

Price Discovery When the Benchmark Is Closed: Evidence from 24/7 Perpetual Futures

Xiaojian Li* Jian Yuan†

Abstract

Perpetual futures on stocks, commodities, and currencies trade around the clock, but their benchmark markets close. For about half of the time the benchmarks are closed, the venue sets the reference price that anchors funding, margins, and liquidations from its own order book. We ask how much this local price anticipates the external price that later replaces it, and what limits it. In all seven asset classes, a 1% local price change over a weekend maps into a 0.7%–0.9% change in the reference price once the external price returns. The shortfall reflects noise from a thin order book that the benchmark’s closure makes thinner: a holiday closure halves top-of-book depth and cuts order-book trading by 78%. A few large accounts trading aggressively bring in most of the information. Because news accumulates the longer the external price is unavailable while noise does not, the weight on the local price that best predicts the reference price rises from about 0.1 in brief maintenance breaks to 0.85 over weekends and holidays.

JEL classification: G14, G13, G18.

Keywords: price discovery, market closure, reference price, perpetual futures, 24/7 trading, liquidity.

1 Introduction

When the market that sets an asset’s price is closed, how much of the price can the markets that remain open discover? Research has asked this question for pre-opening, after-hours, and overnight trading, where the off-hours price is later replaced by the primary market’s. It now arises for derivatives that never close, and there the off-hours price is not replaced; it is used. Perpetual futures on stocks, commodities, and currencies trade around the clock, while the cash and futures markets that anchor their values, their benchmarks, keep their own calendars. When the benchmark’s price is unavailable, the trading venue must still publish a reference price, which it calls the *oracle*, to set the funding payments between long and short positions and to value positions for margin and liquidation. It can hold the last external price, the last price it received from the benchmark, which grows stale as

*IE University. Email: xiaojian.li@student.ie.edu.

†The Chinese University of Hong Kong. Email: jianyuan@cuhk.edu.hk.

news accumulates, or follow the local price in its own order book, which stays current but is set by the few traders who remain. We find that the local price anticipates most of the benchmark’s move: over weekends, a 1 percent change in the local price maps into a 0.7 to 0.9 percent change in the reference price once the external price returns. The local price is noisy, however, and the benchmark’s closure itself thins the order book that produces it. A reference price should therefore follow the local price more closely the longer the external price has been unavailable.

The setting separates the loss of the external price from the closure of the benchmark. The venue, TradeXYZ, lists perpetual futures on Hyperliquid, and its documented rules specify when the oracle follows the external price, which we call external pricing, and when it follows an internal rule that updates from the order book, which we call internal pricing. The two states do not follow the benchmark’s trading hours. Of all minutes in which a contract’s benchmark is closed, 53 percent remain in external pricing, as on U.S.-equity weeknights, and 47 percent are in internal pricing, as over weekends, holidays, and non-U.S. overnight closures. Periods of internal pricing, or episodes, therefore isolate the problem in the first paragraph: the reference price must be set without the external price. Our sample covers 103 contracts on U.S. and international equities, exchange-traded funds, equity indices, commodity futures, spot metals, and foreign exchange from December 15, 2025 through August 31, 2026, which executed about \$179 billion of order-book trading while their benchmarks were closed.

Most of a local price change is information rather than noise. In a simple signal-extraction framework (Section 3), the local price equals the price the benchmark would set plus transitory noise, and the slope of a regression of the post-resumption oracle on the local price change measures the signal share of that change. The weekend slopes of 0.7 to 0.9 therefore mean that roughly four-fifths of a typical local price change is information that the external price later confirms. The local price also forecasts: adding it to a model based only on information available before the episode lowers out-of-sample mean squared error by 17 to 73 percent in the asset classes with enough history to estimate the model. The remaining fifth behaves like noise. Slopes lie below one in every asset class, and in U.S. equities, international equities, and ETFs, local price changes partly reverse after resumption, by 20 to 36 percent of the change before it.

The local price is set in a thin order book, and the benchmark’s closure makes it thinner. Over weekends, displayed depth confirms that a \$10,000 purchase and sale would each execute within 10 basis points of the local price in 44 percent of minutes, from about a quarter for international equities and ETFs to three-quarters for commodity futures, spot metals, and foreign exchange. To isolate the effect of the closure itself, we use the fact that holidays close some benchmarks while others remain open in the same hour. Relative to contracts whose benchmark is open, a holiday closure halves top-of-book depth and cuts order-book trading by about 78 percent, consistent with liquidity providers withdrawing when they cannot hedge in the closed benchmark (Ho and Stoll, 1981; Grossman and Miller, 1988). In the framework, thinner order books mean more noise, and across the seven asset classes, implied noise is lower where order books are costlier to move.

A few large accounts bring the information into the local price, and when the benchmark closes, accounts trade less and supply less liquidity in the affected contract. Because every execution carries a persistent account identifier, we can follow about 400,000 accounts. Over

weekends, the local price moves 2.9 basis points in the direction of takers’ trades, which execute against resting orders, within 30 minutes of the benchmark’s reopening, and the 1 percent of takers with the largest notional account for 95 percent of this information. Programmatic market makers, who trade through software rather than by hand, supply 76 percent of maker notional, lose 3.0 basis points to these trades, and recover it from the spread. When a holiday closes the benchmark, the same accounts, in the same hour, become 0.76 percentage points less likely to trade the affected contract than contracts whose benchmarks remain open, about a fifth of their usual 3.9 percent rate, and they post less liquidity that is filled.

The weight that best predicts the reference price rises with episode length, as the framework implies: little news accumulates in short episodes, so local noise dominates, whereas over long episodes the last external price grows stale. We evaluate rules that add a fraction λ of the local price change to the last external price, choosing λ each month from outcomes already observed. The selected weight is about 0.1 for maintenance breaks but 0.8 to 0.9 for overnight, weekend, and holiday episodes, and the frozen rule ($\lambda = 0$) beats full pass-through ($\lambda = 1$) only in short episodes. The selected rule improves on the venue’s oracle by 4.2 basis points (95 percent interval -6.3 to -2.0), or 3.5 basis points when the venue’s documented price bounds are imposed on it. Relative to the frozen rule the gain is 43 basis points, compared with 2 basis points relative to full pass-through.

