

Predicting Kalshi BTC 15-Minute in Its Last 5 Minutes

Volume 1 — Internet Research: settlement mechanics, the core mathematics, features, fees, and live data sources

Every factual claim sourced; evidence graded A (peer-reviewed) / B (official primary) / C (practitioner measurement) / D (blog/marketing).

Headline: Kalshi settles on the CF Benchmarks BRTI — both strike and settlement are 60-second averages of 1-second prints.

Predicting Kalshi's BTC 15-Minute Binary (KXBTC15M) in Its Last 5 Minutes

External Knowledge, Mathematics, and Data Sources — A Complete Research Report

Research date: 21 October 2026 (America/Los_Angeles). All web sources below were read on this date unless a source itself carries an earlier publication date, which is stated. **Scope as commissioned:** Everything publicly available that could improve predicting the outcome of Kalshi series **KXBTC15M** — a binary that pays Yes (\$1) if Bitcoin finishes *at or above* a strike ("floor strike" / "price to beat") set at the window open, No (\$0) otherwise — **specifically during the last 5 minutes** of each 15-minute window (5:00 to 1:00 remaining, and, where the mathematics demands it, inside the final 60 seconds). The commissioning agent already holds a local dataset of every window at 1-minute resolution; this report is deliberately about *external* knowledge: settlement mechanics, mathematics, features, venues, and live data the local dataset does not contain. **Not financial advice.** This is a measurement and modelling report. Several of the best-measured findings below are that the obvious strategies *lose after fees*.

How to read the evidence grades used throughout

Grade	Meaning	Examples in this report
A — Peer-reviewed / formal working paper	Journal article or university working paper with stated data and methods	Bürge, Deng & Whelan on Kalshi (CESifo WP 12122); Cont, Kukanov & Stoikov on order-flow imbalance; Easley, López de Prado & O'Hara on VPIN; Nicola on intraday Bitcoin (Ledger); Takaishi on 1-minute Bitcoin (Physica A); Bergsli et al. on Bitcoin volatility forecasting
B — Primary / official	The contract, the index methodology, the exchange's own API documentation	Kalshi CRYPTO15M contract terms; CF Benchmarks RTI Methodology v17.0; Coinbase / Binance / Kraken / Deribit / Chainlink documentation
C — Practitioner measurement study	Non-peer-reviewed, but publishes sample size, period, method, and a reproducible public dataset	DegenHedge's 25,746-window settlement studies; OddsShopper's 2,200-window flip studies; Imbalance Labs' rank-IC study (vender — see caveat); Gravia's lead-lag measurement; BotForKalshi's self-measured taker study
D — Blog / marketing / folklore	Claims without auditable data, or secondary explainers	TradingView indicator scripts, CVD "divergence" academy articles, GitHub bot READMEs, whale-alert commentary

Where a Grade C/D claim is the *only* evidence for something, this report says so explicitly and treats it as a hypothesis to backtest, not a fact.

A note on a date discrepancy in the commission itself

The commission states the local dataset covers "every window since the series launched (Aug 13, 2026)". **No public source found confirms 13 August 2026 as the series launch, and several contradict it as a launch date.**

DegenHedge (Grade C, using Kalshi's own public API) analysed 25,746 settled KXBTC15M windows from **10 January to 8 October 2026**, and notes 2,451 *earlier* markets existed that used a point-price strike and were excluded because **on 10 January 2026 both the strike and the settlement became 60-second averages**

(<https://degenhedge.com/research/are-kalshi-bitcoin-15-minute-markets-rigged/>). A separate field note found Kalshi's historical API returns 21,560 KXBTC15M markets back to **10 December 2025**

(<https://pondletter.com/kalshi-free-history-api/>). Treat 13 August 2026 as the local dataset's start, not the series' launch — and note the practical consequence: **any backtest that mixes pre-10-January-2026 windows (point-price strike) with later windows (60-second-average strike) is mixing two different contracts.** The settlement mathematics in §1–§2 applies only to the averaging contract.

Executive Summary — the ten findings that matter most

- 1. Settlement is an average of averages.** Both the strike and the expiration value are the *simple average of 60 one-second CF Benchmarks Bitcoin Real-Time Index (BRTI) prints* — the strike from the 60 seconds before the window opens, the expiration value from the 60 seconds before it closes. Yes wins on `expiration_value >= floor_strike`

(ties pay Yes). Primary source: Kalshi CRYPTO15M terms (https://assets.kalshi.com/contract_terms/CRYPTO15M.pdf). This single fact changes the correct probability formula in the last minute by a factor of up to ~3 in standard-deviation terms (§1.5, §2.4).

2. **BRTI is an order-book index, not a trade-price index, and Binance is not in it.** The current methodology builds BRTI every second from the *consolidated order books* of its constituent exchanges, with an exponentially-weighted mid-price-volume curve. The current BRTI constituent list is Bitstamp, Coinbase, Kraken, Gemini, LMAX Digital, Bullish and Crypto.com — **not Binance, Bybit or OKX**

(<https://docs.cfbenchmarks.com/CME%20CF%20Constituent%20Exchanges.pdf>). Most secondary explainers listing "Bitstamp, Coinbase, Gemini, itBit, Kraken" are out of date. A model fed only Binance prices is modelling a *leader* of the settlement index, not the index.

3. **The correct last-5-minutes model is $\Phi(\text{lead} / (\sigma \cdot \sqrt{\tau_{\text{eff}}}))$ with $\tau_{\text{eff}} = \tau - 40$ seconds** for $\tau \geq 60$ s, and a *running-average* version inside the final minute (§2.4). At 60 s remaining the averaging rule cuts the effective standard deviation to 58% of the naive point-settlement value; at 30 s remaining, to 29% of it, *if you track the locked-in running average*.

4. **Kalshi's own price is already a very good forecast in the last 5 minutes.** In 15,183–15,414 windows logged ~3 minutes before close, the side won within ~1–2 points of its ask in almost every price band, usually slightly *under* it (<https://degenhedge.com/research/are-kalshi-bitcoin-15-minute-markets-rigged/>). At ~3 minutes, 99¢ asks lost 0.97% (price implies 1.00%), 98¢ lost 2.24%, 97¢ lost 2.34%, 90–94¢ lost 8.71% vs 7.81% implied (<https://degenhedge.com/research/can-you-lose-a-99-cent-kalshi-contract/>). **No price band in these data beat its own price after fees.**

5. **The fee is the strategy.** Kalshi taker fee = $\text{round_up}(0.07 \times C \times P \times (1-P))$. That is 1.75¢/contract at 50¢, 0.63¢ at 90¢, 0.0693¢ at 99¢ *in size* — but a **one-contract order at 99¢ rounds up to a full 1¢ fee, i.e. 100% of the possible profit, making its breakeven win rate 100%**. Breakeven win probability for a hold-to-settlement taker buy is simply $P + \text{fee_per_contract}$ (§4.4, Table 4.2).

6. **Buying the leading side as a taker has been measured, and it loses.** A practitioner who ran it with real money reports, over 6,298 windows, gross edge **+0.67¢/contract** for buying the leader, against an average taker fee of **1.55¢** — net negative at every entry time tested, including the last minute; the same source reports no maker fee on this series, making resting limit orders the only execution style that is not structurally negative, subject to adverse selection (Grade C, self-reported: <https://www.botforkalshi.com/blog/kalshi-bitcoin-bot>).

7. **The strongest published microstructure signal at exactly this horizon is order-book imbalance — and it is small.** On ~1.5M one-minute bars per asset (Hyperliquid perpetuals, vendor data), BTC top-5 order-book imbalance has rank-IC **+0.1129 for the next 1 minute**, decaying to +0.0508 (5-min forward), +0.0244 (15-min), +0.0105 (60-min); the extreme deciles differ by only **2.25 bps** in next-minute mid-return. The authors themselves state the raw signal "sits inside the taker fee" (<https://imbalancelabs.com/research/orderbook-imbalance-predictive-power>).

8. **Cross-exchange lead-lag is measured in milliseconds, not minutes.** In a 70.6-hour, five-venue perpetual-BBO measurement, Binance led on exchange timestamps (74 ms ahead of Kraken/Deribit, 43 ms ahead of Bybit, 2 ms ahead of OKX), but *delivery* to a server in Ireland reversed the usable order: Kraken arrived first (median delivery age 7.6 ms vs Binance 110.1 ms) and was first mover in 38.0% of 11,965 clean moves (<https://gravia.trade/learn/who-moves-bitcoin-first>). At a 1–5-minute horizon this matters only as "use the fastest-arriving constituent feed as your spot proxy and another venue to confirm", not as a multi-second forecasting edge.

9. **Polymarket's near-identical product was exploitable precisely because it did *not* average — and in August 2026 it was changed to copy Kalshi.** A Stanford/SMU study of Polymarket 5-minute BTC found 821 accounts made \$8.2M in likely-manipulated windows settled on single snapshots; Polymarket responded by moving 5-minute markets to a 30-second Chainlink TWAP and 15-minute/4-hour markets to a **60-second TWAP** (<https://www.coindesk.com/business/2026/08/07/how-a-five-second-trick-let-traders-drain-millions-from-polymarket>). Any "Polymarket edge" article written before August 2026 about last-second snapshot sniping is describing a dead mechanism.

10. **Favourite-longshot bias is real on Kalshi in aggregate, but it does *not* create a free lunch in KXBTC15M's last 5 minutes.** Bürgi, Deng & Whelan (Grade A, >300,000 Kalshi contracts) find cheap contracts win far less than

priced (a 5¢ contract winning ~2% loses ~60% of stake) and expensive contracts slightly more (95¢ winning ~98%), with average return **−20% pre-fee, −22% post-fee** across the exchange (https://www.ifo.de/sites/default/files/docbase/docs/cesifo1_wp12122.pdf). But the KXBTC15M-specific measurements in finding 4 show last-5-minutes favourites priced almost exactly right — the aggregate bias lives mostly in other categories, longer horizons, and taker flow.

1. Settlement Mechanics (most important)

1.1 The official contract language, verbatim in substance

Kalshi's contract terms file for this family, **CRYPTO15M** (Grade B, primary), states (https://assets.kalshi.com/contract_terms/CRYPTO15M.pdf):

> **Underlying**: "The Underlying for this Contract is the spot price of one \<cryptocurrency\> in U.S. dollars at \<time\>, according to a simple average of the CF \<cryptocurrency\> \<index\> for the 60 seconds prior to \<time\>, after Issuance and before \<date\>. Revisions to the Underlying made after Expiration will not be accounted for in determining the Expiration Value." > **Source Agency**: "The Source Agency is **CF Benchmarks**." > **\<index\>**: "refers to a specific CF Benchmarks Index (e.g. 'Bitcoin Real-Time Index') specified by the Exchange." > **Comparators**: "'at least X' means X or greater"; "Above X" is strictly greater, etc. > **Payout Criterion**: encompasses the Expiration Values where the index is \<above/below/between/exactly/at least\> \<price\>; **"If no data is available or incomplete on the Expiration Date at the Expiration Time, then affected strikes resolve to No."** > **Settlement Value**: \$1.00. **Position Accountability Level**: \$25,000 per strike, per Member. **Minimum tick**: \$0.01 in the CRYPTO15M file (the older CRYPTO file, https://assets.kalshi.com/contract_terms/CRYPTO.pdf, states \$0.001 — see §1.6 on tapered deci-cent pricing, which reconciles the two).

For KXBTC15M each individual market's own rules (exposed in Kalshi's API as `rules_primary/rules_secondary`, and as `floor_strike`, `expiration_value`, `result` on the settled record) instantiate the template as, in the wording quoted by two independent Grade C sources from the live rules:

> "At the last minute before expiration, **60 RTI prices are collected. The official and final value is the average of these prices, rounded to the nearest 2 decimal places.**" — quoted in <https://www.oddsshopper.com/articles/prediction-markets/are-kalshi-15-minute-bitcoin-markets-rigged/>

DegenHedge states the same rule and verifies it against outcomes (§1.4): "The market resolves Yes if the simple average of 60 seconds of BRTI before the close is at least the same 60-second average before the open" (<https://degenhedge.com/research/are-kalshi-bitcoin-15-minute-markets-rigged/>). PM Countdown independently notes Kalshi types the comparison as `greater_or_equal` in its systems, which is why a dead-flat window resolves Yes (<https://pmcountdown.com/markets/crypto>).

Therefore, precisely:

$$K \stackrel{\text{def}}{=} \frac{1}{60} \sum_{i=1}^{60} \text{BRTI}_{\lfloor t_{\text{open}} \rfloor + i - 1} \text{ (sec)} \quad \text{rounded to } \$0.01 \text{ for BTC}$$

$$S \stackrel{\text{def}}{=} \frac{1}{60} \sum_{i=1}^{60} \text{BRTI}_{\lfloor t_{\text{close}} \rfloor + i - 1} \text{ (sec)} \quad \text{rounded to } \$0.01$$

$$\text{Yes pays } \$1 \text{ iff } S \geq K; \quad \text{otherwise No pays } \$1.$$

Three consequences that are routinely gotten wrong in blog coverage:

- **It is not the BRR.** The CME CF Bitcoin *Reference Rate* (BRR) — a once-daily 4pm London fix built from a one-hour window partitioned into twelve 5-minute intervals, each scored by volume-weighted median trade price, then equally averaged — is a *different* CF Benchmarks product used for CME futures settlement. KXBTC15M uses the **Real-Time Index (BRTI)**, published every second. Secondary explainers that describe KXBTC15M settlement using BRR partition/median language are conflating the two products. The BRR partition description itself is documented in CME's analysis (<http://www.cmegroup.com/education/articles-and-reports/analysis-of-cme-cf-bitcoin-reference-rate>).

- **It is not a single print, and it is not any single exchange.** A one-second spike on one venue moves at most a few of the consolidated-book inputs to one of 60 BRTI prints, each of which is 1/60 of the settlement value.
- **Both ends are averaged.** A model that treats the strike as "the spot price at :00/:15/:30/:45" has the wrong strike by the difference between a point price and a 60-second mean — small in dollars, decisive in the ~1-in-22 windows that settle within \$5 of the strike (\$1.4).

1.2 How and when the strike is determined

- Windows close on the quarter hour (:00, :15, :30, :45), all Kalshi 15-minute crypto series on the same clock; there is no 5-minute Kalshi crypto series (KXBTC5M returns 404) (<https://pmcountdown.com/markets/crypto>).
- The strike for a window is the 60-second BRTI average in the minute *before* the window opens — which is **the identical 60 seconds that form the previous window's expiration value**. Empirically: in DegenHedge's full sample, **the strike equalled the previous window's settlement value to the cent in 25,281 of 25,285 consecutive pairs**; the four exceptions were two records carrying the 13 April 2026 $\times 10,000$ scale/typo error, one differing by \$0.07 (24 April 2026) and one by \$3.15 (19 May 2026), unexplained from public data (<https://degenhedge.com/research/are-kalshi-bitcoin-15-minute-markets-rigged/>).
- Practical implication: **at the moment a window opens, its strike is already fully known and is simply the previous window's published** `expiration_value` (field name `floor_strike` on the new record). There is no strike uncertainty to model in the last 5 minutes; all uncertainty is in the final 60-second average. Kalshi exposes `floor_strike`, `expiration_value` and `result` on every settled market, which is the correct grading data (<https://predictionmarketspicks.com/articles/how-kalshi-15-minute-markets-settle>).
- Settlement is fast: median window marked settled **10.9 seconds** after close, 75.2% within a minute (<https://degenhedge.com/research/are-kalshi-bitcoin-15-minute-markets-rigged/>). There is no post-close information window to trade in; the last trade opportunity is at expiry itself (Last Trading Time = Expiration Time in the template terms).

1.3 What BRTI actually is — the current CF Benchmarks methodology

The governing document is **CME CF Cryptocurrency Real Time Indices, Methodology Guide, Version 17.0, 21 September 2026** (Grade B:

<https://docs.cfbenchmarks.com/CME%20CF%20Real%20Time%20Indices%20Methodology.pdf>). Its relevant provisions, confirmed in the document text:

1. **Input is order books, not trades.** "CME CF Cryptocurrency Real Time Indices are calculated in real time based on the **Relevant Order Books of all Constituent Exchanges**." Each venue's full unmatched limit-order book is retrieved (Retrieval Time defined for both REST request/response and WebSocket feeds).
2. **Consolidation.** At the Effective Time, all venues' books are added to a joint list and aggregated into **one consolidated order book**. Size at any price level above the **order-size cap** is truncated to the cap (the cap is now dynamic; in the original 2016 methodology it was a static 100 BTC per order).
3. **Uncrossing (added v16.8, 29 June 2026).** While the consolidated best bid $\geq$ best ask, the lesser size at those two levels is removed from each side, until the book is uncrossed or one side is consumed. (This step matters for replication: crossed books across venues are common and are *not* naively averaged.)
4. **Curves.** From the consolidated book, cumulative **bid price–volume**, **ask price–volume**, **mid-price–volume** (= average of the bid and ask curves) and mid-spread–volume curves are computed at granularity set by a (now dynamic) spacing parameter.
5. **Exponential weighting.** The mid-price–volume curve is multiplied by the normalised probability density of an **exponential distribution** over cumulative volume — i.e. volume near the top of book counts far more than deep volume — and **BRTI is the mean of that weighted curve**. In the original 2016 methodology document the maximum volume integrated was 5,000 BTC and the design rationale is spelled out: order-book data is "forward-looking", uses the whole book rather than best bid/ask only (robustness), and exponential weighting emphasises "bid and ask prices closest to the global best bid and ask" (<https://www.cmegroup.com/trading/files/bitcoin-real-time-index-methodology.pdf>).

6. **Dissemination.** Approximately **every second**, every day including weekends/holidays; dissemination precision for BRTI **\$0.01**; a maximum dissemination of approximately every 200 ms is available on request for some indices per the specifications tables in the same document family (<https://docs.cfbenchmarks.com/CME%20CF%20Real%20Time%20Indices%20Methodology.pdf>).

7. **Contingencies.** Delayed data, erroneous books/prices, potentially-erroneous data (parameter tightened from 25% to 10%/5% for various indices in 2024–2026 revisions), expert judgement and calculation-failure procedures are all specified in §5 of the methodology. For Kalshi purposes the operative contingency is the contract's own: incomplete data at expiry ⇒ affected strikes resolve **No**.

