: a documented "hat-shaped" seasonality (gold appreciates in Asian hours, depreciates the rest of the day) earned Sharpe up to 1.61 gross in the one backtest found — and underperformed after transaction costs. Per-15-minute magnitude is a few bps: at most a prior tilt, never a trigger.
- Central-bank purchases: monthly, lagged, revised. WGC gold-ETF flows: updated weekly (Monday of the following week). GLD creation/redemption: daily. None of this exists at 15-minute frequency in any form — frequency-disqualified.
- Copper inventories: LME warehouse stocks are daily and T+1 (position at 16:30 on day T published 09:00 on T+1, official LME document). COMEX stocks daily; SHFE warrant stocks after the close. A T+1 daily stock change cannot predict a 15-minute window, full stop.
- China macro (PMI, CNY, credit): daily/monthly releases — same blackout logic as §7 around 01:30-03:00 UTC release times, not a continuous signal. Copper's intraday external drivers in practice are the same CME-feed complex as gold's (6E, ZN, ES) — coincident.

9 · What the ML literature actually found

Study	Setup	Out-of-sample result	Features that carried it
Akyildirim et al., CME BTC futures (peer-reviewed)	Mid-price direction, 5-60 min	Logistic ~54%, XGBoost ~55%; SVM 71% in one split, random forest 87% in-sample = overfit signature	Own-price history only — no external features tested
Multi-timeframe BTC (2026)	15-min bars + 4h/daily features, true 2025 holdout	AUC 0.57-0.61; correcting look-ahead removed ~0.20 AUC of fake performance; net after fees "likely negligible"	4-hour Bollinger position and RSI — the asset's own slower-timeframe technicals
40-asset walk-forward (Frontiers 2026)	1-min data, 1-30 min horizons, technicals only	Only 5 of 40 assets profitable after costs	Technicals; the lesson is costs, not features
MDPI Applied Sciences 2025	Minute limit-order-book + OHLCV + daily macro	82.68% accuracy — only on the 11.99% of minutes traded under a confidence threshold	Order-book features dominate technicals

Synthesis: across the credible literature, out-of-sample 5-15-minute direction accuracy clusters at 52-55% unconditional (AUC $\approx$ 0.57-0.61), before costs, mostly from the asset's own price history — which the Kalshi quote already reflects. The only feature family that repeatedly adds anything beyond own-price is order-book / order-flow (§2). No study was found in which funding, options skew, macro levels, or cross-venue prediction prices improved 15-minute out-of-sample accuracy.

10 · Cross-check against our own data

Everything above was tested independently on our own Kalshi datasets before this report existed. The two agree on every point — this is why the live calls are built the way they are:

- Kalshi's own % is the best probability at every minute from 5 to 1 left (BTC, 6,482 windows: Brier 0.130 / 0.115 / 0.094 / 0.074 / 0.034, well calibrated). No local feature — momentum, path shape, % change, volume, pre-window trend, regime — beat it at any minute.
- Momentum is dead: on 34,982 blind synthetic BTC windows (a pre-Kalshi year) momentum added nothing over the lead alone (79.4 vs 79.5); same null on 25,887 gold windows. Copper is the lone exception — a small real gain (+0.3 pt, $\geq 90\%$ coverage 25.7% $\rightarrow$ 27.9%) — deliberately not wired into the live call.
- Direction-feature hunt (60-day, both assets, holdout splits): cross-asset signals (ETH $\rightarrow$ BTC, BTC $\rightarrow$ Gold) identical accuracy to own-lead; last-3-minute burst worse or reversed out-of-sample; no feature ever flips a call — features only sort confidence. Consecutive windows are independent (P(next UP) 48-53%): no honest pre-open call exists.
- Volatility regime IS a real sorter (the one "external-ish" input that survived): BTC \$50-100 leads held 89.8% in low-vol vs 78.6% in high-vol; σ -normalized buckets are cleaner for BTC (AUC 0.729 vs 0.709); the copper high-vol shave (-3 on the \$0.002-0.005 band in the top vol tercile) is already wired live. For gold, σ -normalizing does NOT help (AUC 0.7331 vs raw 0.7385).
- Settlement-averaging is real and already in the math: leads of \$25+ collapse between 2 and 1 minutes left at exactly the rate the 60-second-average formula predicts — the final-minute trap zone ($|\text{lead}|$ under $\sim \$10$ on BTC) flips 44.9% at 5 minutes left, 32.2% at 1 minute.
- Order-book imbalance: our Vol 1 research measured the same +0.113 rank-IC at 1 minute that this report's vendor study found (+0.1129) — and the same conclusion: extreme deciles move ± 1.1 bp next-minute, inside the taker fee. Probability tilt only, never a direction flip.