These errors matter because every open position is revalued at once when the reference price jumps to the external price. Over our 260-day sample, the frozen rule would have left \$1.1 billion more of such revaluation than the venue’s oracle, about 70 percent more, and resumptions with large reference-price errors account for most of the forced trading, mainly liquidations, that follows. Regulators have raised the same concern. As it considered 24/7 trading and perpetual futures in 2025 and 2026, the Commodity Futures Trading Commission stated that a perpetual’s reference “must be reliable at every funding interval, without interruption, for as long as the contract remains active” (Commodity Futures Trading Commission, 2025a,b, 2026a,b,c).

Our first contribution is to research on price formation outside regular trading hours. Pre-opening, after-hours, and overnight trading incorporate information, typically with lower volume and noisier prices (French and Roll, 1986; Biais, Hillion, and Spatt, 1999; Cao, Ghysels, and Hatheway, 2000; Barclay and Hendershott, 2003, 2004, 2008); price discovery in the round-the-clock Treasury market concentrates in particular sessions (He et al., 2009); and returns earned inside and outside the regular session differ systematically (Lou, Polk, and Skouras, 2019; Hendershott, Livdan, and Rösch, 2020; Boyarchenko, Larsen, and Whelan, 2023). Weekend cryptocurrency returns even predict Monday equity returns (Mourey, Shahrou, and Şoiman, 2025). In these settings the off-hours price is eventually replaced by the regular-session price, so its noise is temporary. In ours, the off-hours price feeds the reference price that sets funding and margin while it lasts, so the share of it that is signal matters in its own right. We measure that share, and show that the mapping slope, the reversal after resumption, and the optimal weight are three views of the same ratio of news to noise, in the tradition of separating informative from transitory price components (Roll, 1984; French and Roll, 1986).

Second, we contribute to research on how liquidity supply depends on the ability to manage inventory risk (Ho and Stoll, 1981; Grossman and Miller, 1988) and on trading

when related markets are closed or halted (Brock and Kleidon, 1992; Gerety and Mulherin, 1992; Lee, Ready, and Seguin, 1994; Corwin and Lipson, 2000; Barclay and Hendershott, 2004). Halts and closures in that work usually coincide with news or with the close of the market being studied. Asynchronous national holidays instead close only the benchmark, on a schedule set long in advance, while the derivative keeps trading. The closure halves depth and cuts trading, and the same accounts trade less in affected contracts than in contracts whose benchmarks remain open, consistent with liquidity providers that cannot hedge in the closed benchmark. Continuous trading therefore does not make liquidity continuous: a derivative’s liquidity inherits its benchmark’s calendar.

Third, we contribute to research on the design of financial benchmarks and oracles. Benchmark research studies how noise and manipulation incentives in the market that sets a reference price affect the contracts tied to it (Kumar and Seppi, 1992; Abrantes-Metz et al., 2012; Duffie and Stein, 2015; Griffin and Shams, 2018; Duffie and Dworczak, 2021; Brugler, Khomyn, and Putniņš, 2025), and regulatory principles require benchmarks to be anchored in observable transactions (International Organization of Securities Commissions, 2013; Financial Stability Board, 2014). Research on perpetual futures and blockchain oracles studies funding mechanisms, price discovery in crypto derivatives, and the reliability of on-chain price feeds (Shiller, 1993; Alexander et al., 2020; Angeris et al., 2023; He et al., 2024; Cong et al., 2025; Akerer, Hugonnier, and Jermann, 2026). These literatures largely take the benchmark market as open. When it is closed, the only observable transactions are in the derivative’s own order book, and the design question becomes how much to rely on them. Our answer depends on the state: the weight on local prices should rise with episode length, as the selected weights do, and venues should disclose when external pricing stops and resumes, which we could reconstruct only from published schedules.

Our evidence covers one venue over eight and a half months. The rule comparison measures forecast accuracy on realized price paths, so it holds trading behavior fixed, and it imposes the venue’s price bounds at their documented values; Section 8 discusses both choices. Section 2 describes the setting, Section 3 the framework, and Section 4 the data; Sections 5–8 present the results in the order above, and Section 9 concludes.

2 Institutional Setting

The venue’s reference price follows an external price on a schedule that differs from the benchmark’s trading hours. When it does not, it follows a rule that places a weight below one on the local price, and that weight varies with the state of the order book.

2.1 The venue and its contracts

TradeXYZ lists perpetual futures contracts that reference equities, equity indices, commodities, metals, and exchange rates. The contracts provide price exposure without ownership or delivery rights, and they remain tradable when the referenced cash or futures market is closed. They trade on Hyperliquid’s order-book, position, and settlement infrastructure, while TradeXYZ specifies the contracts and their reference-price rules. Our sample contains only TradeXYZ contracts and excludes Hyperliquid’s native cryptocurrency perpetuals and

contracts listed by other issuers on Hyperliquid.¹

Two prices govern the economics of a position. The *oracle* is the contract’s reference price. It anchors funding, an hourly transfer between long and short positions that responds to the contract’s premium over the oracle; a positive rate requires longs to pay shorts.² The *mark price* is the valuation used for margin, liquidation, and unrealized profit and loss. It is a robust median of local order-book prices and an oracle-based component that adds a smoothed local basis to the oracle (Internet Appendix Section A.1). A stale or distorted oracle therefore affects funding directly and margin and liquidation through the mark price. Throughout the paper, the *benchmark* is the contract’s principal cash or futures market, and a *benchmark closure* is a scheduled period in which the benchmark is closed while the perpetual keeps trading.

2.2 Two clocks

The benchmark and the oracle run on different clocks. The *benchmark clock* records whether the benchmark is scheduled to trade. It follows the relevant exchange calendar, including holidays, early closes, local time zones, and daylight-saving-time transitions; for the three foreign-exchange and four spot-metal contracts, it follows TradeXYZ’s published OTC reference hours, and results excluding these seven contracts are similar (Section 8). The *source clock* records whether the contract specification calls for the oracle to follow an external price or the venue’s internal rule. We refer to the two source states as *external pricing* and *internal pricing*. We reconstruct both clocks minute by minute from exchange calendars and the venue’s documented session rules; the data record the oracle but not a flag for each switch, so source states and switch times are those the schedules specify.