Constituent exchanges for BRTI (current list). The CME CF Constituent Exchanges document (Grade B, version dated 15 June 2026: <https://docs.cfbenchmarks.com/CME%20CF%20Constituent%20Exchanges.pdf>) lists for the CME CF Bitcoin Real Time Index (ticker BRTI):

Constituent exchange	Added	Notes
Bitstamp	14 Nov 2016	
Coinbase	14 Nov 2016	Includes Coinbase, Coinbase Pro and Coinbase Prime flow per the document's footnote in the family
Kraken	14 Nov 2016	
Gemini	30 Aug 2019	
LMAX Digital	3 May 2021	
Bullish Exchange	30 Dec 2024	
Crypto.com	31 Mar 2025	
itBit (Paxos)	14 Nov 2016 / 1 Feb 2019 entries	The extract shows itBit suspended 25 Jan 2019 "(BRR only)" and a 27 Jul 2026 date in the BRR block — i.e. treat itBit as a <i>reference-rate</i> constituent historically, not a current BRTI constituent, unless CF Benchmarks' live list says otherwise
Bitfinex, OKCoin (HK)	14 Nov 2016	Both suspended April 2017

Not constituents: Binance, Bybit, OKX, Deribit, Bitstamp-only feeds, Chainlink, Pyth. CF Benchmarks' CEO has publicly explained Binance's exclusion on regulatory/integrity grounds despite its volume (reported in https://coinstats.app/zh/news/97938ccec4c5eebeec3c1f2252cc63346d00ed327f95eeca1df32634ffebf64c_How-Kraken%E2%80%99s-index-wizard-is-besting-Coinbase-and-Binance-in-the-%2458bn-Bitcoin-ETF-race/ — Grade D reporting of a Grade B fact; the constituent list itself is the authority).

Correction of a widespread error: several current explainers state BRTI "uses a volume-weighted median across Bitstamp, Coinbase, Gemini, itBit and Kraken, published once per second" (e.g. <https://bitcoincalculator.tools/learn/cf-benchmarks-brti-explained>). That mixes the **BRR's** volume-weighted-median-of-trades method with the **RTI's** order-book method and an outdated constituent list. For building a BRTI *proxy*, the order-book method is what must be replicated: consolidated top-of-book-weighted mid across the seven venues above. In practice nobody outside licensees replicates the exponential volume curve exactly; the achievable proxy is a size-capped, spread-filtered, multi-venue mid — and its measured error is quantified in §1.4.

1.4 How well can an external observer reproduce the settlement value? (measured)

Study (Grade C)	Proxy used	Measured agreement with Kalshi's official value
DegenHedge, 25,746 windows, Jan–Oct 2026 (https://degenhedge.com/research/are-kalshi-bitcoin-15-minute-markets-rigged/)	Coinbase/Kraken spot read at the close	Median absolute distance from official settlement value \$7.80 ; spot read pointed the wrong way on 4.37% of all windows and on 39.2% of windows finishing within \$5 of the strike . Authors state plainly: "We cannot recompute the BRTI."
PredictionMarketPicks, 3,320 windows, 24 Aug–27 Sep 2026 (https://predictionmarketpicks.com/articles/is-kalshi-bitcoin-rigged)	Their own independent recording of "the settlement index"	Matched Kalshi's published settle to a median of \$0.13 across 2,238 settles — i.e. a properly-built BRTI-like feed is ~60× closer than a raw spot read. (Their feed construction is not fully disclosed; treat the \$0.13 as "achievable with a real index feed", not with candles.)
DegenHedge strike study, summarised at https://degenhedge.com/research/	Official values	1 in 22 windows settled within \$5 of the strike; 1 in 117 within \$1 ; in windows that close, the side ahead ~3 minutes out lost 44.1% of the time.

Study (Grade C)	Proxy used	Measured agreement with Kalshi's official value
PredictionMarketPicks, same sample	Official values	8.8% of windows decided by less than \$10 of Bitcoin.
OddsShopper, 2,200 windows in 11 samples, Sep–Oct 2026 (https://www.oddsshopper.com/articles/prediction-markets/are-kalshi-15-minute-bitcoin-markets-rigged)	Kalshi 2-minutes-left trade price vs outcome	Windows finishing \$50+ from the line flipped late only 12 times in 1,195 (~1%) ; windows finishing inside \$50 flipped 216 in 1,005 (~21%) ; median finishing margin across samples \$33–\$83 depending on volatility regime.

Modelling translation: your probability model must be built on (a) the *published strike*, not a chart open; (b) a *multi-venue mid-price proxy* for BRTI, not a single exchange's last trade; and (c) an explicit BRTI-proxy error term. A \$7.80 median proxy error (single-spot) is larger than the entire finishing margin in ~1 window in 12. Reducing proxy error is plausibly worth more than any directional feature in §3.

1.5 What the 60-second average implies mathematically in the last 1–2 minutes

This is the single highest-value section of the report. Full derivation in §2.4; the operational facts:

Let τ = seconds remaining until expiry, X_t = current BRTI (proxy), K = strike, σ = BRTI volatility in **dollars per $\sqrt{\text{second}}$** .

- $\tau \geq 60$ (i.e. **5:00 to 1:00 remaining — most of the commissioned window**): none of the settlement average is locked yet, but the average will be taken over a window that *starts* $\tau-60$ s from now. The settlement value's conditional standard deviation is $\sigma \cdot \sqrt{\tau - 40}$, not $\sigma \cdot \sqrt{\tau}$. At $\tau=300$ the reduction is only 7% ($\sqrt{260}/\sqrt{300}$); at $\tau=120$ it is 18%; at $\tau=60$ it is 42%.
- $\tau < 60$ (**final minute**): a fraction $(60-\tau)/60$ of the settlement average is **already observed and irreversible** if you are recording BRTI each second. Define the running sum of the elapsed settlement-minute prints. The conditional expectation of the settlement value is

$E[\text{settlement value}] = \frac{1}{60} \sum (\text{sum of elapsed prints}) + \frac{\tau}{60} X_t$ and the remaining standard deviation collapses as $\sigma \cdot \sqrt{3/2} / (60\sqrt{3})$.

- **How much is "known" — quantified:**

Seconds left τ	Share of settlement prints already locked	Remaining SD ($\times \sigma$)	Remaining SD vs naive point-settlement $\sigma\sqrt{\tau}$	Move in remaining prints' average needed to erase a \$1 gap in $E[S]-K$
60	0%	4.472 σ	58%	\$1.00
45	25%	2.905 σ	43%	\$1.33
30	50%	1.581 σ	29%	\$2.00
15	75%	0.559 σ	14%	\$4.00
10	83%	0.304 σ	10%	\$6.00
5	92%	0.108 σ	5%	\$12.00

The last column is the intuition OddsShopper states verbally: "The final 15 seconds are only a quarter of the average, so closing a gap in the running average takes a move about four times the size of the gap. You can watch Bitcoin rip \$200 through the line in the last stretch and still lose, because 45 of the 60 samples were already locked in on the other side" (<https://www.oddsshopper.com/articles/prediction-markets/are-kalshi-15-minute-bitcoin-markets-rigged/>). Conversely — and this is the tradable side of the same fact — **a contract can be a near-lock on the running average while the live spot price is visibly on the other side of the strike**, and the market (dominated in the final minute by programs repricing off the index) will usually, but not instantaneously, reflect that.

- **Speed in the final minute is a documented structural feature, not colour.** In the OddsShopper sample's most-contested window, the final minute printed 864–1,281 fills per 10-second bucket versus 92 fills in the *entire* final minute of a decided window the night before — "Programs reading the public index reprice the contract before you can" (same source). DegenHedge's feed business is built on the same observation and still discloses its model side won *below* breakeven (§4.3). **Any last-60-seconds strategy that requires human-speed execution should be assumed dead on arrival; the averaging edge in the table above is real arithmetic, but capturing it is a latency competition.**

1.6 Pricing granularity near the extremes

Kalshi's records for these markets list a **"tapered deci-cent" price structure: 0.1¢ steps below 10¢ and above 90¢, whole cents in between**; in Kalshi's public 1-minute history Mar–Oct 2026, 83.8% of Yes asks $\geq 90¢$ were not whole cents (<https://degenhedge.com/research/can-you-lose-a-99-cent-kalshi-contract/>). Consequences: (a) "99¢" in older studies means 99.0–99.9¢; at 99.9¢ a win pays 0.1¢ gross. (b) The template minimum tick of \$0.01 (CRYPTO15M file) and \$0.001 (CRYPTO file) are both "true" for different price regions. (c) Any backtest on 1-minute candles that rounds prices to whole cents misprices exactly the 90¢+ region the commission asks about.

2. The Core Math

2.1 The baseline model: driftless Brownian motion and the probit formula

Assumptions (state them in any backtest): over horizons of 1–5 minutes, (i) drift is negligible — even a strong 30%-annualised drift contributes $\mu \cdot \tau \approx 0.3 \times (300/31,536,000) \approx 0.0003\%$ over 5 minutes, versus a diffusion move of order $\sigma\sqrt{\tau} \approx 0.05\text{--}0.15\%$; (ii) log-price (equivalently, for these tiny moves, price) follows a Brownian motion with locally constant volatility σ ; (iii) increments are independent and Gaussian. Then, for a contract settling on the *point* price at expiry:

$$P(\text{Yes} \mid \text{mathcal F}_t) = P(X_T \geq K) = \Phi\left(\frac{X_t - K}{\sigma\sqrt{\tau}}\right) = \Phi\left(\frac{\text{lead}}{\sigma\sqrt{\tau}}\right),$$

where Φ is the standard normal CDF, lead is in the same units as $\sigma\sqrt{\tau}$ (use dollars throughout for BTC: at BTC $\approx \$82,000$, 1 bp = \$8.20). This is the continuous-time limit of the binary-option / Black–Scholes digital formula with zero rates and zero drift, and it is the formula implemented, e.g., by published TradingView fair-value scripts for these markets ("models the 15-minute candle ... as a binary derivative subject to Brownian Bridge dynamics, calculating probability from Time to Expiry and Volatility" via a normal CDF — Grade D implementation, correct textbook formula: <https://www.tradingview.com/script/bKp6f7yZ-Polymarket-15m-BTC-Probability-10x/>).

Terminology correction: the commission calls this a "Brownian-bridge" formula. Strictly, a Brownian *bridge* is a Brownian motion conditioned on its endpoint; the formula above is the *unconditioned* first-passage/terminal-value probability of a Brownian motion (a martingale). The bridge becomes literally relevant in §2.4, where conditioning on the already-observed part of the settlement average pins part of the path. The distinction matters only for which variance to use — which is the whole game at this horizon.

Worked calibration of magnitudes. Take σ from 1-minute data (§2.2): suppose the last hour's 1-minute BRTI-proxy returns give $\sigma_{\text{■■■■}} = \$25$ (≈ 3.0 bps at \$82,000 — a calm-to-typical minute in Oct 2026 conditions). Then $\sigma = 25/\sqrt{60} = \mathbf{\$3.23 \text{ per } \sqrt{\text{second}}}$, and:

τ	$\sigma\sqrt{\tau}$ eff (averaging-corrected, §2.4)	Lead \$25 $\rightarrow$ P(Yes)	Lead \$50 $\rightarrow$ P(Yes)	Lead \$100 $\rightarrow$ P(Yes)
300 s	\$52.0	68.5%	83.2%	97.3%
180 s	\$38.2	74.4%	90.5%	99.6%
120 s	\$28.9	80.7%	95.8%	99.97%
60 s	\$14.4	95.9%	99.97%	$\approx 100\%$

Read this table against §1.4's empirical flip rates: a \$50 lead with 2 minutes left is *not* a lock — the model says $\sim 4\text{--}16\%$ loss probability depending on τ and σ regime, and the measured data agree (windows finishing inside \$50 flipped $\sim 21\%$ of the time from their 2-minute price side). **Most retail "it was at 99¢ and lost" stories are this table being ignored: at $\tau=120\text{--}180$ s, a 99¢ price requires roughly a \$75–\$95 lead at $\sigma_{\text{■■■■}}=\$25$, and leads that size are routinely spent by one volatile minute.**

2.2 Estimating σ at minute scale

σ is the only free parameter in the baseline model, and at 1–5-minute horizons its estimation error dominates every other refinement. The estimators, in the order this report recommends testing them:

(a) Close-to-close realised volatility (the baseline)

For 1-minute log returns $\$r_i = \ln(X_i/X_{i-1})$ over a lookback of n minutes:

$\hat{\sigma}_{\text{1min}}^2 = \frac{1}{n-1} \sum_{i=1}^n (r_i - \bar{r})^2 \approx \frac{1}{n} \sum r_i^2 \quad (\bar{r} \approx 0 \text{ at this scale}), \quad \hat{\sigma}_{\text{per-sec}} = \frac{X_t \cdot \hat{\sigma}_{\text{1min}}}{\sqrt{60}}$

Choices that must be backtested, not assumed: lookback n (15 / 30 / 60 / 240 minutes are the sensible grid); and **time-of-day deseasonalisation** — Bitcoin intraday volatility has a strong human-activity cycle (documented at 5-minute frequency in the literature surveyed at https://www.researchgate.net/figure/Descriptive-statistics-of-BTC-e-at-the-5-minute-frequency_tbl1_319584824), and DegenHedge's time-of-day study found the 10:00 ET hour moves most and Saturdays are calmest for this exact series (summarised at <https://degenhedge.com/research/>). A σ estimated at 03:00 ET and used at 10:00 ET is a miscalibration, not a forecast error.

(b) EWMA (RiskMetrics-style) — the practical favourite at this horizon

$\hat{\sigma}_t^2 = \lambda \hat{\sigma}_{t-1}^2 + (1-\lambda) r_{t-1}^2$

updated every minute (or every 5–15 seconds on sub-minute returns). λ per minute in the 0.90–0.97 range gives half-lives of ~7–23 minutes — matched to volatility clustering (§2.3). EWMA's advantage over a flat lookback is that it reacts within 1–2 minutes to a vol burst, which is exactly when the probit model's errors are largest in dollars.

(c) Range-based estimators: Parkinson and Garman–Klass (use the OHLC you already have)

Per bar of length Δ (here $\Delta = 1$ minute), with H, L, O, C the bar's high/low/open/close:

$\text{Parkinson (1980):} \quad \hat{\sigma}_P^2 = \frac{1}{4 \ln 2} \left(\ln \frac{H}{L} \right)^2 \approx 0.361 \cdot \left(\ln \frac{H}{L} \right)^2$

$\text{Garman–Klass (1980):} \quad \hat{\sigma}_{GK}^2 = \frac{1}{2} \left(\ln \frac{H}{L} \right)^2 - \frac{1}{2} \left(\ln \frac{C}{O} \right)^2 \approx 0.5 \cdot \left(\ln \frac{H}{L} \right)^2 - 0.386 \cdot \left(\ln \frac{C}{O} \right)^2$

averaged over n bars and divided by Δ to get variance per unit time. Under idealised continuous Brownian motion with no drift, Parkinson is ~5× and Garman–Klass ~7–8× more efficient (lower estimator variance) than close-to-close for the same number of bars — the standard textbook result, and the Garman–Klass form is the one used as a volatility proxy in peer-reviewed Bitcoin forecasting work (<https://doi.org/10.3390/jrfm14070337>). **Caveats at 1-minute scale:** (i) H/L from *trade* prices embed bid–ask bounce, biasing range estimators *up* when the spread is a non-trivial fraction of the minute's range (use mid-price H/L if available); (ii) with only 1-minute bars you cannot see intra-minute path — GK on 1-minute bars estimates the minute's variance well, but the *settlement-minute* calculation in §2.4 wants σ per second, which is just the same variance rescaled only if the Brownian assumption holds within the minute.

(d) Realised variance from sub-minute sampling, and why not 1-second

$RV_{[t,t+h]} = \sum_{j=1}^n r_j^2 \quad \text{over sub-intervals } j$

Theory (Andersen–Bollerslev) says $RV \rightarrow$ integrated variance as sampling refines; practice says microstructure noise (§2.6) dominates below a threshold. The Bitcoin HAR literature explicitly adopts **5-minute returns for RV construction because finer sampling degrades fit and forecasting accuracy due to noise**, citing Liu et al. (2015)'s cross-asset result that 5-minute RV generally outperforms other frequencies (<https://WWW.MDPI.COM/2227-7072/14/4/81>). For this application the useful compromise is: estimate σ for the *probit denominator* from 1-minute returns (matching your dataset) or 15–30-second mid-price returns with a noise-robust estimator (two-scale RV / pre-averaging — standard references in §2.6), and *validate* the choice by calibration (predicted vs realised win rates in z -bins), not by $RV R^2$.

(e) What not to use as σ

- **Daily σ rescaled by $\sqrt{\text{time}}$.** Intraday vol seasonality + clustering make a daily number wrong by 2–3× at specific hours.
- **DVOL directly.** Deribit DVOL is a 30-day implied vol (§3.6). Its per-minute rescaling ($DVOL / \sqrt{365 \cdot 24 \cdot 60}$) is a *prior* for the day's regime, not a minute-scale σ . On 20 Oct 2026 DVOL stood at 38.28% annualised while 14-day realised was 20.2% by one measure and 25.5% by another — a variance-risk-premium gap of ~13–18 points that would inflate every probit denominator by ~1.5–1.9× if used raw (<https://www.tokenpost.com/news/investing/36192>).

2.3 Volatility clustering and GARCH/HAR at intraday scale

Stylised fact (Grade A). Nicola, "On the Intraday Behavior of Bitcoin", *Ledger* 6 (2021) 58–80: autocorrelation of **squared** 1-minute returns = **0.3639** and of absolute returns = 0.3459; at 5 minutes, 0.2592 / 0.3899; at 15 minutes, 0.2163 / 0.3812 — "strong evidence of volatility clustering at all timeframes considered" (<https://www.readkong.com/page/on-the-intraday-behavior-of-bitcoin-ledger-journal-4532086> — journal text via a reprint host; the article is *Ledger* Vol. 6, DOI 10.5195/LEDGER.2021.213). Volatility is therefore *forecastable* at exactly the horizon of interest, even when direction is not: a large last minute predicts a large next minute.

GARCH(1,1) at minute scale:

$$\sigma_t^2 = \omega + \alpha \epsilon_{t-1}^2 + \beta \sigma_{t-1}^2, \text{ where } \epsilon_t = r_t - \mu$$

with $\alpha + \beta$ typically 0.95–0.99 on 1-minute BTC (persistence), $\alpha \approx 0.05$ –0.15. Student-t innovations ($\nu \approx 3$ –6) are standard for BTC and matter again in §2.5. For *forecasting* comparisons on Bitcoin specifically, Bergsli, Lind, Molnár & Polasik (2021, *Research in International Business and Finance*, DOI 10.1016/j.ribaf.2021.101540) find EGARCH and APARCH best among GARCH variants, and **HAR models on realised variance beat daily-data GARCH, with the superiority strongest for short-term forecasts** (<https://ntnuopen.ntnu.no/ntnu-xmlui/handle/11250/3003405?show=full>).