11 · The build-ready pull list

Exact feeds, in wiring order. Everything except #2 is free and needs no account or API key.

#	Feed and endpoints	Cost	Use rule
1	Crypto L2 + trades websockets — Binance spot & USD-M futures: @depth20@100ms, @aggTrade; Bybit v5: orderbook.200, publicTrade; Coinbase Advanced Trade WS; Kraken v2 WS; Bitstamp; Gemini	Free	Rolling 1-3 min top-5 OBI + perp-vs-spot CVD, BTC & ETH. Constituent venues (Coinbase/Kraken/Bitstamp/Gemini) = settlement-relevant book; Binance perp = leading book. Act only when implied prob shift $\geq \sim 4$ pts vs the Kalshi quote.
2	CME Globex via Databento (MDP 3.0): GC, HG, BTC, ETH, ES, NQ, ZN, 6E — one connection	$\sim \$199/\text{mo}$ + usage	Gold: model GC flow/imbalance (the price-discovery venue), not spot. Copper: HG book + session awareness. Macro complex (ZN yields, 6E dollar proxy, ES/NQ risk) as coincident confirmation, mainly in release minutes.
3	Settlement bundle — Kalshi Trade API (market data public; WS needs key handshake); Pyth Hermes websocket/REST (Metal.XAU/USD + copper feed); Polymarket CLOB websocket (no auth) + Gamma API	Free	Track the settlement series itself in the last 5 minutes (RTI running 60-s average; Pyth candle for metals). Polymarket as second book: only fresh-both-sides gaps $> \sim 4\text{¢}$.
4	Derivatives flow — Binance futures liquidation stream (@forceOrder; throttled to latest/symbol/1,000 ms) + @markPrice (funding/mark); REST openInterest polling; Bybit v5 tickers (OI + funding). Optional: CoinGlass API v4	Free native; CoinGlass from $\$29/\text{mo}$	Cascade gate on #1: continue with a running cascade, fade nothing into one, skip crowded fragile windows (extreme funding + rising OI + thin book). Never a standalone trigger.
5	Deribit public API/WS — volatility index data (DVOL), book summaries (OI, mark IV $\rightarrow 25\Delta$ skew), expiry calendar	Free	σ calibration for fair-probability checks ($\sigma_{15m} \approx \text{DVOL} \times 0.53\%$); avoid-flags for 08:00 UTC expiries (delivery TWAP 07:30-08:00 UTC) and macro-release windows (08:30 ET data, 14:00 ET FOMC). Not direction.

12 · Popular signals the evidence says are useless at 15 minutes

- Funding-rate level — best evidence is a mixed, asset-inconsistent effect at 72 hours; at 15 minutes the variable is a constant. Bottom-decile hit rate $\approx 50\%$ over 622 episodes.
- Max pain / pinning — no mechanism, unknowable dealer positioning, venue itself reports limited pinning evidence.
- DVOL / GVZ level or change as direction — volatility indices; direction-free by construction.
- 25-delta skew flips as a trigger — a days-weeks regime gauge; intraday moves respond to spot.
- Liquidation heatmaps as prediction — projections of OI onto modelled levels; their own explainers say they cannot predict direction. (Live liquidation flow is different — pull list #4.)
- Open-interest level / "OI rising = bullish" — contemporaneous description; no published next-15-min effect.
- Central-bank gold purchases, gold ETF flows — monthly / weekly-lagged; cannot exist at 15-minute frequency.
- LME / COMEX / SHFE inventories — daily, T+1. Structurally incapable of informing a 15-minute window.

- London-fix folklore, Asian-session drift as a trigger — real seasonality, a few bps, negative after costs in the only backtest found.
- Technical indicators (RSI/MACD/Bollinger) on the asset itself at 15 min — 52–55% gross out-of-sample, negative after costs; and already in the Kalshi price.
- Naïve Polymarket-Kalshi arbitrage / "Polymarket leads" — median fresh gap 0.78¢ vs a 1.75¢ one-leg fee; big apparent gaps are stale quotes or settlement basis.
- DXY / real-yield levels as a gold signal — daily betas (–0.35 to –0.43) that break by regime; only release-minute changes matter, as coincident inputs.
- Fear & Greed index, Coinbase-premium anecdotes, options put/call ratios — daily/weekly composites with no measured sub-hour predictive content.