Crossing the two clocks gives the four states in Table 1, and the two clocks differ substantially. Of the 15.4 million benchmark-closed contract–minutes in our sample, 53 percent remain in external pricing and 47 percent are in internal pricing. The split varies sharply across products (Figure 1, Panel B). U.S. equities, ADRs, and ETFs are in external pricing in about 62 percent of their benchmark-closed minutes, because their documented external-pricing schedule runs continuously on weekdays, well beyond the benchmark’s regular session; they switch to internal pricing mainly over weekends and holidays. Commodity futures, spot metals, and foreign exchange, by contrast, are internally priced in nearly all of theirs.

The two clocks can diverge sharply within a single contract. Panel A of Figure 1 shows two contracts over one winter week. AAPL’s regular exchange session runs from 14:30 to 21:00 UTC, but its oracle remains in external pricing from Monday 01:00 to Saturday 01:00 UTC and is internally priced only over the weekend. For WTI crude oil, the two clocks nearly coincide: the daily maintenance break and the weekend closure of the futures market are also internal-pricing episodes. A benchmark closure and the loss of the external price are therefore different events. The analysis treats benchmark-closed minutes in external pricing

¹TradeXYZ contracts carry the `xyz:` prefix under Hyperliquid’s HIP-3 framework for third-party perpetual markets. The institutional description follows Hyperliquid’s documentation on HIP-3, funding, and margining and TradeXYZ’s documentation on oracle prices, external prices, the oracle time constant, and contract specifications.

²See Hyperliquid’s *Funding* documentation for the payment interval and the role of the oracle in the funding calculation.

Table 1: Benchmark and Oracle-Source States

Benchmark / source	External Price field	Oracle	Share of minutes (%)	Typical periods
Open / external pricing	Updates	External price	26.05	Regular benchmark sessions
Closed / external pricing	Updates	External price	38.87	U.S. equity and ETF weeknights; some equity-index closures
Closed / internal pricing	Fixed at last external price	Impact-price rule, equation (1)	34.26	Weekends, holidays, futures maintenance, lunch breaks, non-U.S. overnight closures
Open / internal pricing	Fixed at last external price	Impact-price rule	0.03	Contract-specific clock boundaries

Notes: Shares are computed over all 21,025,440 contract-minutes in the sample; the remaining 0.79 percent have an unknown benchmark state or an unverified source state. The External Price field is the venue’s record of the last external price. In every state, the mark price is a robust median of local order-book prices and an oracle-based component; Internet Appendix Table IA.1 defines each price object and its role.

and those in internal pricing as different regimes, and it studies the absence of the external price directly in the latter.

2.3 How the oracle updates during internal pricing


During external pricing, the oracle equals the external price. During internal pricing, the venue holds the last external price in a field it calls the External Price, and the oracle starts from that value and updates from *impact prices*: the average execution prices of a configured notional on the bid and ask sides of the order book. Let S_{t-} and S_t denote the oracle immediately before and after an update at time t . The documented rule is

$$\begin{aligned} \text{IPD}_t &= \max(P_{\text{impactBid},t} - S_{t-}, 0) - \max(S_{t-} - P_{\text{impactAsk},t}, 0), \\ S_t &= S_{t-} + \alpha_t \text{IPD}_t, \end{aligned} \tag{1}$$

where the impact-price displacement IPD_t measures how far the local impact prices lie outside the current oracle and $0 \leq \alpha_t < 1$ is set by elapsed time and a smoothing time constant.

Three features of this rule matter for the analysis. First, it has a dead zone. When the oracle lies between the impact bid and the impact ask, $\text{IPD}_t = 0$ and the oracle does not move; it adjusts only when local impact prices move outside it, and then only toward the nearer boundary. Second, it adjusts partially. Because $\alpha_t < 1$, a local move enters the oracle gradually. Third, it depends on depth. If one side of the order book cannot fill the configured notional, the documented rule sets that side’s contribution to zero; if neither side can, the oracle does not move. In a thin order book, the oracle therefore stays close to the last external price. In the terms of the framework in Section 3, the oracle’s rule places a weight on the local price change that is below one, varies with the state of the order book, and falls toward zero when the order book cannot fill the impact notional. Section 8 compares this rule with fixed weights.

Figure 1: Benchmark and Oracle-Source Clocks



Notes: Panel A shows a pre-specified ordinary week (January 26 to February 1, 2026, in UTC) with no exchange holiday or daylight-saving transition. Dashed lines mark scheduled resumptions of external pricing. Panel B reports the share of benchmark-closed contract-minutes in each source state, pooled across all contracts ($N = 15,379,020$) and by asset class; each contract-minute receives equal weight. Unverified source states account for 0.02 percent of pooled minutes and are shown in red. Source states are reconstructed from documented rules and schedules.

Contract-level price bounds, which the venue calls Discovery Bounds, further limit how far the oracle and mark price can move from the last external price. Their widths depend on contract parameters, and some contracts allow a limited number of resets.

When external pricing resumes, the documentation states that the oracle returns to the external price on the next update. We call this scheduled switch a *resumption* and distinguish it from *benchmark reopening*, the scheduled open of the benchmark itself. The post-resumption oracle therefore provides a common target against which both local prices and alternative weights can be evaluated.

3 A Signal-Extraction Framework

When external pricing stops, the venue must decide how much of the local price change to pass into its reference price. The best fixed weight is the signal share of the local price change: it rises with the time since the last external price and falls with the noise in the order book that sets the local price. The framework takes liquidity and news intensity as given. Its role is to give the empirical objects in Sections 5–8 a common interpretation and to state four predictions that the evidence can confront.

3.1 Setup

Consider one contract and one internal-pricing episode that begins at c and ends at resumption o . All prices are in logs. Let v_t denote the external price that the source would report if it were live, and normalize the last external price so that $\log F = v_c$.

Accumulating news. While external pricing is off, v_t follows a martingale with independent increments,

$$\text{Var}(v_t - v_c) = \sigma^2 \tau_t, \quad \tau_t \equiv t - c, \quad (2)$$

so the staleness of the last external price grows with episode length, and σ^2 measures the intensity of value-relevant news during the episode.