HAR, adapted intraday (the testable form for this problem). Corsi's HAR-RV regresses future RV on lagged RV at multiple horizons. Translated from days to the last-5-minutes problem — forecast the RV of the *next* 1–5 minutes from trailing blocks:

$$\widehat{RV}_{t \rightarrow t+5} = c + \beta_1 RV_{t-5:t} + \beta_2 RV_{t-25:t} + \beta_3 RV_{t-125:t} + \epsilon_t$$

where $RV_{t-k:t}$ denotes average per-minute RV over the trailing block. Fit by OLS on your local windows (this is a regression your 1-minute dataset *can* support, unlike most of §3). Its output replaces $\hat{\sigma}$ in the probit formula.

This is a confidence/width model, not a direction model — see §3.8.

2.4 Settlement averaging changes the effective variance — full derivation

Let the settlement window be $[T-60, T]$, $S = \frac{1}{60} \int_{T-60}^T X_s ds$ (the discrete sum over 60 one-second prints is the same to $O(1/60)$), X a Brownian motion with volatility σ ($\$/\sqrt{s}$), current time t , $\tau = T - t$.

Case $\tau \geq 60$. Write $a = \tau - 60$ (time until the averaging window starts). Then

$$S - X_t = \frac{\sigma}{60} \int_a^{a+60} W_u du, \text{ where } \text{Var} = \frac{\sigma^2}{60^2} \int_a^{a+60} \int_a^{a+60} \min(u, v) du dv = \frac{\sigma^2}{60^2} \text{Big}(60^2 a + \frac{60^3}{3}) = \sigma^2 \frac{(a+20)^2 - 40^2}{60}$$

(The identity uses $\min(a+x, a+y) = a + \min(x, y)$ and $\int_0^L \int_0^L \min(x, y) dx dy = L^3/3$ — the same $t^3/3$ variance as for integrated Brownian motion, cf.

<https://www.quantopia.net/interview-questions-vii-integrated-brownian-motion/>, and the same 1/3 factor that appears in the variance of time-averaged rates generally, e.g. <https://www.markowitzoptimizer.pro/wiki/CompoundRateVariance>.) Hence:

$$\boxed{P(\text{Yes})} = \Phi\left(\frac{X_t - K}{\sigma \sqrt{\tau - 40}}\right), \text{ where } \tau \geq 60$$

Case $\tau < 60$. Let $A_{\text{obs}} = \sum$ of the BRTI prints already taken in the settlement minute (count $60 - \tau$), each fully known. Then $E[S | \mathcal{F}_t] = \frac{A_{\text{obs}}}{60} + \frac{\tau}{60} X_t$ and

$$\text{Var}(S | \mathcal{F}_t) = \frac{\sigma^2}{60^2} \tau = \frac{\sigma^2}{60} \frac{\tau^3}{3} = \frac{\sigma^2}{60} \frac{\tau^3}{10800}, \text{ where } \boxed{P(\text{Yes})} = \Phi\left(\frac{E[S | \mathcal{F}_t] - K}{\sigma \sqrt{\tau/3}}\right), \text{ where } \tau < 60$$

The two branches agree at $\tau=60$: $\sigma\sqrt{20} = \sigma \cdot 60^{3/2} / (60\sqrt{3}) = 4.472\sigma$. ✓

Equivalent "required move" form (often better for live display and backtests). Inside the final minute, Yes needs the *average of the remaining prints* to exceed

$$X_{\text{req}} = \frac{60K - A_{\text{obs}}}{\tau}$$

Define the **required-move lead** $\ell_{\text{req}} = X_t - X_{\text{req}}$. Note $\ell_{\text{req}} = \frac{60}{\tau} \text{Big}(E[S | \mathcal{F}_t] - K)$, so $P = \Phi\left(\ell_{\text{req}} \sqrt{3} / \sigma \sqrt{\tau}\right)$ — i.e. *versus a point-settled contract, the same displayed lead is worth $\sqrt{3} \approx 1.73\times$ more in z-units at every instant of the final minute, *on top of the shrinking of τ .* This is

the precise sense in which "part of the settlement is locked in": not merely that $(60-\tau)/60$ of the weight is known, but that the unknown part's variance falls as τ^3 , not τ .

Strike-side averaging needs no parallel correction in the last 5 minutes: K is a realised constant (§1.2). It *does* matter for anyone reconstructing K historically from point prices — don't.

Rounding. S and K are rounded to \$0.01. At BTC $\approx$ \$82,000 this is 0.0012 bps — utterly negligible for probability, except that it makes exact ties ($\$S=\$K \Rightarrow \text{Yes}$) measure-zero in practice: DegenHedge found **zero exact ties in 25,317 windows**, but 1 in 117 within \$1 (<https://degenhedge.com/research/are-kalshi-bitcoin-15-minute-markets-rigged/>, strike study summarised at <https://degenhedge.com/research/>).

2.5 Jump risk and fat tails at 1–5-minute horizons

The Gaussian probit understates the probability of large short-horizon moves for BTC. Evidence:

- **1-minute returns are severely fat-tailed (Grade A).** Takaishi (2018, *Physica A* 506:507–519), on 1-minute Bitstamp BTC: "the 1-min return distribution is fat-tailed, and kurtosis largely deviates from the Gaussian expectation... convergence to [Gaussian] is very slow" (<https://ideas.repec.org/a/eee/phsmap/v506y2018icp507-519.html>). A multifractal study of 1-minute Bitstamp returns 2013–2020 reports sample min/max 1-minute returns of **–30.65% / +35.29%** (crash-era data, but the leptokurtosis result is robust across sub-periods) and concludes multifractality comes "mainly from fat-tails" (<https://www.mdpi.com/2227-7390/9/17/2088/>).
- **Jumps are a distinct, testable component (Grade A).** Scaillet, Treccani & Trevisan's high-frequency jump analysis of Bitcoin models log-price as diffusion + jumps and *observed* price as efficient price + microstructure noise, finding significantly negative 1–3-lag autocorrelations at tick frequency from the noise (https://www.academia.edu/108497250/High_Frequency_Jump_Analysis_of_the_Bitcoin_Market). Merton's (1976) jump-diffusion is the canonical parametric form: $dS/S = (\mu - \lambda \kappa)dt + \sigma dW + dJ$, with Poisson intensity λ and random jump sizes; it exists precisely to generate the excess kurtosis a pure diffusion cannot at short horizons (summary: <https://www.nottldr.com/FinanceArchitect/jump-diffusion-models-capturing-the-fat-tails-markets-act-ually-exhibit-0gw3y3o>).

Operational corrections to the probit, in increasing complexity:

1. **Student-t probit (recommended first test).** Replace Φ with the CDF of a Student-t with ν degrees of freedom, rescaled to unit variance:

$$P = T_{\nu} \left(\frac{\text{lead}}{\sigma \sqrt{\tau_{\text{eff}}}} \cdot \sqrt{\frac{\nu-2}{\nu}} \right)$$
 With $\nu = 4$, a nominal $z = 3.0$ (Gaussian $P = 99.87\%$) becomes $\approx 98.9\%$ — i.e. the model stops calling things "locks" that empirically lose $\sim 1\%$ of the time. Compare this against the measured 99¢ loss rate of 0.97% (§4.2): the market's own prices behave much more like the t-version than the Gaussian version at the extremes.

2. **Jump-mixture probit.** With jump intensity λ (per second), jump mean μ_J and jump SD σ_J (in dollars):

$$P = \sum_{k=0}^{\infty} e^{-\lambda \tau_{\text{eff}}} \frac{(\lambda \tau_{\text{eff}})^k}{k!}; \Phi \left(\frac{\text{lead} + \mu_J}{\sqrt{\sigma^2 \tau_{\text{eff}} + \lambda \sigma_J^2}} \right)$$
 Truncate at $k = 3-4$. Calibrate λ , σ_J from your own 1-minute data by counting $|r| > 4\sigma$ minutes (a Lee–Mykland-style threshold in spirit). Even λ calibrated to a few jumps per day materially fattens the 90–99¢ region and barely moves the 40–60¢ region — which is the empirically correct shape.

3. **Empirical / isotonic calibration on z .** Compute z for every historical window at fixed τ snapshots (300/240/180/120/60 s), bin by z , and map $z \rightarrow$ realised win frequency with isotonic regression. This absorbs fat tails, σ -estimation bias, and BRTI-proxy error simultaneously. **Given the commissioning agent already has every window at 1-minute resolution, this is feasible immediately and is the benchmark every parametric version must beat out-of-sample.**

What fat tails do *not* justify: widening σ uniformly. That overstates uncertainty near the strike (where the Gaussian is roughly right and the market is most liquid) to fit the tails. The t / mixture / isotonic forms widen *tails relative to centre*, which is the actual empirical pattern.

2.6 Bid–ask bounce and microstructure noise in 1-minute prices

Roll (1984) mechanics. If the observed price is the efficient price plus a bounce between bid and ask (half-spread $c = s/2$), observed returns acquire negative first-order autocovariance $-\lambda c^2/2$ (i.e. $-\lambda s^2/4$), and realised variance from n observations is biased upward by $\lambda \approx 2n\lambda c^2$ where λc^2 is the noise variance. The crypto volatility literature states the same point: "The noise arises from bid-ask bounces, staleness, rounding errors... [and] reveals itself as autocorrelations in high-frequency returns" (<https://arxiv.org/html/2109.12142v2>).

Measured for BTC. Nicola (Grade A) finds first-order autocorrelation of **1-minute returns** = **−0.1445**, significant, and interprets it as systematic mean reversion (see §2.7 for the competing interpretation). At tick frequency Scaillet et al. find negative autocorrelations at lags 1–3 attributable to noise with serial dependence. **The two interpretations are not fully separable from trade prices alone** — which is itself the actionable point:

- **Use mid-prices, not last-trade prices, for both lead and σ .** A lead computed from a last trade that happened to print at the bid understates the true lead by $\sim$ half the BRTI-constituent spread, and an RV computed from trade prices overstates σ . Both errors push z toward zero — i.e. toward 50% — which, if the market is pricing off mids, manufactures phantom "edges" on the favourite side in backtests.
- **The backtest trap is quantified in the pairs-trading literature (Grade A):** executing a microstructure signal with *no* delay showed 20.5 bps/round-trip; with a realistic one-period delay, 3.8 bps; by five periods the alpha had vanished (<https://www.mdpi.com/411108>). Any last-5-minutes backtest that signals and fills on the *same* 1-minute bar's close is assuming the 20.5 bp version of reality.
- BRTI itself mitigates bounce *for settlement* (it is built from order books/mids, §1.3), so the settlement series is smoother than any single venue's trade series. Your proxy should mimic that: consolidated mid, not consolidated last-trade.

2.7 Mean reversion vs momentum at 1–5 minutes — what the evidence actually says

This is the most over-claimed topic in the practitioner literature. The defensible summary:

Claim	Best evidence	Verdict
1-minute BTC returns are negatively autocorrelated	Nicola (Grade A): $\rho = -0.1445$ at 1 min (2015–2018 Bitstamp-era data). HarbourFrontQuant's summary of the same paper notes strategies fading 5-min moves $> 2\sigma$ were profitable <i>in that sample</i> , with explicit lookback-bias caveats (https://harbourfrontquant.substack.com/p/trading-strategies-based-on-autocorrelation-properties)	Real in-sample, old sample, and confounded with bid-ask bounce (§2.6). A -0.14 autocorrelation on <i>trade</i> prices can be produced by bounce alone. On <i>mid</i> prices the true reversion component is smaller. Test on mids before believing any of it.
Window-to-window (15-min) results revert	DegenHedge (Grade C, 25,746 windows, 2026): Up-after-Up 48.00% vs Up-after-Down 51.61% ($z = -5.80$); after 4 identical results the 5th matched only 44.6%. But: the effect halved in the second half of the sample (gap 5.4 pts $\rightarrow$ 1.9 pts), Kalshi's price 1 minute into the window already reflected $\sim$ half of it, and at the ask the favoured side cost <i>more</i> than it won (Yes ask 51.41¢ after a Down window vs 51.24% win rate) (https://degenhedge.com/research/are-kalshi-bitcoin-15-minute-markets-rigged/)	Real, decaying, and already priced + smaller than the fee. Not a tradable last-5-minutes signal; at most a tiny prior tilt at window open.
Intraday momentum <i>and</i> reversal both appear at 30–60-min scales	Xu et al., intraday predictability in crypto using 5-min data (2013–2021): of efficient half-hour pairs, 12 show momentum, 21 show reversal (https://www.researchgate.net/profile/Yahua-Xu/publication/361202793_Intraday_return_predictability_in_the_cryptocurrency_markets_momentum_reversal_or_both/links/62a83068416ec50bdb24a18f/Intraday-return-predictability-in-the-cryptocurrency-markets-Momentum-reversal-or-both.pdf)	Mixed by time-of-day and pair — no stable unconditional sign. Do not import a generic "BTC momenta at 5 min" prior.
Returns concentrate at the turn of the 15-minute candle	Shanaev et al. (Grade A, <i>Heliyon</i>): +0.58 bps/min disproportionately in minutes 0/15/30/45, other minutes negative on average, since mid-late 2020, across 7 exchanges, robust out-of-sample (https://www.sciencedirect.com/science/article/pii/S2405844023014433)	A genuine timing anomaly at exactly this series' boundaries — worth including as a <i>minute-of-window fixed effect</i> in calibration, and a warning that the settlement minute's first seconds are not drawn from the same distribution as a random minute. 0.58 bp = \$0.48 at \$82,000 — small versus $\sigma_{\text{BTC}} \approx \$15\text{--}40$, so it shifts z only in near-tie windows.
"Momentum in the last 5 minutes predicts the close" (blog form)	No Grade A/C source found measuring a <i>mid-price, cost-adjusted</i> 1–5-min momentum effect in BTC that survives spread+fees. The measured trade-flow effects (§3.2) are contemporaneous-to-1-minute, not 5-minute trend continuation.	Unsupported as stated. Treat any such backtest as bid-ask bounce or lookahead until proven otherwise on mids with a 1-bar execution delay.

3. Predictive Features at 1–5-Minute Horizons

For each feature: definition precise enough to code, best evidence with effect size, horizon, and the fees/spread verdict. A consolidated comparison table is at §3.9.

3.1 (a) Order-book imbalance / depth imbalance

Definitions. For the top k levels of a (preferably consolidated, BRTI-constituent) book at time t :

$$OBI_k = \frac{\sum_{i=1}^k p^{bid}_i q^{bid}_i - \sum_{i=1}^k p^{ask}_i q^{ask}_i}{\sum_{i=1}^k p^{bid}_i q^{bid}_i + \sum_{i=1}^k p^{ask}_i q^{ask}_i} \in [-1, 1]$$

(the notional form used by Imbalance Labs for $k=5$), and Cont–Kukanov–Stoikov **order-flow imbalance (OFI)** over an interval, which accumulates *changes* at the best bid/ask (additions, cancellations, trades) rather than the static snapshot:

$$OFI_{[t, t+\Delta]} = \sum \{\text{events}\} \cdot \Delta \cdot \text{bid size at best} - \Delta \cdot \text{ask size at best} \cdot \text{signed}, \quad \Delta \cdot P \approx \frac{OFI}{2D},$$

where D is average market depth — the Grade A result (Cont, Kukanov & Stoikov 2014, *Journal of Financial Econometrics*, on NYSE stocks) is that price changes over short intervals are **linear in OFI with slope inversely proportional to depth**, robust across stocks and timescales, and stronger than the volume–price relation (summary: <https://electronicTradingHub.com/leveraging-limit-order-book-imbalances-for-profitable-trading-a-deep-dive-into-research-and-practical-tools/>). Multi-level extensions add a reported 5–15 points of R^2 for 1-second returns beyond level-1 in equity-style analyses (<https://delphicalpha.substack.com/p/reference-guides-order-book-dynamics> — Grade D summary of Grade A literature).

Crypto effect size (Grade C, vendor). Imbalance Labs, Hyperliquid L2, ~1.5M 1-minute snapshots/asset (<https://imbalancelabs.com/research/orderbook-imbalance-predictive-power/>):

- BTC rank-IC of OBI vs forward mid return: **+0.1129 @1 min**, +0.0508 @5 min, +0.0244 @15 min, +0.0105 @60 min. (Note: these are ICs by *forward horizon* on 1-minute bars.)
- By *bar resolution* (8 most liquid assets averaged): IC +0.0498 (1-min bars) → +0.0224 (5-min) → +0.0113 (15-min) → +0.0030 (60-min): "the signal roughly halves at every step down in resolution... a half-life measured in minutes."
- BTC decile means, next-minute mid-return: **−1.129 bps (most ask-heavy) to +1.118 bps (most bid-heavy)**, monotonic across all 10 deciles — total spread **2.25 bps/min**. At BTC \$82,000, 2.25 bps ≈ **\$18.45** of expected next-minute mid move between extreme deciles; the *typical* (interquartile) imbalance carries a fraction of that.
- Vendor conflict-of-interest caveat: the study sells the dataset it measures. The monotonic deciles and the decay pattern are, however, exactly what the Grade A OFI literature predicts, and the authors' own conclusion is bearish on naive use: "the raw signal sits inside the taker fee; harvesting it requires real execution — maker placement, queue position, adverse-selection control."

Translation to the binary. OBI does not move $P(\text{Yes})$ directly; it shifts the *expected path* of X over the next 1–2 minutes. The clean way to use it in the probit frame: add a short-horizon drift term $\mu_{OBI} \cdot \tau_{eff}$ to the lead, with μ_{OBI} calibrated by regressing forward 1-minute mid returns on OBI in bps (your calibration, on constituent-venue books, will differ from Hyperliquid perps). Even at the extreme decile, the shift over the *remaining* window is ~\$10–20 — material only when $|z| \gg 1$, which is also where the fee (in probability points) is largest. **Verdict: real, Grade A mechanism + Grade C crypto size; horizon 1–2 minutes; does NOT survive as a taker strategy on its own; plausible as (i) a tie-breaker adjustment to fair value and (ii) a maker quoting input.**

Spoofing caveat. Displayed depth on centralised venues can be cancelled; the Imbalance Labs study argues its DEX depth is "genuine committed capital, not phantom spoofing" — the inverse warning applies to CEX books. Prefer OFI (which counts events, including cancels with the correct sign) over static OBI snapshots, and time-weight imbalance over the last 15–60 s rather than snapshotting once.