13 · Could not verify — read before building on it

- BRTI-constituent vs non-constituent imbalance: no published study exists. The constituent-book argument is mechanical inference from the index methodology, not measured evidence. This is the single most valuable original backtest available: OBI on Coinbase+Kraken+Bitstamp+Gemini books vs Binance books, against the RTI's own next-1-5-minute changes.
- Polymarket vs Kalshi 15-minute lead-lag: no dataset-backed study found either way — only the general matched-pair gap audit.
- GC / HG order-flow imbalance effect sizes: OFI logic is established for equities/FX/crypto; no published GC- or HG-specific forward IC at 1–15 minutes was found.
- Kalshi metals settlement identity: rests on one detailed secondary source, which notes Kalshi names a Pyth index symbol absent from Pyth's public catalogue; measured at 99.5% verdict agreement for gold. Copper's exact Pyth feed equivalence was not independently measured in any source found — verify in Kalshi's copper contract terms before building on it.
- Regime staleness: several effect sizes predate the spot-ETF / current regime (CME-vs-spot BTC literature is 2018–2021; gold discovery shares end 2014). Treat magnitudes as upper bounds — every replication wave in this literature has found smaller effects than the one before.

14 · Sources

All read 2026-10-21. [P] peer-reviewed · [W] working paper/preprint · [V] vendor/practitioner · [O] official · [N] news. Full notes: [research report kalshi-15m-external-direction-data \(2026-10-21\)](#).

1. [V] Imbalance Labs, "Does Orderbook Imbalance Predict Crypto Returns? A 24-Asset Rank-IC Study" — imbalancelabs.com/research/orderbook-imbalance-predictive-power
2. [O] Kalshi fee schedule summaries — kalshiview.com/blog/kalshi-fees-explained-what-you-actually-pay-per-trade ; allium.so/blog/kalshi-fees-why-the-cost-peaks-near-50-cents
3. [V] PredictionMarketPicks, Kalshi 15-minute markets & settlement (calibration audit; Pyth verification) — predictionmarketpicks.com/kalshi/15-minute-markets
4. [V] PM Countdown, Kalshi crypto settlement (60-s RTI averages, tie → YES) — pmcountdown.com/markets/crypto
5. [O] CF Benchmarks, CME CF Constituent Exchanges — docs.cfbenchmarks.com ; [O] CME, BRTI methodology — cmegroup.com
6. [P] Hauptfleisch, Putniņš & Lucey, "Who sets the price of gold? London or New York," *Journal of Futures Markets*
7. [W] Figuerola-Ferretti et al., "Copper Price Discovery on Comex, the LME and the SHFE, 2001–2013," *UC3M* ; Li et al., *SHFE copper pricing (USC)*

8. [P] BTC futures/spot price-discovery literature: Corbet (2018); Baur & Dimpfl (2019); Kapar & Olmo (2019); Fassas; Akyildirim; Alexander et al.
9. [V] Lead-lag study, Binance / Hyperliquid / Lighter (600–800 ms; 100 ms) — weex.com
10. [V] PredictMarketCap, Polymarket-Kalshi matched-pair gap audit — predictmarketcap.com
11. [W/N] Dai, Jia & Yu, "Settlement Manipulation in Prediction Markets" (Stanford/SMU, Jul 2026), via Bloomberg / Bitcoin Magazine
12. [O] Deribit Insights, DVOL whitepaper — insights.deribit.com ; Deribit API docs — docs.deribit.com ; [N] Deribit on pinning & 08:00 UTC expiries
13. [V] MarketTrace, "What happens after extreme funding rates? 789 days of data" — markettrace.ai
14. [O] Binance, liquidation (forceOrder) stream docs — developers.binance.com ; [V] CoinGlass API v4 ; [V] liquidation-heatmap limits — alphafinance.io
15. [P] Andersen, Bollerslev, Diebold & Vega, macro announcement effects ; [O] Fed IFDP 823 literature review — federalreserve.gov
16. [V] Mercatus Data, gold real-yield / DXY daily correlations — mercatusdata.com/metals/gold
17. [O] LME, warehouse stock movements (daily, 09:00 T+1) — lme.com ; [O] WGC, gold ETF holdings & flows — gold.org
18. [W] CBS thesis, "Gold Price Dynamics Around the Clock" (Asian-hours appreciation; Sharpe 1.61 gross, negative net)
19. [P] Akyildirim et al., "Forecasting mid-price movement of Bitcoin futures using machine learning" — PMC8296834
20. [W] "Multi-Timeframe Feature Engineering for Bitcoin Market Prediction" (AUC 0.57–0.61; look-ahead +0.20 AUC) — Preprints 202603.0994
21. [P] Gaona et al. 2026, Frontiers in Applied Mathematics and Statistics (40-asset walk-forward)
22. [P] MDPI Applied Sciences 2025, confidence-threshold LOB study (order-book features dominate; 82.68% on 11.99% traded)
23. [V] Databento, CME Globex futures data — databento.com/futures ; [V] crypto websocket comparison — coinmarketcap.com/api
24. [V] Signed order flow & BTC price formation (contemporaneous R^2) — tradingstrategy-ai/docs ; OFI forward- R^2 decay explainer — electronictradinghub.com