Noisy local prices. The local price equals value plus a stationary, mean-zero deviation, the local noise, that is independent of news,

$$m_t = v_t + \eta_t, \quad \text{Var}(\eta_t) = \omega^2, \quad \text{Cov}(\eta_t, \eta_{t-k}) = \gamma(k). \quad (3)$$

The deviation η_t collects inventory pressure, transitory order-flow imbalances, stale quotes, and any strategic displacement of the order book (Roll, 1984; Hasbrouck, 1991). Thinner order books and fewer active liquidity providers allow larger and more persistent deviations, so the noise process is the counterpart of the liquidity measured in Section 6.

The local price change at evaluation point $u \leq o$ is $X(u) = m_u - m_c = (v_u - v_c) + (\eta_u - \eta_c)$, as in equation (8). The post-resumption oracle is $Y = v_o - v_c + \zeta$, where ζ is measurement error in this target, such as that from averaging over the first 30 minutes, uncorrelated with the local path. The local noise accumulated by u is

$$\psi_u^2 \equiv \text{Var}(\eta_u - \eta_c) = 2 [\omega^2 - \gamma(\tau_u)]. \quad (4)$$

3.2 The optimal weight on local prices

For the rule family evaluated in Section 8, $\log P_\lambda(u) = \log F + \lambda X(u)$, the forecast error is $\lambda X(u) - Y$, and mean-squared error is minimized at

$$\lambda^*(u) = \frac{\text{Cov}(Y, X(u))}{\text{Var}(X(u))} = \frac{\sigma^2 \tau_u}{\sigma^2 \tau_u + \psi_u^2}. \quad (5)$$

The optimal weight is the signal share of the local price change. It equals the population slope of a regression of Y on $X(u)$, so a slope below one need not indicate biased local prices or information lost at resumption; it arises whenever local prices contain noise, as measurement error attenuates a regression slope.

The two polar rules have transparent losses:

$$\begin{aligned} \text{MSE}(0) &= \sigma^2 \tau_o + s_\zeta^2, \\ \text{MSE}(1) &= \sigma^2(\tau_o - \tau_u) + \psi_u^2 + s_\zeta^2, \\ \text{MSE}(\lambda^*) &= \sigma^2(\tau_o - \tau_u) + \lambda^*(u) \psi_u^2 + s_\zeta^2, \end{aligned} \quad (6)$$

where $s_\zeta^2 = \text{Var}(\zeta)$. The frozen rule bears the full staleness cost, which grows linearly with episode length; full pass-through bears the local noise instead. The frozen rule beats full pass-through if and only if $\sigma^2 \tau_u < \psi_u^2$. The best fixed weight improves on the two polar rules by $\lambda^* \sigma^2 \tau_u$ and $(1 - \lambda^*) \psi_u^2$, respectively. The venue's oracle does not use a fixed weight: it applies the state-dependent, depth-sensitive rule described in Section 2, which Section 8 compares with the fixed weights in equation (5).

The same weight determines how far a trader can move the reference price. A trader who displaces the local price by d moves the rule-based reference price by λd . Whose trades move the local price, the question of Section 7, therefore matters for the reference price in proportion to λ .

3.3 Mapping to the empirical objects

The estimates correspond to the framework's objects with two adjustments. The mapping regression in Section 5 also controls for the initial basis between the local price and the oracle, which in the framework equals η_c .³ Its slope is therefore an upper bound on $\lambda^*(u)$ for the rule family above. The rule comparison in Section 8 uses absolute rather than squared error. Under joint normality, the conditional median of Y given $X(u)$ equals its conditional mean, so the same weight minimizes both losses.

The weight and the forecast gain measure different things. News arriving between u and o and target measurement error raise the variance of Y and lower the incremental out-of-sample R^2 in equation (11), but they leave $\lambda^*(u)$ unchanged. A market can therefore show a slope close to one together with a modest forecast gain, or the reverse.

³With the basis control, the population slope is $\sigma^2 \tau_u / [\sigma^2 \tau_u + \text{Var}(\eta_u | \eta_c)]$. Because $\psi_u^2 - \text{Var}(\eta_u | \eta_c) = [\omega - \gamma(\tau_u)/\omega]^2 \geq 0$, this slope weakly exceeds $\lambda^*(u)$, with equality when the initial deviation persists fully, $\gamma(\tau_u) = \omega^2$.

3.4 Predictions

Equations (5) and (6) yield four predictions, listed in the order in which Sections 5, 6, and 8 test them.

- P1.** *Traces of noise.* Local noise produces two patterns. The mapping slope lies below one (Table 3). And if deviations persist for several minutes, local price changes partially reverse after resumption, even when the change is measured a few minutes before the boundary (Table IA.8). Internet Appendix Section A.2 derives the reversal slope.
- P2.** *Information accumulation.* Within an episode, $\lambda^*(u)$ rises as the episode proceeds whenever accumulated noise ψ_u^2 grows less than proportionally with elapsed time, as it does when η_t follows a stationary autoregressive process. The mapping slope and the forecast gain should therefore strengthen from early to late evaluation points. Figure 2 tests this prediction.
- P3.** *Liquidity.* The weight depends on accumulated news relative to local noise, not on either alone. Holding news intensity fixed, events that thin the order book raise ψ_u^2 and lower λ^* ; a deep order book implies a high weight only if the episode also carries substantial news. Section 6 tests this prediction.
- P4.** *Episode length.* Across episodes, the loss advantage of full pass-through over the frozen rule, $\sigma^2\tau_u - \psi_u^2$, rises with episode length. The frozen rule should dominate in short episodes and local updating in long ones, and the weight selected from past outcomes should rise with episode length. Section 8 tests this prediction by closure type, which sorts episodes by length.

4 Data and Sample Construction

Our data follow 103 TradeXYZ contracts in seven asset classes from December 15, 2025 through August 31, 2026, covering 10,043 benchmark closures and \$179 billion of order-book trading. Three linked panels follow the same closure at the level of the contract-minute, the closure, and the account. During these closures the order book is thin: the median contract-minute has no order-book trading, and median top-of-book depth is \$1,960.