3.2 (b) Trade flow / taker buy–sell imbalance / CVD

Definitions. For a trailing window W (30 s, 60 s, 300 s are the testable choices), over trades with signed volume v_j (+ if taker-buy, – if taker-sell; Binance's `isBuyerMaker` flag and Coinbase's `side/maker` fields give this directly, no Lee–Ready inference needed):

$$\text{TFI}_W = \frac{\sum_{j \in W} v_j}{\sum_{j \in W} |v_j|} \in [-1, 1], \quad \text{CVD}_t = \sum_{j \leq t} v_j, \quad \Delta \text{CVD}_W = \text{CVD}_t - \text{CVD}_{t-W} \text{ (not the level).}$$

Evidence. The strongest honest statement in the literature is about *contemporaneous* explanatory power: signed order flow explains a large share of BTC price *variation* (a practitioner replication of the peer-reviewed price-formation literature reports ~60% of daily return variation explained by signed flow on Binance futures, down from ~85% in earlier samples — Grade D report of Grade A work:

<https://github.com/tradingstrategy-ai/docs/raw/refs/heads/master/articles/bitcoin-price-formation-signed-order-flow.pdf> via the thread summary in search results). Contemporaneous explanation is **not** forecasting: flow moves price *as it trades*, so by the time a 1-minute CVD is printed, most of its price effect is in the lead already. The residual forecasting content is the same order-flow persistence that underlies OFI (§3.1) at a 1-minute-ish horizon.

Folklore warning (Grade D). "CVD divergence" (price makes a new high, CVD doesn't $\Rightarrow$ reversal) is taught in exchange academy articles (e.g. <https://phemex.com/academy/what-is-cumulative-delta-cvd-indicator>, <https://www.gate.com/crypto-wiki/article/trading-and-profiting-with-cvd-cumulative-volume-delta-in-cryptocurrency-markets-20260114>) with **no published effect sizes, no out-of-sample tests, and no cost accounting**. A more candid practitioner piece states the opposite of the folklore: "CVD doesn't [predict]. CVD is a measure of participation, not prediction... the winners are not always the most aggressive; they are the participants who absorb aggression" (<https://affiliateprofitsmedia.com/cumulative-volume-delta-cvd/>). Absorption is real market microstructure, but it means the *sign* of the CVD $\rightarrow$ price mapping flips depending on whether aggressive flow is informed or is being faded by a large passive counterparty — unobservable from CVD alone.

Verdict: TFI over 30–60 s is a legitimate *component* (it is essentially the trade-event part of OFI) and should be tested jointly with OBI, with which it will be highly collinear. Standalone CVD levels/divergences: **unsupported**. Horizon: ≤ 1 –2 min. Size: expect $\leq$ the OBI decile spread. Fees: same verdict as OBI.

3.3 (c) Order-flow toxicity / VPIN

Definition (Easley, López de Prado & O'Hara 2012, *Review of Financial Studies* 25:1457–1493 — Grade A). Partition the trade tape into buckets of equal *volume* V (original calibration: $V = 1/50$ of average daily volume, averaged over $n = 50$ buckets). For bucket τ with buyer-initiated volume V^B_τ and seller-initiated V^S_τ :

$$\text{VPIN} = \frac{1}{n} \sum_{\tau=1}^n \frac{|V^B_\tau - V^S_\tau|}{V}.$$

Crypto adaptation: buckets must be rescaled to the horizon — for a last-5-minutes model, V should be a small fraction of *hourly* volume and $n \approx 20$ –50, giving a VPIN that updates every few seconds; the daily calibration is meaningless intraday. Trade signing is exact on crypto venues (maker/taker flags), removing the bulk-volume-classification error of the equity original (noted at

<https://dev.to/followsm/detecting-orderbook-toxicity-polymarket-divergence-building-a-microstructure-circuit-breaker-ef9> — Grade D source for an implementation point).

The academic dispute is the finding. Andersen & Bondarenko (2014, *Journal of Financial Markets* 17, "VPIN and the Flash Crash" — Grade A critique) found VPIN peaked *after* the Flash Crash and had **no incremental predictive power for future volatility once trading intensity was controlled for**; the apparent predictiveness is "mechanical" via volume. Easley et al. conceded timing nuances while maintaining VPIN measures adverse-selection conditions for market makers (both sides summarised at <https://electronictradinghub.com/vpin-and-real-time-order-toxicity-what-your-execution-stack-cannot-see-before-the-fill/>).

Verdict for this problem: VPIN is a **volatility/regime variable, not a direction variable** — by construction it is an *absolute* imbalance, so it cannot sign a trade. Its only defensible use here is as an alternative σ -regime gate (§3.8), where it competes with the much simpler "trailing 5-minute RV percentile", which captures the same trading-intensity channel Andersen & Bondarenko identify. **Do not pay complexity budget for VPIN unless it beats RV-percentile as a vol gate out-of-sample.** No crypto-specific, cost-adjusted VPIN forecasting study at 1–5-minute horizons was found.

3.4 (d) Cross-exchange lead-lag

What leads, in the academic literature. Hasbrouck information-share / Gonzalo–Granger component-share studies of BTC disagree by era and sampling: 1-minute and 5-minute studies find spot leading CME futures (Corbet et al. 2018; Baur & Dimpfl 2019), daily studies find futures leading (Kapar & Olmo 2019), and later work (Fassas et al. 2020; Akyildirim et al. 2020; Alexander et al. 2020) finds futures prominent — surveyed at <https://www.sciencedirect.com/science/article/pii/S0165176523001775>. None of these minute-scale information shares answer the operational question, which is about *milliseconds of arrival*.

The operational measurement (Grade C, the best public one found). Gravia, 22–25 August (year as published on page; 70.6 hours), five perpetual BBO feeds recorded on one AWS Ireland (eu-west-1b) server (<https://gravia.trade/learn/who-moves-bitcoin-first>):

Venue	Feed	Engine location	Median delivery age to Ireland
Kraken	PF_XBTUSD book	London	7.6 ms
Deribit	BTC-PERPETUAL BBO	Equinix LD4, Slough	14.8 ms
Bybit	BTCUSDT orderbook.1	not disclosed	85.6 ms
Binance	BTCUSDT bookTicker	Tokyo	110.1 ms
OKX	BTC-USDT-SWAP bbo-tbt	Tokyo	118.9 ms

- On **exchange timestamps**: Binance leads Kraken and Deribit by **74 ms**, Bybit by 43 ms, OKX by 2 ms (Binance≈OKX, both Tokyo).
- On **arrival time in Ireland**: **Kraken leads Binance by 30 ms**, Deribit by 18 ms over Binance; full arrival order Kraken → Deribit → Binance → OKX → Bybit, all ten pairs passing stability checks (same sign across random 3-hour blocks and across 20/50/100 ms move definitions).
- **First-mover share** of 11,965 confirmed ≥ 1 bp moves: **Kraken 38.0%**, Binance 17.7%, Deribit 16.9%, Bybit 16.5%, OKX 10.9%.
- Author's own caveat, which this report endorses: "A 20 to 40 ms head start can easily be eaten by how often the prediction market updates, your queue position, order-entry latency, fees, and how long your code takes to act." One server, one region, one 3-day window, perps not spot, mids only.

Translation to KXBTC15M (this is the part most commentary gets wrong):

1. The venues that *lead* (Binance/OKX perps) are **not in BRTI**; the venues that *settle* (Coinbase, Kraken, Gemini, Bitstamp, etc.) are slower sources of the same information. The lead-lag that matters is therefore **Binance-perp mid → BRTI-constituent mids → BRTI print**, a two-hop chain. In stressed seconds the BRTI proxy built from constituents will lag a Binance-triggered fair value — a real, recurring, *small* staleness that Kalshi market makers also face and price.
2. At 60–300 s horizons, a 30–110 ms head start is worth exactly nothing for *probability estimation* (your $\sigma\sqrt{\tau}$ is \$15–50; a 100 ms head start captures $\sim\sigma\sqrt{0.1} \approx \$1$ of expected move). It is worth something only for *execution* in the final seconds (\$1.5) and for choosing which feed's timestamp to trust when feeds disagree.
3. **Coinbase's role is special**: it is both a BRTI constituent and the venue most retail charts show. The measured \$7.80 median gap between a Coinbase/Kraken spot read and the official settlement (\$1.4) is the number to beat by using a *consolidated constituent mid*.

Verdict: Use Kraken + Coinbase (constituents, fast, free) as the primary BRTI proxy; use Binance perp BBO as a *confirmation/leading indicator* feature — e.g. `binance_lead_bps = (Binance perp mid - constituent mid)` in bps, which is a genuine 0.1–2-second-ahead signal that decays long before 1 minute. Do not build a "Binance leads by X seconds" feature: at second-plus scales the venues are contemporaneous for practical purposes, and published minute-scale lead-lag (e.g. CEX leading DEX by 5 minutes in a Zenodo study, https://zenodo.org/records/17085264/files/price_discovery_paper.pdf — Grade C/D) concerns DEX staleness, irrelevant to BRTI.

3.5 (e) Perpetual futures: funding, open interest, liquidations

Funding rate — mechanics first. Funding is a periodic (typically 8-hourly on Binance/Bybit; Deribit/OKX variants differ) payment anchoring perp to spot. Within any single 15-minute window the *next* funding rate is essentially a known constant, printed in advance, changing only at settlement boundaries. It therefore **cannot contain 1–5-minute directional information except at extremes**, where it is a crowding gauge: sustained extreme positive funding = crowded longs = liquidation fuel *if* price drops. Practitioner dashboards state the same: "sustained extremes mean one side is crowded, and crowded trades unwind violently" — a risk statement, not a direction forecast (<https://www.bitpredict.io/predictions/bitcoin/next-hour>). No Grade A/C study was found showing funding *changes* predict 1–5-minute BTC direction after costs. **Verdict: regime/crowding filter at most; never a last-5-minutes trigger.**

Open interest changes. ΔOI over 1–5 minutes mixes new positions, closes, and liquidations; the standard 2×2 reading (price↑+OI↑ = new longs; price↑+OI↓ = short covering; etc.) is *descriptive of the move already happening*. An honest data-source caveat from a live OI tracker: "Open interest cannot reveal traders' entry prices, leverage or liquidation levels. Funding rates describe the cost of holding a perpetual, not individual positions. [The] liquidation heatmap models possible levels; it does not observe actual positions" (<https://bitcoinmargin.com/open-interest>). **Verdict: ΔOI is a useful confirmation feature inside a vol/cascade gate (rising OI + rising RV = fragile), not a signed predictor.**

Liquidations — the one derivatives flow that is genuinely fast and signed, but contemporaneous. Exchange liquidation streams (Binance `forceOrder`, Bybit `allLiquidation` @500 ms, OKX `liquidation-orders`) report forced closes as they happen. A burst of long liquidations *is* market selling: it pushes price down *during* the burst and is a component of the jump/cascade dynamics in §2.5 and Nicola's mean-reversion explanation ("excessive use of leverage leading to liquidation cascades"). Two tradable nuances: (i) liquidation bursts are **momentum at 1–10-second scale and exhaustion/reversal candidates at 1–5-minute scale** — the cascade ends when the leveraged inventory is cleared, which is one mechanism behind the negative 1-minute autocorrelation; (ii) as a *feature in the binary model*, trailing-60 s net liquidation imbalance belongs in the same regression as TFI/OFI and will be collinear with them. **Verdict: include as a cascade indicator in the vol gate and as a candidate signed feature at 30–60 s; expect it to be episodic (most windows have ~zero liquidations, so it is a sparse, high-variance feature).** Sources for the data itself: §5.

3.6 (f) Deribit options: DVOL, ATM IV, 25-delta skew

What they are. DVOL is Deribit's 30-day forward implied-volatility index for BTC, built VIX-style from the strip of out-of-the-money options bracketing 30 days, interpolated between the two nearest expiries; the raw value is computed **every second** and smoothed with an **EMA over the last 240 points** (Deribit's own methodology whitepaper, via <https://insights.deribit.com/industry/dvol-deribit-implied-volatility-index/> — Grade B; the URL was returned without a full link in one search session, so verify on Deribit Insights before citing externally). A current explainer confirms the construction and the conversion rule $DVOL \div \sqrt{365} \approx$ daily 1σ move, $DVOL \div 20$ as a rule of thumb (<https://coinlab.info/en/articles/bitcoin-dvol-implied-volatility/>).

Horizon mismatch — the decisive point. DVOL prices the *next 30 days*. Its EMA-240s smoothing means even its *changes* are 4-minute-scale. A same-day news source states the obvious correctly: "DVOL gives no indication of whether Bitcoin will rise or fall" (<https://www.tokenpost.com/news/investing/36192>). The 25-delta risk reversal ($RR25 = IV(25\Delta \text{ call}) - IV(25\Delta \text{ put})$) measures the *price of crash insurance over weeks*: e.g. −3.19 vol points for November expiry on 19 Oct 2026 (<https://www.tokenpost.com/news/investing/35733>) — information about positioning and fear on a multi-week horizon, with no published evidence of 1–5-minute predictive content, and a mechanical reason to expect none: options market makers re-hedge on spot moves, so at minute scale causality runs spot → IV, not IV → spot.

The one legitimate use: DVOL (or the front-expiry ATM IV, which is more responsive than DVOL) as a **slow prior for σ** — specifically, a sanity ceiling/floor for the day's RV regime and a variance-risk-premium adjustment: $\sigma_{\text{prior}, 1\text{min}} = X_t \cdot \frac{DVOL}{\sqrt{365 \cdot 24 \cdot 60}} \cdot \kappa$, with $\kappa \approx RV/IV \approx 0.55\text{--}0.70$ estimated from history (cf. the 20.2–25.5% RV vs 38.28% IV readings above). Even then, trailing RV (§2.2) will dominate it in any honest horse race at 5-minute horizon. **Verdict: too slow for direction; marginal for σ prior; skew: do not use at this horizon.**

3.7 (g) Stablecoin flows / Coinbase premium

Coinbase premium — two different objects are conflated under this name.

1. *The CryptoQuant-style Coinbase Premium Index*: an hourly/daily series of (Coinbase BTC-USD – Binance BTC-USDT)/Binance × 100, interpreted as US institutional demand. Its published use cases are daily/weekly moving-average crossovers (<https://coinmarketcap.com/community/articles/67003894b20f2815b3a379f9/>) and multi-day sentiment calls (e.g. premium –0.079, a one-month low, discussed as a *weeks-scale* signal at <https://bitriss.com/bitcoin-coinbase-premium-drops-to-monthly-low-after-clarity-act-vote-251976>). **At that sampling, it has no conceivable 1–5-minute content.**

2. *The instantaneous cross-venue spread* (Coinbase mid – Binance mid, in bps, sampled per second): this is a real microstructure variable — but it is exactly the lead–lag/staleness variable of \$3.4 wearing another name. If Coinbase (a BRTI constituent) trades rich to Binance for several seconds, either Coinbase-specific buying is pulling BRTI up, or Binance is lagging a move. Test it as `cb_premium_bps_30s_mean` — as a *spread* feature, not as the published index.

Stablecoin flows. Exchange stablecoin reserves/inflows are measured on-chain or from exchange wallet labels, arrive with block-confirmation and labelling delays of minutes-to-hours, and their own proponents describe reaction times of "days... [or] several weeks" (XWIN Research via <https://bitriss.com/stablecoin-liquidity-displays-clear-uptrend-when-will-bitcoin-price-follow-149782>). **Verdict: no role at 1–5 minutes. None.** The same applies to ETF flows (daily), exchange BTC reserves (hourly+), and funding-settlement calendars, except that *scheduled* events (CPI/FOMC prints, funding timestamps, options expiry at 08:00 UTC on Deribit, CME opens) belong in the model as **volatility-regime dummies**, because they reliably change σ , not direction.

3.8 (h) Realised-volatility regime as a confidence filter

This is the feature the evidence most supports — precisely because it makes no directional claim.

Mechanism. The probit model already divides by σ ; a *regime filter* is a second, discrete layer: only trade (or trade larger) when the model's preconditions hold. Concretely, define the trailing RV percentile:

`$$Regime_t = \text{percentile of } RV^{\{(30)\}}_t \text{ within the trailing 30 days at the same time-of-day}, \$`

and gate on it. Why it helps even if direction is unforecastable:

- **Calibration is regime-dependent.** In high-RV regimes the Gaussian probit understates tails most (\$2.5), Kalshi spreads widen, and the averaging-window dynamics (\$1.5) are dominated by single prints; in low-RV regimes, $|z|$ explodes for trivial dollar leads and the model calls 99%+ on windows that are actually decided by BRTI-proxy error (\$7.80 median for a spot proxy — larger than a low-vol \$ lead).
- **The edge-to-fee ratio is regime-dependent.** The fee in cents is fixed by price; the *probability* edge from any mispricing scales with how far fair value is from the market price, which is more often large when vol is mis-estimated by the market — i.e. at regime *transitions*.
- **Volatility itself is the forecastable part** (\$2.3: squared-return autocorrelations 0.26–0.39 at 1–5 min). A model that knows σ well and direction not at all can still make money as a *maker* by quoting around a well-calibrated fair value, and can avoid the worst losses as a *taker* by standing down when its z is unreliable.

Verdict: implement as (i) time-of-day-deseasonalised RV percentile, (ii) a hard stand-down when trailing-1-minute RV exceeds, say, the 98th percentile or when a liquidation cascade is in progress (\$3.5), and (iii) position-size scaling $\propto 1/\text{regime}$. Report all backtest results split by regime tercile — a feature that only "works" in one tercile and is marketed as unconditional is the most common form of self-deception in this literature.