4.1 Data sources

The data combine four venue records. The best-bid-and-offer feed gives the standing inside quotes and displayed quantities, and one-minute snapshots of the first 20 price levels on each side of the book measure displayed depth for Section 6. The contract-state feed reports the oracle, mark price, open interest, and funding rate. Account-level executions identify the trader, the maker or taker role, the direction, and the notional amount, and they distinguish voluntary trades from forced trades executed on or outside the order book. Initial position snapshots and executions over the 28 days before the sample begins complete the account histories. The sample begins with the one-minute order-book archive; Internet Appendix Section A.3 describes contract coverage.

The seven asset classes are commodity futures, foreign exchange, equity indices, international equities, spot metals, U.S. equities (including American depositary receipts), and U.S. exchange-traded funds.

4.2 Benchmark closures and episodes

Two schedules define the units of analysis. The benchmark’s schedule defines benchmark closures. For each contract, we convert the benchmark’s published trading hours, including exchange holidays, early closes, lunch breaks, maintenance breaks, and daylight-saving transitions, into a minute-by-minute indicator in UTC. A *benchmark-closure window* is a maximal interval during which this indicator is closed. We classify windows into five closure types: overnight, weekend, holiday, lunch break, and maintenance break.

The oracle’s documented schedule defines episodes. Each minute carries the oracle’s scheduled state, external or internal pricing (Section 2), and an episode is a maximal interval of internal pricing, from the last external price at c to resumption at o , as in Section 3. Episodes and benchmark-closure windows usually overlap but need not coincide. On U.S.-equity weeknights, for example, the benchmark is closed while the oracle remains in external pricing; we therefore analyze these closures separately from internally priced weekends, holidays, and breaks. Event time is measured from scheduled boundaries.

4.3 Analysis samples

Contract-minute panel. One observation is contract j in UTC minute t . It contains the executions over $(t - 1 \text{ minute}, t]$ and the latest quotes, oracle, and contract state observed by t , so order flow is paired with the market state at the end of the minute. For a valid two-sided order book, the local price (the midquote) and the quoted spread are

$$P_{jt}^{mid} = \frac{P_{jt}^{ask} + P_{jt}^{bid}}{2}, \quad Spread_{jt} = 10,000 \frac{P_{jt}^{ask} - P_{jt}^{bid}}{P_{jt}^{mid}}. \quad (7)$$

The panel also records top-of-book depth, one-minute local price changes, the oracle and mark price, the basis between the local price and the oracle, quote age, order-book trading, trade counts, and the number of active makers. Order-book trading is the dollar value of voluntary and forced executions on the order book, counting each execution once on its taker side. The panel contains 21,025,440 contract-minutes, of which 15,226,677 fall in complete benchmark-closure windows.

Contract-closure-window panel. Each observation follows one contract through a complete benchmark-closure window. It records the local price path during the closure, the oracle over the 30 minutes after benchmark reopening, and window averages of spreads, depth, trading, quote updates, and maker participation. Analyses of episodes rebuild prices from the minute panel at each episode’s own boundaries. The panel contains 10,043 contract-windows.

Account-contract-closure-window panel. One observation is account i in contract j during benchmark-closure window w . The population includes accounts that traded the

contract in the preceding 28 days or held an opening position, with zero activity recorded when they do not trade, and accounts that first trade during the window. Characteristics measured over the preceding 28 days and fixed at the start of the closure include specialization in the contract, share of the contract’s activity, number of contracts traded, the share of fills that carry a marker of programmatic trading (Section 7), and the share of notional traded as maker. Markouts value each execution at the local price 30 minutes later. The panel contains 41,860,740 observations for 400,709 accounts.

Each analysis draws its sample from these panels. Section 5 uses weekend episodes; its mapping regressions use a paired sample of 1,114 episodes (Internet Appendix Table IA.3). Section 6 uses 1,846 weekend episodes with displayed-depth data and 13 holiday events on four benchmark calendars. Section 7 uses the account panel around the same holiday events and on weekends. Section 8 evaluates rules on 3,749 episodes in 102 contracts over 32 resumption weeks.

4.4 Summary statistics

Table 2 describes the sample. U.S. equities account for 64 of the 103 contracts and 6,395 of the 10,043 benchmark-closure windows, so pooled statistics mostly reflect U.S. equities; Section 5 therefore estimates by asset class. Overnight closures are the most frequent type, with 6,707 windows. Weekends are fewer, with 1,774 windows, but last 63 hours on average.

The order book is thin during benchmark closures (Panels B and C). The median contract-minute has no order-book trading, and median top-of-book depth is \$1,960. The median closure window carries \$1.58 million of order-book trading, against a mean of \$17.9 million, so a few windows carry most of the trading. Section 3 links such thinness to local noise, and Section 6 measures it directly.

Few accounts trade in any given window (Panel D). An account trades voluntarily in 8.75 percent of account–contract–window observations and supplies liquidity in 3.26 percent, so the account analyses in Section 7 separate the decision to trade from the amount traded.

Table 2: Sample Composition and Summary Statistics

Panel A: Sample Composition

Sample group	Contracts	Benchmark-closure windows	Observed benchmark-closure minutes	Mean scheduled duration (hours)	Order-book trading (\$bn)
<i>A1: By Asset Class</i>					
Commodity futures	4	618	410,609	11.07	14.24
Foreign exchange	3	68	195,720	47.97	0.47
Equity indices	4	585	846,476	24.16	56.64
International equities	11	948	1,175,839	20.68	25.63
Spot metals	4	586	334,440	9.51	3.54
U.S. equities/ADRs	64	6,395	10,846,286	28.40	68.09
U.S. ETFs	13	843	1,417,307	28.23	10.68
<i>A2: By Closure Type</i>					
Holiday	77	323	1,592,610	82.18	11.15
Lunch break	9	374	24,900	1.11	0.14
Maintenance	7	865	51,900	1.00	1.85
Overnight	96	6,707	6,880,317	17.26	114.60
Weekend	103	1,774	6,676,950	62.73	51.56

Panel B: Contract-Closure-Window Panel

Variable	Unit	<i>N</i>	Mean	Std. dev.	P25	Median	P75
Closure duration	hours	10,043	25.38	22.05	17.50	17.50	17.50
Local price change to reopening	bp	10,034	13.62	271.49	-85.54	4.58	102.69
Oracle change after reopening	bp	10,035	5.29	309.89	-111.82	0.44	116.88
Mean quoted spread	bp	10,042	13.38	22.01	2.25	5.85	14.59
Mean top-of-book depth	\$thousand	10,042	15.87	54.17	2.94	6.46	13.87
Order-book trading	\$million/window	10,043	17.85	62.46	0.35	1.58	6.62
Trading-minute share	percent	10,043	52.32	34.30	19.36	49.62	87.90
Minutes with a quote update	percent	10,043	80.17	23.44	69.14	90.29	98.38
Effective maker count	count	10,036	10.06	7.64	4.89	8.27	13.22

Table 2: Sample Composition and Summary Statistics (continued)

Panel C: Contract–Minute Panel

Variable	Unit	<i>N</i>	Mean	Std. dev.	P25	Median	P75
Quoted spread	bp	15,223,095	14.62	33.53	1.77	5.35	14.64
Top-of-book depth	\$thousand	15,223,095	15.10	125.00	0.59	1.96	7.13
Order-book trading	\$thousand/minute	15,226,677	11.77	108.61	0.00	0.00	1.03
Order-book trades	count/minute	15,226,677	9.59	40.44	0.00	0.00	5.00
Active executed makers	count/minute	15,226,677	2.60	6.04	0.00	0.00	3.00
Absolute one-minute price change	bp	15,223,082	3.23	8.85	0.00	0.70	3.16
Local price–oracle basis	bp	15,223,095	2.89	27.77	-3.97	2.11	8.86
Quote age	seconds	15,223,366	145.33	961.17	0.93	5.79	30.63

Panel D: Account–Contract–Closure–Window Panel

Variable	Unit	<i>N</i>	Mean	Std. dev.	P25	Median	P75
Pre-closure population member	percent	41,860,740	98.04	13.86	100.00	100.00	100.00
Voluntary trading in window	percent	41,860,740	8.75	28.26	0.00	0.00	0.00
Account notional, all executions	\$thousand	41,860,740	8.60	416.64	0.00	0.00	0.00
Maker activity	percent	41,860,740	3.26	17.76	0.00	0.00	0.00
Maker notional	\$thousand	41,860,740	4.30	219.90	0.00	0.00	0.00
Pre-closure contract specialization	percent	40,552,884	34.33	37.84	3.12	15.11	62.83
Pre-closure share of contract activity	percent	41,860,066	0.02	0.39	0.00	0.00	0.00
Pre-closure contracts traded	contracts	40,552,884	8.74	11.39	2.00	5.00	11.00
Pre-closure programmatic share	percent	40,552,884	24.23	41.37	0.00	0.00	48.54
Pre-closure maker propensity	percent	40,552,884	18.03	27.41	0.00	0.00	32.73
Maker 30-minute markout	bp	1,364,418	3.85	84.01	-8.72	0.48	13.50
Taker 30-minute markout	bp	3,171,962	-2.11	73.44	-13.10	-0.74	9.75

Notes: The sample runs from December 15, 2025 through August 31, 2026. A benchmark-closure window is a maximal interval during which the contract’s benchmark is closed under its published schedule (Section 2). Panel A counts contracts, windows, and minutes by asset class and by closure type; contract counts are not additive across closure types. Mean scheduled duration is the equal-weighted mean of the scheduled time from the last pre-closure boundary to benchmark reopening. Order-book trading is the dollar value of voluntary and forced executions on the order book, counting each execution once on its taker side; liquidations executed outside the order book are excluded. Panel B weights closure windows equally. The local price change runs from the last pre-closure boundary to one minute before benchmark reopening, and the oracle change averages the oracle over the 30 minutes after reopening. Price changes in Panel C are changes in the local price. Panel C uses the 15,226,677 contract–minutes in complete closure windows, including the minute of benchmark reopening. Panel D includes accounts that traded the contract in the preceding 28 days or held an opening position, with zero activity when they do not trade, and accounts that first trade in the window. Pre-closure characteristics are measured over the preceding 28 days; the programmatic share is the share of fills that carry a client order identifier (Section 7). Markouts value each execution at the local price 30 minutes later and are reported for account–windows with positive matched notional, weighted equally. Observation counts vary across variables with the availability of quotes, endpoints, and account histories. Internet Appendix Section A.3 reconciles the minute counts with Figure 1 and the closure calendar.

5 Information in Local Prices

About four-fifths of a weekend local price change is information about the post-resumption oracle. The mapping slope lies between 0.68 and 0.93 in every asset class, adding the local price lowers out-of-sample forecast error by 17 to 73 percent in every class with enough history to train the forecasting model, and both measures strengthen as the episode proceeds. The remaining fifth is noise, and part of it reverses after resumption. The tests use weekend episodes: long, regularly scheduled periods of internal pricing during which the benchmark is closed. U.S.-equity weeknights are excluded because their oracle remains in external pricing.

5.1 Measurement design

For contract j and weekend episode w , let c_{jw} denote the boundary at which internal pricing begins and o_{jw} the scheduled resumption. Let $M_{jw}(u)$ be the local price at evaluation point u of the episode, the level of m_u in Section 3, and $O_{jw}(c)$ the last oracle before internal pricing begins, which equals the last external price. The local price change is

$$X_{jw}(u) = 10,000 [\log M_{jw}(u) - \log M_{jw}(c)], \quad (8)$$

evaluated at 25, 50, and 75 percent of the scheduled episode and at the final minute before resumption, which is the main specification. The target is the post-resumption oracle, the price at which positions are revalued when external pricing resumes. With $\bar{O}_{jw}^{+30}$ the mean oracle over minutes 1 through 30 after o_{jw} ,

$$Y_{jw} = 10,000 \left[\log \bar{O}_{jw}^{+30} - \log O_{jw}(c) \right]. \quad (9)$$

The sample requires fresh prices, at most three minutes old, at all four path points and in all 30 minutes of the target window, with the initial prices and target observations dated in external pricing. The same requirement applies to the 30 minutes after benchmark reopening, so that Internet Appendix Section A.4 can compare the two endpoints on the same episodes. The resulting sample contains 1,114 weekend episodes (Internet Appendix Table IA.3); retained episodes have fresher quotes and narrower spreads than excluded ones (Internet Appendix Table IA.4).