3.9 Feature comparison table (consolidated)

Feature	Precise definition (testable)	Best evidence (grade)	Effect size	Horizon	Directional?	Survives Kalshi taker fee + spread alone?
Settlement running average	\$2.4: $E[S]$	F_t, X^*_req	Contract itself (B) + derivation	Changes SD by up to 3.4× in final minute	0–60 s	Yes (mechanical)
BRTI-proxy quality	Consolidated constituent mid vs single-spot	DegenHedge / PMP (C)	Proxy error \$7.80 → ~\$0.13 median	All	Reduces noise	Yes — prerequisite for everything else
OBI■ / OFI	\$3.1 formulas	CKS (A, equities); Imbalance Labs (C, crypto)	BTC rank-IC 0.113 @1m; extreme deciles ±1.1 bps next-min	1–2 min	Yes	No (authors' own verdict); use as fair-value tilt / maker input

Feature	Precise definition (testable)	Best evidence (grade)	Effect size	Horizon	Directional?	Survives Kalshi taker fee + spread alone?
TFI / Δ CVD 30–60 s	\$3.2 formulas	Price-formation lit. (A, contemporaneous)	Contemporaneous R^2 large; <i>forecast residual</i> ■ OBI	≤ 1 –2 min	Yes	No standalone; collinear with OFI
VPIN	\$3.3 formula	Easley et al. (A) vs Andersen–Bondarenko (A)	Disputed even for vol forecasting	—	No (absolute value)	N/A — vol gate candidate only
Cross-venue lead (Binance perp – constituent mid)	Spread in bps, per-second	Gravia (C)	Leads by 2–74 ms (event time); arrival reverses by venue	0.1–2 s	Yes, fleeting	No at minute scale; execution/proxy use only
Funding rate	Next printed rate; extremes only	Practitioner consensus (D)	No measured 1–5-min effect	8 h	No	No
Δ OI 1–5 min	% change, price-conditioned	Descriptive (C/D)	Unquantified for forecasting	5–60 min	Ambiguous sign	No
Liquidation imbalance 60 s	(short liqs – long liqs)/total, \$	Exchange streams (B data); cascade mechanism (A)	Episodic; large during cascades, ~0 otherwise	10 s–5 min	Yes during cascades (then reversal-prone)	Unproven; test as gate + signed feature
DVOL / front ATM IV	\$3.6; DVOL EMA-240 s	Deribit (B)	IV–RV gap ~13–18 vol pts (Oct 2026)	30 days (DVOL)	No	N/A — σ prior only
RR25 skew	IV(25 Δ C) – IV(25 Δ P)	Market data (B)	–3.2 vol pts (Nov, Oct 2026)	Weeks	No at this horizon	No
Coinbase premium (published index)	Hourly Coinbase–Binance %	CryptoQuant commentary (D)	Discussed at days–weeks scale	Hours–days	Weak, slow	No
Coinbase–Binance spread (instant)	bps, 30-s mean	= \$3.4 (C)	ms-scale	Seconds	Yes, fleeting	No standalone
Stablecoin / ETF / on-chain flows	Reserve changes	Proponents concede days–weeks lags (D)	None at minutes	Days–weeks	No	No
RV regime percentile	\$3.8 formula	Vol clustering (A)	$p(\text{sq. rets})$ 0.26–0.39 @1–5m	5–60 min	No — confidence/size	Yes, as a <i>filter</i> on other trades
RSI / MACD / SMA / VWAP composites	Standard TA on 1-min candles	GitHub READMEs only (D)	No audited effect sizes	—	Claimed	No evidence; see §6.3

4. Prediction-Market-Specific Research

4.1 Efficiency and calibration of Kalshi (all markets)

Bürgi, Deng & Whelan, "The Economics of the Kalshi Prediction Market", CESifo Working Paper No. 12122 (September 2025; Kalshi Research) — **Grade A**, transaction-level data on **over 300,000 contracts**

(https://www.ifo.de/sites/default/files/docbase/docs/cesifo1_wp12122.pdf; non-technical summary:

<https://Cepr.org/voxeu/columns/economics-kalshi-prediction-market>):

- "Kalshi's contract prices are **informative** and **improve in accuracy as markets approach closing**" — the calibration curve is close to the 45° line in aggregate.
- **Favourite–longshot bias is strong and pervasive**: across politics, entertainment, economic releases, and across volume/transaction-size subsamples. Cheap contracts win *far* less than priced; expensive contracts win slightly *more*.
- Returns arithmetic: a 5¢ contract winning 2% **loses 60%** of money invested; the other side (95¢ winning 98%) earns a small positive rate. "Even though Kalshi's market involves zero-sum profits (prior to fees), the **average pre-fee return on a contract is still minus 20%**" (because the *average contract bought* is a cheap one — a composition effect every backtest must replicate or explicitly avoid). **Post-fee average return: –22%**.
- Mechanism (their model, supported by taker-side data): **Makers** post offers seeking positive expected return after fees but are slightly over-optimistic; **Takers** accept them. The model predicts, and the data confirm, *distinct and stronger* FLB patterns for takers — loss rates for takers rise as price falls, because fee incidence and maker profits fall hardest on cheap contracts. (A secondary Grade D source quantifies this as "takers lose 32% on average, makers 10%" — <https://beatpoly.com/learn/why-traders-lose.html> — that specific split could not be verified in the working paper text retrieved; **cite the –20%/–22% and the qualitative maker/taker asymmetry, not the 32/10 split, unless verified in the paper's tables.**)

What this does and does not imply for KXBTC15M last-5-minutes: the aggregate FLB is *not* evidence that 95¢ KXBTC15M contracts in their last 5 minutes are underpriced — §4.2 measures that exact cell and finds them priced within ~0.5 points. FLB on Kalshi is concentrated in (a) cheap contracts, (b) taker flow, (c) less repetitive/less arbitrated categories. A 96-times-a-day, BRTI-arbitrated binary is close to the *most* efficient end of Kalshi's distribution. **Importing the exchange-wide FLB as a prior for this series is an ecological fallacy; the series-specific measurements govern.**

4.2 Calibration of KXBTC15M specifically in the last 5 minutes (the core empirical facts)

DegenHedge logged the ask for its model's side at a median **170 s** and separately **310 s** before close, graded on Kalshi's official result (Grade C; methods and CSVs published):

At ~3 minutes left (15,183 windows, 24 Mar–8 Oct 2026)

(<https://degenhedge.com/research/can-you-lose-a-99-cent-kalshi-contract/>):

Ask	Windows	Lost	Realised loss rate	Price-implied loss rate
90–94¢	2,123	185	8.71% (95%: 7.6–10.0)	7.81%
95¢	560	36	6.43% (4.7–8.8)	5.00%
96¢	841	27	3.21% (2.2–4.6)	4.00%
97¢	939	22	2.34% (1.6–3.5)	3.00%
98¢	1,693	38	2.24% (1.6–3.1)	2.00%
99¢	2,884	28	0.97% (0.7–1.4)	1.00%

At ~5 minutes left the same rows read: 90–94¢ lost 7.08%, 95¢ 4.62%, 96¢ 3.67%, 97¢ 2.47%, 98¢ 2.48%, 99¢ 0.56% (different windows/counts; see source). The settlement study's summary across all price groups at ~3 min (15,414 asks):

"in most groups the side won within a point or two of what it cost, usually a little under, which is the spread a buyer pays...

no price group in our data beat its own price after fees"

(<https://degenhedge.com/research/are-kalshi-bitcoin-15-minute-markets-rigged/>).

Independent corroboration (Grade C): PredictionMarketPicks graded 2,764 windows (29 Aug–27 Sep 2026) by 10¢ bands — "contracts near 46¢ won 45.8% of the time, contracts near 92¢ won 91.4%" — and concludes "a fair price is a fair price, and then you pay the fee" (<https://predictionmarketpicks.com/kalshi/15-minute-markets>). Their separate study: exactly 50.0% of 3,320 windows settled Up (<https://predictionmarketpicks.com/articles/is-kalshi-bitcoin-rigged/>).

Losing favourites, anatomised (same DegenHedge study): the 88 sides priced 97–99¢ at ~3 min that lost had a median lead of **\$73.18** (0.106%) with a median 169 s left, then reverted to settle a median **\$18.07 on the wrong side**; 53/88 settled within \$25 of the strike; in **15 of the 88, a Coinbase/Kraken spot read after the close still showed the bought side ahead** — the 60-second average decided otherwise. Winning sides at the same prices settled a median \$85.77 from the strike. **Lesson for model-building: your model's job in the 90¢+ region is not finding direction — it is estimating the probability of a \$70–90 lead being spent in ~3 minutes, i.e. exactly the σ -and-tails problem of \$2.2/\$2.5.**

4.3 How accurate is the market price vs a spot-lead model?

The two Grade C operations that actually ran spot-lead models against this market both disclose the same result:

- **DegenHedge's own feed:** "won **91.18%** at a **91.72¢** average entry from March 24 to October 8, 2026, **below breakeven before fees**" over 11,803 graded signals; at 99¢ entries specifically, won 99.06% — "0.06¢ per contract ahead before fees and 0.94¢ behind after a 1¢ fee" (<https://degenhedge.com/research/can-you-lose-a-99-cent-kalshi-contract/>). Their framing deserves quoting because it is the honest equilibrium statement: "*Fair is not the same as favourable. A fairly priced market with a fee is a market where buyers lose a little on average.*"
- **BotForKalshi (practitioner, real-money, self-reported):** buying the leading side over 6,298 windows: **gross +0.67¢/contract, average taker fee 1.55¢ $\Rightarrow$ net negative at every entry time tested, including the last minute**; across eleven price bands the leading side "settled slightly better than its quoted price implied — roughly one point of genuine underpricing" that the fee more than consumes (<https://www.botforkalshi.com/blog/kalshi-bitcoin-bot>).

Interpretation. A spot-lead probit model and the market price contain *almost the same information* in the last 5 minutes (both are, to first order, $\Phi(\text{lead}/\sigma\sqrt{\tau})$ computed by many participants). The measured residual — the market underpricing the leader by ~0.5–1 point gross — is the *sum* of (a) genuine model improvements available (averaging correction, better σ , better BRTI proxy) and (b) the spread you must cross. It is ~1¢ wide and the taker fee alone is 0.07–1.75¢. **Therefore the research question is not "find a better direction model" but "find the sub-populations where the residual is widest"**: plausibly (i) the first seconds of the settlement minute, before index-reading programs finish repricing (§1.5); (ii) regime transitions where the market's implicit σ is stale (§3.8); (iii) near-tie windows where BRTI-proxy differences dominate (§1.4); (iv) as a *maker*, quoting inside the spread instead of crossing it (§4.4). Each is testable on Kalshi's free 1-minute candle/bid-ask history (§5.6) plus your own BRTI-proxy recording.

4.4 Kalshi fees, spread, and the edge needed to break even

Fee formulas (Grade B schedule as reproduced consistently across Grade C sources; verify per-series multiplier before trading, as schedules are revised — a July 7, 2026 schedule is referenced at <https://launchpoly.com/blog/polymarket-kalshi-arbitrage-guide>):

$\text{Taker fee} = \text{round_up}(0.07 \times C \times P \times (1-P))$, $\text{Maker fee (flagged series only)} = \text{round_up}(0.0175 \times C \times P \times (1-P))$, \$

rounded up **per order** to the next cent; no fee to hold to settlement; cancelling is free. BotForKalshi reports KXBTC15M specifically carries **no maker fee** (resting orders free) — treat as Grade C and verify in the live fee schedule for the series before relying on it, since maker-fee flagging is series-specific and changeable.

Table 4.2 — Breakeven win probability for a taker buy held to settlement ($EV = q \cdot \$1 - P - \text{fee} \geq 0 \Rightarrow q^* = P + \text{fee}$; fees shown per contract in size, i.e. large C so rounding vanishes, and for $C = 1$):

Price P	Taker fee/contract (size)	Breakeven q^* (size)	Taker fee, $C=1$ (rounded up)	Breakeven q^* , $C=1$
50¢	1.750¢	51.75%	2¢	52.00%
65¢	1.593¢	66.59%	2¢	67.00%
75¢	1.313¢	76.31%	2¢	77.00%
90¢	0.630¢	90.63%	1¢	91.00%
92¢	0.515¢	92.52%	1¢	93.00%
95¢	0.333¢	95.33%	1¢	96.00%
97¢	0.204¢	97.20%	1¢	98.00%
98¢	0.137¢	98.14%	1¢	99.00%
99¢	0.069¢	99.07%	1¢	100.00%
99.9¢	0.007¢	99.91%	1¢	100.90% (impossible)

Sources for the formula and the rounding behaviour: <https://www.oddsshopper.com/articles/prediction-markets/kalshi-not-for-dummies-five-advanced-lessons-on-prediction-markets-you-may-not-h> (table of fees by price, and the explicit warning that 1-contract orders on penny contracts pay a full cent) and <https://degenhedge.com/research/can-you-lose-a-99-cent-kalshi-contract/> (99¢ arithmetic). DegenHedge's dedicated fee study (11,803 signals) confirms "a one-contract order at 99¢ pays a 1¢ fee, so its breakeven is 100%" and found **no entry-price bucket beat its own price by more than chance**; the 90–92¢ bucket fell short by 2.80 standard errors (summarised at <https://degenhedge.com/research/>).

Add the spread. Table 4.2 assumes you buy at P = the ask and that P is the price in the fee formula. If your model compares fair value to the *mid*, the half-spread (typically 1¢ mid-market, wider in the final minute for near-tie windows, and 0.1¢-granular above 90¢) is an additional cost of the same order as the fee. **Round-trip (enter and exit as taker) roughly doubles the fee line**: PredictionMarketPicks puts in-and-out near 50¢ at "about 3.5¢ a contract before the spread" (<https://predictionmarketpicks.com/kalshi/15-minute-markets>).

The break-even edge statement for the commissioned price region: at 90–98¢, a taker needs a *true* probability **0.6–1.4 points above the ask** (in size) just to cover the fee, before any spread/impact — against a market whose measured calibration error in that region is ■0.9 points (§4.2) and usually in the *wrong* direction for the buyer. This is the quantitative reason every measured taker strategy in §4.3 is net negative, and why the ranked list in §7 is dominated by calibration, proxy-quality, execution-style and filtering improvements rather than new directional signals.

4.5 Polymarket BTC Up-or-Down (5m / 15m): settlement, the manipulation study, and the August 2026 TWAP change

Current settlement (global platform). Polymarket's own market pages state the 15-minute market "resolves based on whether Bitcoin's price at the end of the 15-minute window is greater than or equal to its price at the start of that window...

The resolution source is the Chainlink BTC/USD data stream"

(<https://polymarket.com/event/btc-updown-15m-1760308200>). The page template describes point-to-point comparison; the *implementation* of that comparison changed in August 2026 (below) — always read the live market's Rules section, because pre- and post-change markets are different contracts in effect.

The Stanford/SMU study and what it actually found (reported Grade C, underlying work Grade A-style working paper). CoinDesk's report (7 Aug 2026,

<https://www.coindesk.com/business/2026/08/07/how-a-five-second-trick-let-traders-drain-millions-from-polymarket/>):

- Researchers from **Stanford University and Singapore Management University** examined ~two months of Polymarket **5-minute** BTC contracts, which then settled on a **single instant price snapshot** via Chainlink.
- They found **unusually large orders on Binance in the final seconds before settlement, followed by rapid price reversals** — the signature of temporary price pressure, not information.
- **821 accounts made \$8.2 million** in windows classified as likely manipulated; **excluding market makers, 93% of the losses in those windows fell on retail traders**; "a bet the market treated as near-certain was overturned one time in three" in those windows. Other reporting of the same study puts the estimated net transfer at ~\$1.28M over its specific sample window (<https://www.cointribune.com/en/new-research-raises-concerns-over-polymarket-bitcoin-contracts/>) — the \$8.2M (gross profits of flagged accounts in flagged windows) and \$1.28M (estimated net wealth transfer) are different quantities; do not conflate them.
- The paper **did not prove intent** nor that the Binance orders came from the same persons holding Polymarket positions. Its structural point: "An asset-price contract settles on a financial price, and that price can be moved by trading the underlying market itself."
- **Polymarket's response:** replaced single-price snapshots with **time-weighted average prices via Chainlink Data Streams — 30-second average for 5-minute markets, 60-second average for 15-minute and 4-hour markets —** plus \$1M of liquidity rewards for August. CoinDesk notes this "mirrors safeguards used by rival platform Kalshi". A Kalshi spokesperson added that its 60-second average over regulated exchanges makes brief distortion "significantly harder and more expensive" than on snapshot platforms.

Why this matters for a Kalshi model (and its limits):

1. **Kalshi was never snapshot-settled** (in the averaging era, \$1) — the specific exploit (push spot for 1–5 seconds across a strike) requires moving a *60-second consolidated-book average across seven venues*, costing roughly 60× the duration and a multiple of the venues. No public study documents successful settlement manipulation of KXBTC15M; DegenHedge's integrity checks (\$1.2/\$1.4) are the relevant negative evidence.
2. **Latency/bot dynamics transfer directly.** Polymarket 5-minute trader-population data (Bitquery, Grade C): in one measured week, **34,482 wallets** traded markets that opened and settled inside the week; **31% finished ahead**; the population was level to within 0.1% *before* fees and **down \$1.62M after \$1.69M of fees**; the fifth of wallets playing >100 markets carried ~3/5 of all losses; median result above 500 markets was ≈ -\$100 (<https://bitquery.io/investigations/polymarket-five-minute-casino>). The population structure — many retail takers, few systematic makers/bots, fees as the main net flow — is the same ecology as Kalshi's (\$4.1's maker/taker asymmetry).
3. **Polymarket fees are a moving target — do not hard-code.** The original 15-minute-crypto taker fee (Jan 2026) was $\text{fee} = C \times 0.25 \times (p(1-p))^2$ (100 shares @50¢ ⇒ \$1.56) with proceeds funding maker rebates (<https://blog.upay.com/polymarket-added-taker-fees-to-its-15-minute-crypto-markets/>). The later (30 Mar 2026) general rollout uses $\text{fee} = C \times \text{feeRate} \times p(1-p)$ with **crypto feeRate = 0.072** (100 @50¢ ⇒ \$1.80), makers paying zero and receiving 20% of taker fees in crypto (https://swaphunt.dev/articles/polymarket-fee-structure-math?utm_source=medium&utm_medium=article, <https://marketmath.io/blog/polymarket-fees-explained>). Which formula applies is per-market (feesEnabled and the market's fee schedule). For cross-venue arbitrage comparisons, note both venues' crypto taker curves peak at 50¢ at ≈1.75–1.80¢/contract — **there is no fee arbitrage between Kalshi and Polymarket crypto; apparent price gaps**

must first clear ~2× that plus two spreads.

4. **Chainlink access for replication** — see §5.4. Note the asymmetry: on Polymarket you can (with credentials) subscribe to the *exact* settlement stream; on Kalshi you cannot subscribe to BRTI as a retail API product (§5.1) and must build a proxy. This is a genuine structural difference in model-building cost between the two venues.

4.6 Kalshi vs Polymarket for the same 15 minutes — comparison table

Dimension	Kalshi KXBTC15M	Polymarket BTC Up-or-Down 15m (global)	Polymarket US (QCX) 15m (launched 22 Sep 2026)
Settlement source	CF Benchmarks BRTI (order-book index, 7 venues)	Chainlink BTC/USD Data Stream	Chainlink USD feed (per launch reporting)
Settlement statistic	Average of 60 × 1-s prints before close vs same before open	Since Aug 2026: 60-s TWAP (Chainlink) at close vs open reference; before: single snapshot	Reported as averaging 60 readings from the last minute before open and close (https://defirate.com/prediction-markets/crypto/)
Tie	Yes (>=)	Up (>=) per market pages	Up
Taker fee (crypto)	round_up(0.07 · C · P · (1 - P))	Per-market: 0.25 · (p(1 - p)) ² legacy 15m, or 0.072 · p(1 - p) general crypto — check market	0.0695 × P × (1 - P) per contract (https://defirate.com/prediction-markets/crypto/)
Maker	No fee on this series per Grade C report; 0.0175 formula on flagged series	Zero + rebate share (20% crypto)	Not established in sources read
Identity / access	CFTC DCM, KYC, USD	Global, USDC, on-chain CLOB (Polygon)	Regulated US app, card/bank funding
Public data	Free unauthenticated REST history incl. 1-min bid/ask candles + trade tape (§5.6); WS requires RSA auth	Gamma/CLOB APIs public; on-chain settlement record	App-centric

(Polymarket US launch facts: https://cryptobriefing.com/polymarket-15-minute-bitcoin-markets-us/.)

5. Live Data Sources

Selection criterion from the commission: free or public APIs, real-time, 1-second to 1-minute granularity, that provide something **Coinbase 1-minute candles do not**. Coinbase 1-minute candles give you one venue's trade-price OHLCV, finalised after the minute ends. Everything below improves on at least one of: *latency* (per-trade/per-quote, not per-minute), *price type* (mid/book, not last trade), *venue coverage* (BRTI constituents or leaders), or *variable type* (depth, flow, liquidations, the settlement stream itself).

5.1 CF Benchmarks BRTI — the settlement feed itself

- **Product:** CME CF Bitcoin Real Time Index (BRTI), administered/calculated by CF Benchmarks Ltd; disseminated ~every second, \$0.01 precision; methodology in §1.3 (https://docs.cfbenchmarks.com/CME%20CF%20Real%20Time%20Indices%20Methodology.pdf).
- **Access reality (important, and a "could not fully verify" item in §8):** BRTI is a **licensed benchmark**. Real-time redistribution is sold via CF Benchmarks/CME market-data licences and resellers (e.g. Kaiko, and Kraken — CF Benchmarks' owner — surfaces some index data). **No free, public, unauthenticated real-time BRTI API was found**, and DegenHedge — who had every incentive — states "We cannot recompute the BRTI" from public data (https://degenhedge.com/research/are-kalshi-bitcoin-15-minute-markets-rigged/). Kalshi's own app displays a live reference price, and Kalshi publishes the *final* strike and settlement value per window via its API (§5.6) — sufficient for grading, not for intra-minute running-average tracking.
- **Practical consequence:** the running-average model of §2.4 must run on a **self-built BRTI proxy**: consolidated mid across the constituent venues' public WebSockets (§5.2), size-capped and spread-filtered to mimic the methodology. Budget the proxy error from §1.4 (\$0.13 median is achievable per PMP with a real index-like feed; \$7.80 is what a lazy spot read costs you) into your z as an additional variance term: $\sigma^2_{\text{total}} = \sigma^2_{\text{model}} + \sigma^2_{\text{proxy}}$, with σ_{proxy} estimated by regressing your proxy's 60-s average against Kalshi's published `expiration_value` over history.
- **Cost:** proxy = \$0 (public WebSockets) + your server. Licensed BRTI = commercial, quote-only; treat any specific price found on reseller marketing pages as unverified.

5.2 Spot exchange WebSockets (BRTI constituents + leaders)

Source	Endpoint / channels (as documented)	Granularity & latency	Cost / auth	What it gives that Coinbase 1-min candles don't
Coinbase Advanced Trade WS (constituent)	ticker (every match), ticker_batch (5 s), level2 (full book updates), market_trades, candles (updates every second, 5-min buckets) — docs: https://docs.coinbase.com/coinbase-business/advanced-trade-apis/websocket/websocket-channels	Per-trade / per-quote, real-time; no auth for market channels	Free, no auth for market data	Per-trade signing (taker side), L2 book for OBI/OFI, mid-price instead of last-trade; the candles channel is <i>not</i> 1-minute (5-min buckets) — build your own 1-s/1-min bars from trades/quotes
Coinbase Exchange WS (constituent)	ticker, level2 + level2_batch (50 ms batches), matches, full (order lifecycle) — https://help.coinbase.com/en/developer-platform/websocket-feeds/exchange	Per-event; level2_batch 50 ms	Free (public channels)	The full channel is the richest public Coinbase feed (received/open/change/done) for OFI-style features
Kraken WS v1/v2 (constituent)	ticker, trade, book (depths 10/25/100/500/1000), spread, ohlc — https://docs.kraken.com/api/docs/websocket-v1/spread/ ; overview https://blog.kraken.com/product/api/unlocked-3-the-market-data-feeds-systematic-traders-use	Per-event; book depth to 1000 levels	Ticker/L2/trades/OHLCV free , no auth; L3 (individual orders) requires auth	Deepest free <i>constituent</i> book; Kraken was also the fastest-arriving perp signal to Europe in §3.4 and is a BRTI constituent for spot — the single best free proxy input
Gemini / Bitstamp / Crypto.com / Bullish / LMAX (constituents)	Each offers public WS market data (Gemini marketdata, Bitstamp live orders/trades, Crypto.com market channels); LMAX is institutional/FIX-oriented	Per-event	Free (retail venues)	Completing the constituent set reduces proxy error most in the seconds when Coinbase/Kraken diverge from the consolidated book; diminishing returns after ~3 venues for a <i>mid</i> proxy
Binance Spot + USD-M Futures WS (NOT a constituent — the leader)	Spot: aggTrade, bookTicker, depth@100ms, kline_1s/1m; Futures: aggTrade, bookTicker, depth@100ms, markPrice@1s, forceOrder (liquidations), funding via REST/WS — new base URLs <code>wss://ftstream.binance.com/public</code>	market	private: https://developers.binance.com/docs/derivatives/usds-margined-futures/websocket-market-streams/Important-WebSocket-Change-Notice	bookTicker per quote change; depth 100 ms; mark price 1 s (or 3 s default)
Bybit WS v5 (not a constituent)	orderbook.1 (10 ms), orderbook.50 (20 ms), orderbook.200 (100 ms), orderbook.1000 (200 ms); publicTrade; tickers (spot 50 ms snapshots, linear 100 ms); allLiquidation (500 ms) — catalogue via https://apis.io/asynccapis/bybit/bybit-asyncapi/ (mirror of Bybit's official AsyncAPI)	10–500 ms by channel	Free, no auth for public	Fastest published book cadence (10 ms top-of-book); second liquidation source for cross-checking Binance
OKX WS (not a constituent)	bbo-tbt (best bid/offer tick-by-tick), books / books5, books-l2-tbt, trades, index-tickers, mark-price, funding-rate, open-interest, liquidation-orders — channel list via https://github-wiki-see.page/m/tiguisigalpa/okx-php/wiki/WebSocket-API (third-party wiki of the official API; verify channel names in OKX docs before coding)	Tick-by-tick for bbo-tbt/books-l2-tbt (requires VIP tier for some tbt channels historically — verify)	Free for most public channels	The only major venue bundling spot index + perp funding + OI + liquidations in one WS ; OKX was co-leader with Binance on event time in §3.4

Engineering notes that affect model quality more than venue choice: (i) timestamp on *arrival* and keep exchange timestamps separately — §3.4's entire result is that they disagree by ~100 ms across venues; (our research window did not test your server's region — re-measure delivery ages from your own host before trusting any lead). (ii) Build 1-second mid bars per venue, then a consolidated mid = median across venues of (mid – venue's trailing 5-min median basis) — basis-correcting before median-filtering is what makes a proxy behave like an index rather than an average of stale quotes. (iii) Record everything at 1 s for the settlement minute even if the model runs at 1 min: the §2.4 running average cannot be

reconstructed from 1-minute candles afterwards.

5.3 Perpetuals / derivatives aggregators

Source	Product / endpoints	Granularity	Cost	Feature vs Coinbase candles
Exchange-native (Binance/Bybit/OKX above)	funding rate, open interest, force Order/allLiquidation/liquidation-orders	Funding: per settlement (8 h) + predicted rate updates ~1 min; OI: ~1 min REST / push on OKX; liquidations: per event / 500 ms	Free	First-party, lowest latency. Always prefer native streams over aggregators for anything used inside 5 minutes.
CoinGlass API v4	Base https://open-api-v4.coinglass.com , header CG-API-KEY; endpoints incl. /api/futures/fundingRate/exchange-list, .../fundingRate/ohlcv-history, .../openInterest/ohlcv-history, .../liquidation/history, taker buy/sell history; intervals m1 m5 m15 m30 h1... (https://dev.to/great-time-flies/how-to-build-a-crypto-derivatives-monitor-with-coinglass-api-and-python-mp4)	1-minute minimum interval on history endpoints	API is a separate paid product, no free tier; from \$29/mo (Hobbyist) scaling to several hundred \$/mo; website dashboards/heatmaps free, enhanced heatmap via Prime ~\$28/mo (secondary reviews: https://simplemarkets.io/blog/post/coinglass-liquidation , https://stage.crypto-investing.page/reviews/coinglass.html — verify on coinglass.com/pricing)	Cross-exchange <i>aggregation</i> in one call (28 exchanges, 2,500+ instruments per CoinMarketCap's acquisition release, https://bitrss.com/coinmarketcap-acquires-derivatives-data-platform-coinglass-which-will-keep-operating-independently-255171) — useful for backfill, regime dashboards and liquidation <i>history</i> ; its 1-minute REST polling is too slow and too coarse to be the live liquidation feed for a last-5-minutes model
CoinGlass website	Liquidation dashboard (1 h/4 h/12 h/24 h), heatmaps, funding/OI pages	Real-time display; 24 h heatmap free	Free tier	Human regime monitoring only

5.4 Chainlink Data Streams (Polymarket's settlement feed; a BRTI-independent reference)

- **What:** Chainlink's low-latency pull-based oracle network. Reports are produced off-chain by a decentralised oracle network, cryptographically signed, and delivered via **REST + WebSocket**; verification happens on-chain only when needed (<https://docs.chain.link/data-streams/architecture>).
- **Access:** Requires **credentials** (API key + user secret, HMAC-signed requests). TypeScript SDK config: `endpoint: https://api.dataengine.chain.link, wsEndpoint: wss://ws.dataengine.chain.link`, with High-Availability mode maintaining ≥ 2 connections (<https://docs.chain.link/data-streams/reference/data-streams-api/ts-sdk>). A public **Discovery endpoint** (GET /api/v1/discovery) lists streams without auth but returns metadata, not report data (<https://docs.chain.link/data-streams/reference/data-streams-api/discovery-endpoint>). The BTC/USD stream ID used in Chainlink's own examples is `0x000359843a543ee2fe414dc14c7e7920ef10f4372990b79d6361cdc0dd1ba782` (same SDK doc).
- **Granularity/latency:** sub-second report generation (Streams is explicitly the "low-latency" product vs the minutes-scale on-chain Data Feeds); your realised latency is WebSocket delivery of signed reports, typically sub-second to ~1 s.
- **Cost:** credentialed/commercial for production use; there is no unauthenticated live BTC/USD report stream. Budget as a paid data line if Polymarket modelling is pursued; for Kalshi-only work it is *optional* — valuable as an independent, manipulation-resistant reference price to cross-check your BRTI proxy (a persistent Chainlink-vs-proxy divergence is a proxy-bug alarm), not as a settlement input.
- **Free Chainlink alternative and its limit:** the on-chain BTC/USD Data Feed (AggregatorV3) is free to read via any RPC, but updates on deviation/heartbeat (minutes-scale) — **too stale for this horizon**; do not confuse the two products in code or in backtests.

5.5 Deribit (options + its own index/perp)

- **Public API (free, no auth for market data):** JSON-RPC over REST/WS at <https://www.deribit.com/api/v2/> / `wss://www.deribit.com/ws/api/v2/`. Relevant methods/channels: `public/get_volatility_index_data` (DVOL history/live), `public/get_book_summary_by_currency / ..._by_instrument` (open interest, 24 h volume, mark IV per instrument — https://docs.deribit.com/api-reference/market-data/public-get_book_summary_by_instrument), `public/get_historical_volatility`, `ticker/book/trades` channels for BTC-PERPETUAL and options, and the Deribit index

price (`index_price` channel) — Deribit's own BTC index is itself a multi-exchange spot index and a usable *secondary* BRTI cross-check.

- **Granularity:** DVOL raw 1 s (EMA-240 smoothed as published, §3.6); option book summaries on demand/push.
- **Feature vs candles:** the entire options surface — front-expiry ATM IV (for the σ prior of §3.6), RR25, term structure, options OI. **At this horizon: regime/prior data only (§3.6 verdict).**

5.6 Kalshi's own API (the market's price, and your grading data)

- **Free, unauthenticated historical REST** (base <https://external-api.kalshi.com/trade-api/v2/> per the field note; the elections subdomain api.elections.kalshi.com serves the same API in other sources — verify which host your client uses):
 - GET `/historical/markets?series_ticker=KXBTC15M&limit=1000` — every market back to Dec 2025, with `open/close/result/settlement` fields.
 - GET `/historical/trades?ticker=<full market ticker>` — full trade tape, both-side prices, count, taker side, **microsecond timestamps**.
 - GET `/historical/markets/{ticker}/candlesticks?...&period_interval=1` — 1-minute candles with last-trade OHLC **plus yes-bid and yes-ask OHLC, volume, and open interest at minute close** (<https://pondletter.com/kalshi-free-history-api/>). The bid/ask candles are the key free resource: they let you backtest *executable* prices and spreads, not just last trades.
- **Live WebSocket** (`wss://api.elections.kalshi.com/trade-api/ws/v2`): channels include market ticker, `orderbook_delta` (snapshot + deltas), and public trades — but the **connection itself requires RSA API-key authentication** (headers `KALSHI-ACCESS-KEY` / `-SIGNATURE` / `-TIMESTAMP`, RSA-PSS/SHA-256 signing), even for public-data channels (<https://www.quantvps.com/blog/kalshi-order-book-api-endpoints-explained>, <https://prediction.com/blog/kalshi-api-getting-started-guide>).
- **Feature vs candles:** Kalshi's book *is* the market's probability — needed for (a) fair-value-vs-market gap measurement, (b) maker strategies, (c) detecting the final-minute repricing speed of \$1.5. A model that never observes the Kalshi book cannot know whether its "edge" is already in the price — and \$4.2–4.3 say it usually is.

5.7 Everything else considered for live data — and rejected for this horizon

Candidate	Why rejected / limited
Mempool data (mempool.space etc.)	Mempool congestion/fees measure <i>on-chain settlement demand</i> , with block times ~10 min and no mechanical link to spot order flow in the next 1–5 min. No study found linking mempool state to minute-scale BTC direction. No role.
Whale alerts / large on-chain transfers	A transfer to an exchange "sometimes signal[s] potential selling... Neither action proves that a trade occurred" and alerts "sometimes trigger speculative reactions despite transfers representing internal wallet reorganisations" (https://www.analyticsinsight.net/bitcoin/what-are-bitcoin-whales-how-large-btc-holders-impact-the-market). Confirmation + labelling delay alone disqualifies them at 5 minutes; the signal content at <i>any</i> horizon is unproven in Grade A/C sources. No role.
Stablecoin exchange flows	§3.7: proponents' own horizon is days–weeks. No role.
CryptoQuant / Glassnode on-chain dashboards	Hourly/daily sampling, paid APIs, on-chain variables. Useful for day-scale regime context in a research notebook; no 1–5-minute feed exists in these products that this research could verify.
News / social sentiment (X, Fear & Greed)	Fear & Greed is daily. Social velocity is unstandardised, bot-contaminated, and no cost-adjusted minute-scale study was found. Scheduled <i>macro releases</i> matter — as a σ -regime dummy (§3.7), sourced from an economic calendar, not from sentiment.
TradingView / charting webhooks	A display layer over the same exchange data with added latency; several "Kalshi/Polymarket predictor" scripts found are Grade D (§6.3). Use raw feeds.

5.8 Recommended live stack (synthesis of §5.1–5.7)

Layer	Build	Latency budget	Cost
BRTI proxy	Coinbase Exchange full/level2 + Kraken book/trade + Gemini or Bitstamp book → 1-s basis-corrected consolidated mid; running 60-s sum always maintained	<200 ms end-to-end achievable on a regional VPS	\$0 data
Leader/confirmation	Binance USD-M bookTicker + aggTrade + forceOrder; Bybit orderbook.50 + allLiquidation	10–110 ms by venue/region (\$3.4) — measure from your host	\$0
Derivatives regime	OKX WS funding-rate/open-interest or native REST polling 1/min; Deribit get_volatility_index_data 1/min	1 min is fine for regime variables	\$0
Market itself	Kalshi WS orderbook_delta (RSA auth) + REST fallback polling	Exchange-bound	\$0 data (account required)
Grading/backfill	Kalshi historical REST (trades + bid/ask candles) nightly; CoinGlass API only if cross-exchange derivatives <i>history</i> is needed	—	\$0 / from \$29/mo

6. Anything Else Credible

6.1 Published strategies and GitHub projects — an honest census

What was found, graded without charity:

1. **DegenHedge's public feed + studies (Grade C, the most credible "strategy" artefact found)** — not a strategy sale: a BRTI-lead model whose operator publishes every graded signal and discloses it wins *below* breakeven before fees (91.18% at 91.72¢ avg, \$4.3). Its value to you is its *datasets and methods* (calibration tables, fee study, settlement study), all linked in \$8.
2. **BotForKalshi's post-mortem (Grade C–, self-reported, no dataset)** — the "buy the leader" taker bot, stopped 2 Aug 2026 after measuring +0.67¢ gross vs 1.55¢ fee (\$4.3). Its prescriptions: rest limit orders (no maker fee on this series per the author), expect adverse selection on fills ("a resting order fills disproportionately when the price is about to move against you" — the classic maker winner's-curse, consistent with Bürgi et al.'s maker model in §4.1).
3. **suislanchez/polymarket-kalshi-weather-bot (GitHub, Grade D)** — Polymarket BTC 5-min strategy = weighted composite of RSI, 1/5/15-min momentum, VWAP deviation, SMA crossover on Coinbase/Kraken/Binance 1-minute candles; trades at "edge > 2%"; fractional Kelly 15%; Brier-score calibration tracking (<https://github.com/suislanchez/polymarket-kalshi-weather-bot>). The *engineering* (calibration logging, fractional Kelly, paper mode) is sound practice; the *signal* is an unaudited TA composite of exactly the type \$2.7/\$6.3 warn against. No verified P&L.
4. **CloddsBot (open-source, MIT; via a bots listing, Grade D)** — multi-venue agent framework with round-based discovery of Polymarket 5m/15m/hourly crypto windows and four automated strategies (<https://polbots.com/>). Framework code is real; no audited strategy performance was found.
5. **Benj4minCup / Polymarket-Trading-Bot-python (Grade D)** — explicitly "strategy descriptions and research materials... not a complete production-ready bot", with a commercial upsell (http://dev.to/benjamin_cup/polymarket-trading-bot-building-market-maker-bot-with-python-a-beginners-guide-ba0). Representative of the genre: treat all such repos as marketing until a graded, timestamped trade log is published.
6. **TradingView "probability" indicators (Grade D implementations of Grade A math)** — e.g. the Brownian/CDF fair-value script in §2.1 and a "strike-to-beat predictor" combining daily Fibonacci levels with WMA/TEMA (<https://www.tradingview.com/script/3aJdDZR4-Kalshi-Polymarket-15m-Predictor-Strike-To-Beat/>). The CDF ones are §2.1 with a chart; the Fibonacci/structure ones have no published validation whatsoever.