For each asset class, we estimate

$$Y_{jw} = \alpha_j + \beta X_{jw}(u) + \Gamma' Z_{jw}^{\text{pre}} + \varepsilon_{jw}, \quad (10)$$

with contract fixed effects α_j and predetermined controls Z_{jw}^{pre} : episode length, the initial basis between the local price and the oracle, and oracle returns over the preceding 30 minutes and 24 hours. The basis control aligns the local-price predictor with the oracle target. In the framework, β is the signal share of the local price change; with the basis control, its population value is an upper bound on the optimal weight $\lambda^*(u)$ in equation (5). We test $\beta = 0$, $\beta = 1$, and equivalence to one within $[0.8, 1.2]$ (Schuirmann, 1987), using a null-imposed wild cluster bootstrap- t over resumption weeks, and report week-bootstrap intervals that resample resumption weeks.

To test real-time predictive content, we compare two forecasting models that differ by one piece of information. The baseline model uses only information available before internal pricing begins: episode length, the initial basis, the two prior oracle returns, and contract and closure-type indicators. The enhanced model adds the local price change observed by the evaluation point, with a separate coefficient for each closure type. Both models are re-estimated at the start of each calendar month, using only episodes whose target windows have already ended, and then forecast subsequent weekend episodes. The forecast gain is

$$R_{\text{OOS,inc}}^2 = 1 - \frac{\text{MSE}_{\text{enhanced}}}{\text{MSE}_{\text{baseline}}}, \quad (11)$$

so a value of 0.60 means that the local price lowers mean squared forecast error by 60 percent. We report one-sided Clark–West statistics for nested models (Clark and West, 2007), together with the number of test weeks, the share of weeks with lower forecast error, and leave-one-week-out ranges.

5.2 Most of a local price change is information

The local price change maps into the post-resumption oracle with a slope of about 0.8 (Table 3, Panel A). The slope is statistically different from zero in every asset class and ranges from 0.675 for foreign exchange to 0.934 for equity indices; for U.S. equities, the largest class, it is 0.832. In the framework, a slope of 0.8 means that about four-fifths of the residual variation in the local price change is signal and one-fifth is noise. Internet Appendix Table IA.2 converts the slopes into implied news and noise standard deviations by asset class.

The slopes lie below one, as prediction P1 implies. The wild bootstrap rejects a unit slope for foreign exchange and U.S. equities, and no class establishes equivalence to one within $[0.8, 1.2]$; for U.S. equities the week-bootstrap interval is $[0.732, 0.936]$.

Part of the noise reverses after resumption, the second trace that prediction P1 implies. Internet Appendix Table IA.8, on a broader resumption sample, regresses the local price change over the 30 minutes after resumption on the local price change up to five minutes before it. The reversal slope is about -0.2 for international and U.S. equities and -0.36 for U.S. ETFs, with week-bootstrap intervals that exclude zero; the slopes for the other four classes are imprecise. Because the reversal is measured after a five-minute gap, it reflects deviations that persist for several minutes rather than bid–ask bounce (Internet Appendix equation (19)). Attenuation and reversal are thus two traces of the same noise component.

Local price dynamics are similar whether or not the external price is available. On a matched sample of 42 contracts, variance ratios (Lo and MacKinlay, 1988) and first-order autocorrelations of the local price are statistically indistinguishable between spells of external and internal pricing (Internet Appendix Table IA.7). The transitory component in local prices is therefore present in both states.

5.3 Local prices improve real-time forecasts

Adding the local price lowers out-of-sample mean squared error by 17 to 73 percent in every asset class with enough history to train both models (Table 3, Panel B): by 58 percent for

Table 3: Information in Local Prices

Panel A: Mapping into the Post-Resumption Oracle

	(1) Commodity futures	(2) FX	(3) Equity indices	(4) International equities
Local price change	0.792*** (0.114)	0.675*** (0.097)	0.934*** (0.146)	0.735** (0.123)
Wild p : $\beta = 1$	0.123	0.013	0.683	0.101
Wild equivalence-test p	0.525	0.841	0.286	0.754
Bootstrap 95% interval	[0.524, 0.986]	[0.518, 0.938]	[0.621, 1.219]	[0.366, 1.249]
Bootstrap draws used	999	999	999	999
Contract fixed effects	Yes	Yes	Yes	Yes
Predetermined controls	Yes	Yes	Yes	Yes
Episodes	103	44	65	40
Contracts	4	3	4	5
Resumption weeks	30	25	33	11
R^2	0.709	0.703	0.836	0.610

	(5) Spot metals	(6) U.S. equities and ADRs	(7) U.S. ETFs
Local price change	0.813*** (0.150)	0.832*** (0.049)	0.807*** (0.096)
Wild p : $\beta = 1$	0.309	0.004	0.073
Wild equivalence-test p	0.468	0.274	0.473
Bootstrap 95% interval	[0.546, 1.186]	[0.732, 0.936]	[0.575, 1.083]
Bootstrap draws used	999	999	999
Contract fixed effects	Yes	Yes	Yes
Predetermined controls	Yes	Yes	Yes
Episodes	90	715	57
Contracts	4	60	11
Resumption weeks	28	31	21
R^2	0.708	0.751	0.630

Notes: The dependent variable is the change in the oracle from the last external price to the mean oracle over the 30 minutes after resumption, in basis points. The regressor is the local price change from the start of internal pricing to one minute before resumption. All regressions include contract fixed effects and predetermined controls: episode length, the initial basis between the local price and the oracle, and oracle returns over the preceding 30 minutes and 24 hours. Standard errors, clustered by resumption week, are in parentheses. Stars test $\beta = 0$ at the 10, 5, and 1 percent levels with a null-imposed wild cluster bootstrap- t over resumption weeks; the same procedure tests $\beta = 1$ and equivalence to one within $[0.8, 1.2]$, using 9,999 weights, or all sign weights with at most 10 weeks. Bootstrap intervals resample resumption weeks (999 draws) and re-estimate the fixed effects in each draw.

commodity futures, 73 percent for equity indices, 17 percent for international equities, and 66 percent for spot metals. The gains do not rest on a few weeks. The enhanced model has lower error in 44 of the 54 test weeks, and dropping any single week leaves every gain positive.