Census conclusion: *No public, audited, profitable last-5-minutes KXBTC15M strategy was found.* The two operators who publish graded records both report gross edges of roughly +0.5–1 point that fees consume. This is itself the most important "anything else" finding: **the prior for any new claimed strategy in this market should be "gross ≈ +0–1¢, net ≈ –1¢ as taker" until its own graded log proves otherwise.**

6.2 Optimal sizing for near-expiry binaries at 90¢+: Kelly, done properly

Derivation for a \$1 binary. Buy at ask P (dollars), true win probability q, fee per contract f (dollars, \$4.4). Net odds on stake: you risk (P + f) to win (1 – P – f). Expected log-growth-optimal fraction of bankroll to *stake* (stake = amount paid for contracts, the standard prediction-market formulation) is:

$$f^* = \frac{q - (P+f)}{1 - (P+f)} \quad \text{(equivalent to } f^* = (p - q_{\text{lose}})/b \text{ with } b = \frac{1-P-f}{P+f} \text{)}$$

Without fees this reduces to the clean form cited by practitioner references: $f^* = (q - P)/(1 - P)$ for a YES at price P with estimate q (<https://infiniti.w3ai.io/intelligence/position-sizing-binary-event-contracts-kelly>), and in odds form $f^* = (p - (1-p))/b$ (<https://www.convexly.app/blog/kelly-criterion-explained>; general theory: https://en.wikipedia.org/wiki/Kelly_criterion).

Worked examples at the commissioned prices (fee in size, from Table 4.2):

Ask P	Model q	Effective cost P+f	Full Kelly f*	Quarter-Kelly	Comment
90¢	92%	90.63¢	14.6%	3.7%	A 2-point edge at 90¢ is a <i>huge</i> Kelly bet — because the downside is capped at the stake and the model claims the loss rate is 8% not 10%
92¢	94%	92.52¢	19.8%	4.9%	
95¢	96%	95.33¢	14.3%	3.6%	
95¢	95.5%	95.33¢	3.6%	0.9%	Halving the claimed edge (1pt → 0.5pt... here 1pt→0.17pt over cost) collapses Kelly — estimation error, not the formula, is the risk
98¢	99%	98.14¢	46.2%	11.6%	Absurd on its face: it says "bet half your bankroll that a 98¢ contract is really 99%". If q is actually 98.0%, f* should be ≈ –7% (don't trade). A ±1-point error in q swings the answer from "half the bankroll" to "nothing"
99¢, C=1	99.5%	100.00¢	0%	0%	The rounding tax (\$4.4) makes single-contract 99¢ buys un-Kelly-able at any q < 100%

Rules that follow from the math + the calibration evidence (§4.2):

- 1. Use quarter-Kelly or less, with hard caps.** Practitioner consensus in the sources read: quarter-to-half Kelly, per-market cap 2–5% of bankroll, per-theme and daily-loss caps on top (<https://infiniti.w3ai.io/intelligence/position-sizing-binary-event-contracts-kelly>). The GitHub bot in §6.1 uses 15% of Kelly. Given that your q comes from a model whose calibration error in the 90¢+ region is measured at ±0.5–1 point (§4.2), and Table 6.1 shows ±1 point is the difference between f*=46% and f*=0, **cap discipline matters more than the Kelly fraction.**
- 2. Shrink q toward the market price before sizing.** If the market is calibrated to ±1 point (§4.2), a model claiming q – P = 3 points in the last 5 minutes is more likely mis-calibrated than right. A defensible shrinkage: $q_{\text{used}} = P + w(q_{\text{model}} - P)$ with $w \approx 0.3\text{--}0.5$ estimated from your own graded log's calibration slope. Kelly on unshrunk model probabilities is the fastest known way to convert a small real edge into ruin.
- 3. Consecutive windows are not independent bets.** §1.2: strikes chain, and §2.7: outcomes mildly anti-correlate; more importantly, *your model errors* correlate across windows within a vol regime (a stale σ loses repeatedly for an hour). Kelly's independence assumption fails in clusters — another argument for caps and regime gates (§3.8) over raw f*.
- 4. Maker vs taker changes the Kelly input, not the formula:** for a resting order, P is your bid, $f \approx 0$ on this series (per §4.4 caveat), but q must be the *fill-conditioned* win probability, which is lower than the unconditional q by adverse selection (§6.1.2). Sizing maker bets on unconditional q overstates f*.

6.3 What does NOT work at this horizon — the negative evidence, collected

1. **Taker "buy the leader / buy the favourite"** — measured twice, net negative both times (§4.3). The direction is right ~1 point more often than priced; the fee is ~1.55¢.
2. **Streak/momentum-of-results betting ("it's due" / "it's on a run")** — the anti-streak pattern is real but (a) halved over 2026, (b) half-priced within the window's first minute, (c) smaller than the fee at the ask (§2.7). DegenHedge's verdict: "Treat 'it's due' and 'it's on a run' as equally weak reasons."
3. **Single-exchange spot as settlement truth** — wrong side in 4.37% of windows overall and 39.2% of close windows (§1.4). Any strategy whose backtest grades itself on Coinbase closes instead of Kalshi's published `expiration_value` has mis-graded ~1 window in 23 — larger than any edge claimed in this report.
4. **TA composites on 1-minute candles (RSI/MACD/SMA/VWAP-deviation)** — no audited effect size exists in any source found; the only implementations are Grade D repos/scripts (§6.1). At 1-minute scale these indicators are smooth functions of the same returns whose mid-price autocorrelation is ≈0-to-negative (§2.7); a composite of lags of a near-martingale is a near-martingale forecast.
5. **Last-second snapshot sniping (the Polymarket pre-Aug-2026 playbook)** — the mechanism was removed by the venue itself (§4.5); on Kalshi it never existed in the averaging era (§1). Articles/bots still selling it are selling a dead edge.
6. **News/sentiment/whale/mempool triggers** — §5.7: latency, labelling, and unproven content, each disqualifying alone.
7. **"99¢ = free money" harvesting as a taker** — §4.2/§4.4: loss rate ≈ price-implied; at C=1 the fee is the entire win. The strategy's Sharpe is dominated by a single number — the fatness of BTC's 3-minute tail (§2.5) — which the Gaussian model underestimates and the market, measurably, does not.

6.4 Two structural ideas that are *not* in the usual lists (and are testable)

1. **Cross-venue settlement-basis trading (Kalshi vs Polymarket, same window).** The two venues settle the "same" 15 minutes on *different indices* (BRTI vs Chainlink) with *different averaging implementations*. In near-tie windows the two settlement values will differ by their basis (historically, BRTI-vs-Chainlink-type bases are a few dollars, occasionally >\$10 in fast minutes — *measure this, do not assume it*). When Kalshi and Polymarket prices for the economically-identical side differ by more than (basis uncertainty + both fees + both spreads), a hedged pair (long the cheap side on one venue, long the opposite side on the other) isolates the basis. This is a *settlement-mechanics* trade, not a direction trade, and it is the natural extension of §1 to §4.6. Risks: the Polymarket fee-formula ambiguity (§4.5.3), funding/withdrawal friction, and non-simultaneous fills. No public graded record of this trade was found — flagging it as an *opportunity to measure*, not a proven strategy.
2. **Quoting the final minute as a maker around the running average.** In the final 60 s, fair value is dominated by the §2.4 running-average arithmetic, which is *computable exactly* from a BRTI proxy, while much of the visible Kalshi flow is still reacting to spot. A maker posting both sides at $\Phi(z_{\text{running}}) \pm \text{margin}$, cancelling on proxy staleness, monetises §1.5 without paying the taker fee — and §4.4/§6.1 say maker execution is the only style not measured as structurally negative. Risks: adverse selection is at its maximum exactly here (your counterparties include the index-reading programs of §1.5), queue position, and the auth/WS engineering of §5.6. This is a market-making research programme, not a signal.

7. Ranked Lists (the deliverable)

7.1 The 5 most promising, testable improvements for a last-5-minutes model

Each is stated with the exact formula/feature, the dataset needed, and the test that would falsify it. Ordered by (expected value × confidence) ÷ implementation cost.

#1 — Replace the point-settlement probit with the settlement-averaging probit (§2.4)

Change: from $P = \Phi\left(\frac{X_t - K}{\sigma\sqrt{\tau}}\right)$ to $P = \Phi\left(\frac{X_t - K}{\sigma\sqrt{\tau - 40}}\right)$ ($\tau \geq 60$), $\text{Iquad } P = \Phi\left(\frac{A_{\text{obs}}/60 + (\tau/60)X_t - K}{\sigma\tau^{3/2}/(60\sqrt{3})}\right)$ ($\tau < 60$), A_{obs} = running sum of 1-second BRTI-proxy prints in the settlement minute. **Why first:** it is not a statistical bet — it is the contract's own arithmetic (§1.1, Grade B). The naive formula is *wrong by construction*, by 18% of SD at $\tau=120$ and 42% at $\tau=60$, and catastrophically inside the final minute. **Data needed:** 1-second BRTI proxy recording (§5.8) for the final-2-minutes extension; the $\tau \geq 60$ branch is testable *today* on the existing 1-minute dataset (use each minute's close as X_t , $\tau \in \{300, 240, 180, 120, 60\}$). **Falsification test:** calibration by z-bin at each τ snapshot, old vs new formula, out-of-sample (train Aug–Sep, test Oct, or walk-forward). If the new formula does not reduce Brier score / log-loss at $\tau=120$ and $\tau=60$, your σ or proxy — not the formula — is the binding problem; proceed to #2/#3.

#2 — Rebuild σ as a deseasonalised, EWMA/HAR intraday estimator, and calibrate the tails (§2.2–2.3, §2.5)

Change: replace any single-lookback close-to-close σ with $\hat{\sigma}_t^2 = \text{EWMA}(\lambda(r_{1\text{min}}^2), \lambda \in \{0.90, 0.94, 0.97\}, \text{divided by the time-of-day seasonal factor } s^2(\text{hour, weekday}))$, horse-raced against (a) Garman–Klass on the same bars, (b) the intraday HAR regression of §2.3, and (c) a DVOL-prior blend (§3.6). Then replace Φ with the Student-t form of §2.5 ($v \in \{3, 4, 5, \infty\}$) or an isotonic $z \rightarrow \text{win-rate}$ map. **Why second:** every measured failure of favourites (§4.2: \$73 median leads spent in 3 minutes) is a σ /tails failure, not a direction failure. Vol clustering at 1–5 min is Grade A (ρ of squared returns 0.26–0.39), so σ is genuinely forecastable at this horizon. **Falsification test:** for each candidate σ , the same calibration-by-z-bin test as #1, reported **split by RV-regime tercile** (§3.8). A σ estimator that wins on average but loses in the top tercile is not an improvement; it is a regime bet.

#3 — Build the BRTI proxy properly: consolidated constituent mid, basis-corrected (§1.3–1.4, §5.2)

Change: replace the single-venue spot input with $X_t^{\text{proxy}} = \text{median}_{v \in \{\text{Coinbase, Kraken, Gemini, Bitstamp, Crypto.com}\}}(\text{mid}_{v,t} - \text{basis}_{v,t})$, $\text{Iquad } \text{basis}_{v,t} = \text{trailing-5-min median of } (\text{mid}_{v,t} - \text{cross-venue median})$, size-capped/spread-filtered per the BRTI methodology's spirit, sampled at 1 s, and add the measured proxy variance to the probit denominator ($\sigma_{\text{total}}^2 = \sigma^2 + \sigma_{\text{proxy}}^2$) with σ_{proxy} fit against Kalshi's published `expiration_value` history). **Why third:** measured proxy error falls from \$7.80 (spot) toward ~\$0.13 (index-like feed) — in close windows (1 in 22 settle within \$5) proxy error is the outcome. It also fixes the strike-side reconstruction for any historical work. **Falsification test:** distribution of (your 60-s proxy average – Kalshi `expiration_value`) over $\geq 1,000$ windows; target: median |error| < \$1 and wrong-side rate in within-\$5 windows ■ 39.2%.

#4 — Add the two microstructure tilts that survive scrutiny — OFI/OBI and TFI — as a bounded lead adjustment (§3.1–3.2)

Change: compute, from Kraken + Coinbase books and Binance perp trades, at each decision time: OBI_5 (§3.1), $\text{Iquad } \text{OFI}_{\{60\text{s}\}} \in \{ \text{§3.1} \}$, $\text{Iquad } \text{TFI}_{\{60\text{s}\}} = \frac{V^{\text{taker, buy}} - V^{\text{taker, sell}}}{V^{\text{taker, buy}} + V^{\text{taker, sell}}}$ (§3.2), regress forward 1-minute *mid* returns (bps) on them over a training sample, and add the fitted expected move, **clipped to $\pm 0.5 \cdot \sigma\sqrt{\tau}$** , to the lead in the #1 formula. Include trailing-60 s liquidation imbalance (§3.5) as a fourth regressor and let the regression kill it if it is noise. **Why fourth:** the mechanism is Grade A and the crypto size is measured (rank-IC 0.113 @1m, extreme deciles ± 1.1 bps/min) — but the honest expected contribution is a few dollars of lead in extreme-flow minutes only, i.e. it moves P by 1–3 points when $|z|$ is small. That is exactly where a *maker* or a *selective taker* can use it; it does not rescue unconditional taking. **Falsification test:** incremental out-of-sample log-loss of the probit *with* vs *without* the tilt, on mid-price features computed with a strict 1-bar execution delay (§2.6's trap). If the improvement is < the fee in the traded sub-population, drop it — the Imbalance Labs authors predict you might have to.

#5 — Trade selection and execution, not signals: regime gate + maker-first execution + shrunk Kelly (§3.8, §4.4, §6.2)

Change (three coupled rules):

- **Gate:** trade only when (a) $|\text{model } P - \text{market ask}| > \text{fee}(P) + \text{half-spread} + 0.5 \text{ pt safety}$, (b) RV regime is not in its top ~2% or in an active liquidation cascade, (c) $\tau \geq 60$ s for taker entries (leave the final minute to the running-average specialists, §1.5).

- **Execute maker-first:** rest at fair value (or 0.1¢ inside) instead of crossing; measure fill-conditioned calibration separately.
- **Size:** quarter-Kelly on the *shrunk* probability $q_{\text{used}} = P + 0.4 \cdot (q_{\text{model}} - P)$, hard cap 2–5% of bankroll per window, daily loss stop.

Why fifth but still top-5: §§4.2–4.4 prove the unconditional taker version of *any* model in this market is net negative; selection + execution is where the measured ~1-point gross residual (§4.3) can survive. It is also the only item on this list that reduces losses even if #1–#4 all fail. **Falsification test:** paper-trade the full rule set against Kalshi's live book (recorded via §5.6 WS) for $\geq 2,000$ windows with timestamped, pre-registered fair values; report net P&L after the *actual* fee schedule and *actual* fills (for maker: only fills that would have occurred). If $\text{net} \leq 0$, the correct conclusion — supported by all Grade C evidence in this report — is that the last-5-minutes KXBTC15M taker/maker retail niche is efficient at the ~1¢ level, and the model's value is defensive (avoiding bad trades) rather than offensive.

7.2 The 5 most commonly claimed things the evidence says do NOT help

#	Common claim	What the evidence says	Key source
1	"Buy the side that's winning late — 90¢+ is free money / a lock"	The leading side is underpriced by only ~0.7–1 pt gross; the taker fee averages ~1.55¢ and the measured strategies are net negative at every entry time. 99¢ asks lose 0.97% = exactly their price; a 1-contract 99¢ buy needs a 100% win rate after fee rounding.	https://www.botforkalshi.com/blog/kalshi-bitcoin-bot ; https://degenhedge.com/research/can-you-lose-a-99-cent-kalshi-contract/
2	"Watch Coinbase / your chart — if it's above the line at the close, you won"	Settlement is a 60-print BRTI average. A spot read is on the wrong side in 4.37% of windows, 39.2% of close windows; 15 of 88 losing 97–99¢ favourites still "looked ahead" on spot after the close.	https://degenhedge.com/research/are-kalshi-bitcoin-15-minute-markets-rigged/
3	"After several Up windows, Down is due (or: ride the streak)"	Anti-streak lean is real (48.0% vs 51.6%) but halved during 2026, is ~half-priced within a minute of the next window, and at the ask costs more than it wins. Both "due" and "run" versions are smaller than the 2¢ fee on a 50¢ contract.	Same DegenHedge study, Check 4
4	"RSI/MACD/VWAP/Fibonacci composites on 1-minute candles predict the next 15 minutes"	No audited effect size exists anywhere in the public record found; the implementations are marketing repos/scripts. 1-minute mid-price direction is near-unforecastable beyond the flow effects of §3.1–3.2, and trade-price "signals" at this scale are largely bid–ask bounce (a 20.5 bp illusion that becomes 3.8 bps with a 1-bar delay and 0 by 5 bars in the pairs literature).	§6.1 census; https://www.mdpi.com/411108
5	"Funding / DVOL skew / Coinbase premium / whale alerts / stablecoin flows give the next 5 minutes' direction"	Each is either (a) unsigned by construction (DVOL, VPIN), (b) sampled at hours-to-weeks (funding settlements, premium index, stablecoin reserves — proponents' own horizon is days–weeks), or (c) unproven and slow (whale transfers may be custody moves). Their legitimate roles — σ prior (DVOL), crowding/regime gate (funding extremes) — are not direction signals.	§3.5–3.7; https://www.tokenpost.com/news/investing/36192 (DVOL "gives no indication of whether Bitcoin will rise or fall")

Runners-up for the does-not-help list: VPIN as a direction signal (§3.3 — it is absolute-valued); CVD "divergences" as taught in academy articles (§3.2 — no published test); last-second snapshot sniping (§4.5 — mechanism abolished Aug 2026); and daily/weekly chart structure (support/resistance, daily Fibonacci) applied to a 15-minute settlement (§6.1.6 — no validation found, and the settlement statistic in §1 makes single-price levels doubly irrelevant).

8. Could Not Verify / Open Questions

1. **Licensed real-time BRTI price and redistribution terms.** CF Benchmarks sells BRTI as a licensed benchmark; no public price list was found, and no free real-time BRTI API could be verified. Any vendor quote should be obtained directly. (§5.1)

2. **Whether KXBTC15M currently carries a maker fee.** The general schedule's maker formula ($0.0175 \times \dots$) applies only to *flagged* series; one Grade C practitioner source states this series has none. This is checkable in Kalshi's live fee schedule / API market object and **must** be checked before implementing §7.1 #5 — the report flags it rather than asserting it.
3. **The Stanford/SMU Polymarket paper's full text.** Findings here are taken from CoinDesk and CoinTribune reporting of the study (Grade C reporting of Grade A work), including the \$8.2M / 821-account / 93%-retail figures and the \$1.28M transfer estimate, which are *different quantities from different passages* — obtain the working paper before citing either number in a publication.
4. **DVOL methodology whitepaper URL.** Deribit Insights hosts the DVOL whitepaper (construction + EMA-240 smoothing are corroborated by multiple secondary sources), but one search session returned the item without a usable full URL, so this report cites the corroborating explainer and Deribit's API docs instead. Verify on Deribit Insights directly.
5. **Coinbase premium / Kimchi premium at minute scale.** No study at 1–5-minute horizons was found in either direction; §3.7's verdict is "no evidence + wrong sampling in all published uses", not a measured null.
6. **Crypto-specific VPIN forecasting at 1–5 minutes.** No study found; §3.3's verdict rests on the equity VPIN dispute (Grade A both sides) plus VPIN's unsigned construction.
7. **The commission's "launched Aug 13, 2026" date.** Contradicted as a *launch* date by public records (§0 note); unresolved whether it denotes a local dataset start, a relaunch, or a different series event. It does not affect any settlement/math conclusion.
8. **Gravia lead-lag year and generalisability.** The page reports "22–25 August" and 70.6 hours from one Ireland server; the *pattern* (Tokyo venues lead on event time, London venues win on European arrival) is mechanically plausible, but re-measurement from the reader's own host/region is required before engineering on it (§3.4).
9. **OKX tick-by-tick channel entitlement and Bybit engine location** — channel names/entitlements should be verified in the venues' current official docs; Bybit's matching-engine location is undisclosed in the sources read.
10. **Kalshi sub-minute price history.** Kalshi's free candles are 1-minute; the trade tape has microsecond timestamps and *can* reconstruct sub-minute last-trade paths, but historical *order-book* depth for KXBTC15M is not in the free API (third-party archives such as DepthFeed sell it — <https://kalshipricedata.com/historical-data> — unverified quality). Backtests of final-minute maker strategies (§6.4.2) therefore require self-recorded live book data going forward.

9. Sources

All URLs copied verbatim from the sources as retrieved on 21 Oct 2026. Grade in brackets.

Primary / official (Grade B)

- Kalshi CRYPTO15M contract terms — https://assets.kalshi.com/contract_terms/CRYPTO15M.pdf
- Kalshi CRYPTO contract terms — https://assets.kalshi.com/contract_terms/CRYPTO.pdf
- CF Benchmarks, CME CF Cryptocurrency Real Time Indices Methodology Guide v17.0 (21 Sep 2026) — <https://docs.cfbenchmarks.com/CME%20CF%20Real%20Time%20Indices%20Methodology.pdf>
- CF Benchmarks, CME CF Constituent Exchanges — <https://docs.cfbenchmarks.com/CME%20CF%20Constituent%20Exchanges.pdf>
- CME/Crypto Facilities, BRTI Methodology Guide v1.0 (2016) — <https://www.cmegroup.com/trading/files/bitcoin-real-time-index-methodology.pdf>
- CME Group, Analysis of CME CF Bitcoin Reference Rate (BRR methodology) — <http://www.cmegroup.com/education/articles-and-reports/analysis-of-cme-cf-bitcoin-reference-rate>
- Coinbase Advanced Trade WebSocket Channels — <https://docs.cdp.coinbase.com/coinbase-business/advanced-trade-apis/websocket/websocket-channels>
- Coinbase Exchange WebSocket feeds (help) — <https://help.coinbase.com/en/developer-platform/websocket-feeds/exchange>

- Binance Developers, WebSocket Change Notice (futures streams incl. bookTicker/depth/aggTrade/forceOrder) — <https://developers.binance.com/docs/derivatives/usds-margined-futures/websocket-market-streams/Important-WebSocket-Change-Notice>
- Kraken Developers, Spread channel — <https://docs.kraken.com/api/docs/websocket-v1/spread/>
- Kraken Blog, API market-data feeds — <https://blog.kraken.com/product/api/unlocked-3-the-market-data-feeds-systematic-traders-use>
- Deribit Docs, public/get_book_summary_by_instrument — https://docs.deribit.com/api-reference/market-data/public-get_book_summary_by_instrument
- Chainlink Data Streams Architecture — <https://docs.chain.link/data-streams/architecture>
- Chainlink Data Streams TypeScript SDK (endpoints, stream IDs, HA mode) — <https://docs.chain.link/data-streams/reference/data-streams-api/ts-sdk>
- Chainlink Data Streams Discovery Endpoint — <https://docs.chain.link/data-streams/reference/data-streams-api/discovery-endpoint>
- CoinGlass Crypto Data API (product page) — <https://www.coinglass.com/CryptoApi>
- Polymarket market page, BTC Up-or-Down 15m (Chainlink resolution, >= rule) — <https://polymarket.com/event/btc-updown-15m-1760308200>

Peer-reviewed / working papers (Grade A)

- Bürgi, Deng & Whelan, The Economics of the Kalshi Prediction Market, CESifo WP 12122 — https://www.ifo.de/sites/default/files/docbase/docs/cesifo1_wp12122.pdf ; summary — <https://Cepr.org/voxeu/columns/economics-kalshi-prediction-market>
- Nicola, On the Intraday Behavior of Bitcoin, Ledger 6 (2021) (journal text via reprint) — <https://www.readkong.com/page/on-the-intraday-behavior-of-bitcoin-ledger-journal-4532086>
- Takaishi, Statistical properties and multifractality of Bitcoin, Physica A (2018) (record) — <https://ideas.repec.org/a/eee/phsmmap/v506y2018icp507-519.html>
- Shanaev et al., Turn-of-the-candle effect in Bitcoin returns, Heliyon — <https://www.sciencedirect.com/science/article/pii/S2405844023014433>
- Bergsli, Lind, Molnár & Polasik, Forecasting volatility of Bitcoin (2021) (record) — <https://ntnuopen.ntnu.no/ntnu-xmlui/handle/11250/3003405?show=full>
- What Explains Bitcoin Volatility? Evidence from an Extended HAR Framework, MDPI — <https://WWW.MDPI.COM/2227-7072/14/4/81>
- Bitcoin volatility forecasting / Garman–Klass proxy usage, JRFM — <https://doi.org/10.3390/jrfm14070337>
- Statistical Arbitrage in Cryptocurrency Markets (bid–ask bounce / execution-delay results), MDPI — <https://www.mdpi.com/411108>
- Periodicity in Cryptocurrency Volatility and Liquidity (microstructure noise), arXiv — <https://arxiv.org/html/2109.12142v2>
- Scaillet et al., High-Frequency Jump Analysis of the Bitcoin Market — https://www.academia.edu/108497250/High_Frequency_Jump_Analysis_of_the_Bitcoin_Market
- Price Appreciation and Roughness Duality / multifractality of 1-min Bitcoin, MDPI — <https://www.mdpi.com/2227-7390/9/17/2088/>
- Price discovery between Bitcoin spot markets and ETPs (literature survey), ScienceDirect — <https://www.sciencedirect.com/science/article/pii/S0165176523001775>

Practitioner measurement studies (Grade C)

- DegenHedge, Are Kalshi 15-Minute Bitcoin Markets Rigged? 25,746 Settlements Checked — <https://degenhedge.com/research/are-kalshi-bitcoin-15-minute-markets-rigged/>

- DegenHedge, Can You Lose a 99-Cent Kalshi Contract? — <https://degenhedge.com/research/can-you-lose-a-99-cent-kalshi-contract/>
- DegenHedge, Research index (fee/breakeven, strike, settlement, time-of-day studies) — <https://degenhedge.com/research/>
- OddsShopper, Kalshi 15-Minute Bitcoin Rigged? We Checked 200 Windows (2,200 total) — <https://www.oddsshopper.com/articles/prediction-markets/are-kalshi-15-minute-bitcoin-markets-rigged>
- OddsShopper, Kalshi NOT For Dummies (fee parabola table) — <https://www.oddsshopper.com/articles/prediction-markets/kalshi-not-for-dummies-five-advanced-lessons-on-prediction-markets-you-may-not-h>
- PredictionMarketPicks, How Kalshi 15-Minute Markets Settle — <https://predictionmarketpicks.com/articles/how-kalshi-15-minute-markets-settle>
- PredictionMarketPicks, Is Kalshi Bitcoin Rigged? 3,320 Windows — <https://predictionmarketpicks.com/articles/is-kalshi-bitcoin-rigged>
- PredictionMarketPicks, Kalshi 15-Minute Markets: Every Series, Live — <https://predictionmarketpicks.com/kalshi/15-minute-markets>
- PM Countdown, Crypto prediction markets: 15-minute and hourly — <https://pmcountdown.com/markets/crypto>
- Imbalance Labs, Does Orderbook Imbalance Predict Crypto Returns? A 24-Asset Rank-IC Study (vendor data — see §3.1 caveat) — <https://imbalancelabs.com/research/orderbook-imbalance-predictive-power>
- Gravia, Who Moves Bitcoin First? Binance Leads, but Kraken Reaches Your Bot Sooner — <https://gravia.trade/learn/who-moves-bitcoin-first>
- BotForKalshi, Kalshi Bitcoin Bot: 15-Minute BTC Setup & Paper Checks (self-reported) — <https://www.botforkalshi.com/blog/kalshi-bitcoin-bot>
- Bitquery, Polymarket Fees and the 5-Minute Markets (wallet population study) — <https://bitquery.io/investigations/polymarket-five-minute-casino>
- CoinDesk, Prediction market Polymarket overhauls rules after study finds fraud (Stanford/SMU study report) — <https://www.coindesk.com/business/2026/08/07/how-a-five-second-trick-let-traders-drain-millions-from-polymarket>
- CoinTribune, New Research Raises Concerns Over Polymarket Bitcoin Contracts — <https://www.cointribune.com/en/new-research-raises-concerns-over-polymarket-bitcoin-contracts/>
- Pond Letter, Kalshi gives away months of history (free API field note) — <https://pondletter.com/kalshi-free-history-api/>
- QuantVPS, Kalshi API Guide: Endpoints, Auth and Order Book — <https://www.quantvps.com/blog/kalshi-order-book-api-endpoints-explained>
- Electronic Trading Hub, VPIN and Real-Time Order Toxicity (incl. Andersen–Bondarenko critique) — <https://electronictradinghub.com/vpin-and-real-time-order-toxicity-what-your-execution-stack-cannot-see-before-the-fill/>
- Electronic Trading Hub, Order Book Imbalance Trading (Cont–Kukanov–Stoikov summary) — <https://electronicTradingHub.com/leveraging-limit-order-book-imbalance-for-profitable-trading-a-deep-dive-into-recent-research-and-practical-tools/>

Blog / marketing / folklore (Grade D — cited to attribute claims, not as evidence)

- Alphascope, 15-Minute Bitcoin Markets: Kalshi Rules & Polymarket — <https://www.alphascope.app/blog/15-minute-bitcoin-prediction-markets>
- BitcoinCalculator, CF Benchmarks BRTI Explained (outdated constituent/method description — see §1.3 correction) — <https://bitcoincalculator.tools/learn/cf-benchmarks-brti-explained>
- CryptoBriefing, Polymarket launches 15-minute Bitcoin price markets on US app — <https://cryptobriefing.com/polymarket-15-minute-bitcoin-markets-us/>
- DeFi Rate, Best Sites for Up/Down Crypto Prediction Markets (Polymarket US fee/averaging details) — <https://defirate.com/prediction-markets/crypto/>
- Upay Blog, Polymarket Added Taker Fees to Its 15-Minute Crypto Markets (legacy fee formula) — <https://blog.upay.com/polymarket-added-taker-fees-to-its-15-minute-crypto-markets/>

- SwapHunt, How Do Polymarket Fees Work? (general crypto fee rate) — https://swaphunt.dev/articles/polymarket-fee-structure-math?utm_source=medium&utm_medium=article
- MarketMath, Polymarket Fees Explained — <https://marketmath.io/blog/polymarket-fees-explained>
- LaunchPoly, Polymarket–Kalshi Arbitrage 2026 (Kalshi fee schedule date) — <https://launchpoly.com/blog/polymarket-kalshi-arbitrage-guide>
- Mario Karras, Kalshi — A Complete Reference for Building Trading Models (fee formulas) — <https://mariokarras.com/explainers/kalshi-trading-models/kalshi-guide-for-claude.html>
- CoinLab, Bitcoin DVOL: implied volatility index explained — <https://coinlab.info/en/articles/bitcoin-dvol-implied-volatility/>
- RegimeRisk, DVOL Explained — <https://regimerisk.com/blog/dvol-explained-what-the-deribit-volatility-index-says-about-bitcoin-risk>
- TokenPost, Bitcoin's 30-Day Implied Volatility Tops 14-Day Realized (DVOL 38.28, 20 Oct 2026) — <https://www.tokenpost.com/news/investing/36192>
- TokenPost, Bitcoin Wallets / 25-delta risk reversal -3.19 (Oct 2026) — <https://www.tokenpost.com/news/investing/35733>
- Phemex Academy, How to Use the CVD Indicator — <https://phemex.com/academy/what-is-cumulative-delta-cvd-indicator>
- Gate Wiki, Trading and Profiting with CVD — <https://www.gate.com/crypto-wiki/article/trading-and-profiting-with-cvd-cumulative-volume-delta-in-cryptocurrency-markets-20260114>
- AffiliateProfitsMedia, Most Traders Misunderstand CVD — <https://affiliateprofitsmedia.com/cumulative-volume-delta-cvd/>
- HarbourFrontQuant, Trading Strategies Based on Autocorrelation Properties of Bitcoin — <https://harbourfrontquant.substack.com/p/trading-strategies-based-on-autocorrelation-properties>
- BitPredict, Bitcoin Price Prediction Next Hour (funding/OI practitioner framing) — <https://www.bitpredict.io/predictions/bitcoin/next-hour>
- BitcoinMargin, Open Interest — Live Tracker (OI limitations) — <https://bitcoinmargin.com/open-interest>
- CoinMarketCap Community / CryptoQuant, Coinbase Premium Index Signals (hourly/daily usage) — <https://coinmarketcap.com/community/articles/67003894b20f2815b3a379f9/>
- AnalyticsInsight, What are Bitcoin Whales? (whale-transfer caveats) — <https://www.analyticsinsight.net/bitcoin/what-are-bitcoin-whales-how-large-btc-holders-impact-the-market>
- Infiniti Intelligence, Position Sizing for Binary Event Contracts: The Kelly Criterion in Practice — <https://infiniti.w3ai.io/intelligence/position-sizing-binary-event-contracts-kelly>
- Convexly, Kelly Criterion for Prediction Markets — <https://www.convexly.app/blog/kelly-criterion-explained>
- Wikipedia, Kelly criterion — https://en.wikipedia.org/wiki/Kelly_criterion
- GitHub, suislanchez/polymarket-kalshi-weather-bot — <https://github.com/suislanchez/polymarket-kalshi-weather-bot>
- POLBOTS, Best Polymarket Trading Bots (CloddsBot listing) — <https://polbots.com/>
- TradingView, Polymarket 15m BTC Probability script — <https://www.tradingview.com/script/bKp6f7yZ-Polymarket-15m-BTC-Probability-10x/>
- TradingView, Kalshi/Polymarket 15m Predictor + Strike To Beat — <https://www.tradingview.com/script/3aJdDZR4-Kalshi-Polymarket-15m-Predictor-Strike-To-Beat/>
- Quantopia, Integrated Brownian Motion ($\text{Var} = t^{3/3}$) — <https://www.quantopia.net/interview-questions-vii-integrated-brownian-motion/>
- MarkowitzOptimizer, Compound Rate Variance (1/3 averaging factor) — <https://www.markowitzoptimizer.pro/wiki/CompoundRateVariance>
- Bybit AsyncAPI catalogue (via APIs.io mirror) — <https://apis.io/asynccapis/bybit/bybit-asyncapi/>
- OKX WebSocket channel list (third-party wiki) — <https://github-wiki-see.page/m/tigusigalpa/okx-php/wiki/WebSocket-API>

- CoinGlass API tutorial (Dev.to, endpoint list) —
<https://dev.to/great-time-flies/how-to-build-a-crypto-derivatives-monitor-with-coinglass-api-and-python-mp4>
- SimpleMarkets, CoinGlass Liquidation Heatmap & Data Explained (pricing) —
<https://simplemarkets.io/blog/post/coinglass-liquidation>
- DepthFeed / KalshiPriceData, Kalshi Historical Data (third-party book archive) —
<https://kalshipricedata.com/historical-data>

End of report. Prepared 21 October 2026. Every factual claim above is tied to a source in §9; derivations in §2.4 and §6.2 are original to this report and can be checked from the stated assumptions. Where sources conflicted (BRTI constituents, Polymarket fee formulas, maker-fee applicability, the series launch date), the conflict is reported in the text rather than resolved silently.