Table 3: Information in Local Prices

Panel B: Out-of-Sample Forecast Gain

	(1) Commodity futures	(2) Equity indices	(3) International equities
Forecast gain (OOS R^2)	0.578	0.731	0.168
Clark–West p	0.010	0.002	< 0.001
Leave-one-week range	[0.548, 0.848]	[0.563, 0.788]	[0.098, 0.233]
Weeks with lower forecast error	16 / 18	11 / 12	4 / 6
Episodes	68	32	19
Test weeks	18	12	6
Forecast months	4	3	1
Contract indicators	Yes	Yes	Yes
Closure-type indicators	Yes	Yes	Yes
Predetermined controls	Yes	Yes	Yes
	(4) Spot metals	(5) U.S. equities and ADRs	(6) U.S. ETFs
Forecast gain (OOS R^2)	0.661	—	—
Clark–West p	0.003	—	—
Leave-one-week range	[0.605, 0.725]	—	—
Weeks with lower forecast error	13 / 18	—	—
Episodes	62	—	—
Test weeks	18	—	—
Forecast months	5	—	—
Contract indicators	Yes	Yes	Yes
Closure-type indicators	Yes	Yes	Yes
Predetermined controls	Yes	Yes	Yes

Notes: The baseline model uses predetermined controls and contract and closure-type indicators; the enhanced model adds the local price change, with a separate coefficient for each closure type. The forecast gain is the incremental out-of-sample R^2 of the enhanced model over the baseline. Both models are re-estimated at the start of each month on episodes whose target windows have ended. Training requires four weeks, 30 observations, 10 observations per effective parameter, and 20 residual degrees of freedom. Forecast gains and Clark–West statistics weight episodes equally; Clark–West p -values are asymptotic, with weekly HAC standard errors of at most four lags. Leave-one-week ranges drop one test week at a time with forecasts held fixed. Dashes indicate cells without enough training history.


The model with contract indicators needs more training history than foreign exchange, U.S. equities, and U.S. ETFs provide. A specification that pools contracts within asset class, which we added after the baseline results (Internet Appendix Table IA.6), covers 558

U.S.-equity weekend episodes over 23 test weeks and yields a gain of 76.5 percent.

5.4 The signal accumulates during the episode

The signal strengthens as the episode proceeds, as prediction P2 implies. Figure 2 repeats the mapping and forecast comparisons at 25, 50, and 75 percent of the episode and one minute before resumption, using at each point only the local price observed by then. Early in the episode, several classes have small or imprecise slopes and little forecast gain. By the final minute, all slopes lie between 0.675 and 0.934, and forecast gains reach 17 to 73 percent.

Figure 2: Information Accumulation during Internal Pricing



Notes: The sample is the weekend sample in Table 3. The four evaluation points are 25, 50, and 75 percent of the episode and one minute before scheduled resumption, and the target is the mean oracle over the 30 minutes after resumption. Panel A reports the coefficient on the local price change from asset-class regressions with contract fixed effects and predetermined controls; the dashed line marks a one-for-one mapping. Error bars are 95 percent week-bootstrap intervals from 999 resamples, with fixed effects re-estimated in each draw; an interval is shown when at least 499 resamples identify the coefficient. Panel B reports the forecast gain from monthly expanding-window forecasts. Foreign exchange and the two U.S. groups have no forecast-gain series because their training history is too short. Internet Appendix Figure IA.1 repeats both panels for the oracle after benchmark reopening.

Local prices therefore carry most, but not all, of the information that the post-resumption oracle reveals. Section 6 turns to the order book that produces this signal and to prediction P3: how thin is it, and does losing the benchmark make it thinner?

6 Liquidity and Market Quality

The local price comes from a thin order book, and losing the benchmark makes it thinner. Over weekends, a \$10,000 purchase and sale can both be confirmed within 10 basis points of the local price in only 44 percent of minutes, and implied noise is lower in asset classes whose order books are costlier to move, as prediction P3 implies. Holidays that close one benchmark while others stay open show that the closure itself withdraws liquidity: top-of-book depth falls by about half and order-book trading by about 78 percent. The closure that makes the last external price stale therefore also thins the order book that produces the local price.

6.1 Weekend order books are thin

Table 4 describes 1,846 weekend episodes in 103 contracts, each weighted equally, using two measures of the displayed order book. The support rate is the share of scheduled minutes in which both a \$10,000 purchase and an equal-sized sale, walked through the displayed levels, would execute within 10 basis points of the local price. The shift cost is the displayed notional needed to move the local price by 25 basis points; we take its fifth percentile within each episode, the most fragile minutes, and report the median across episodes. Because the archive records 20 levels on each side, both measures are lower bounds (Internet Appendix Table IA.9).

Weekend order books are thin and differ widely across products. Across all episodes, the support rate is 44.2 percent, and moving the local price by 25 basis points requires a median of \$10,965 in the most fragile minutes. International equities and U.S. ETFs are the thinnest, with support rates of 21 and 27 percent; commodity futures, spot metals, and foreign exchange are supported in about three-quarters of minutes. Thin order books do not simply reflect news: across weekends, larger oracle changes do not come with cheaper-to-move order books (rank correlation -0.07 ; group correlations vary in sign, final column).

6.2 Implied noise is lower where order books are deeper

Implied noise is lower in asset classes whose order books are costlier to move. Internet Appendix Table IA.2 uses the weekend mapping regressions to split the residual variance of the local price change into two parts: the mapping slope times this variance is the implied news component, and the remainder is the implied noise component. Across the seven classes, the rank correlation of implied noise with the median shift cost is -0.86 , and with the support rate -0.46 . Because implied noise also rises with news volatility, and the comparison rests on seven classes, these correlations are descriptive.

Depth alone does not order the mapping slopes, because news intensity also differs across classes: the rank correlation of the slope with the median shift cost is 0.29 , and with the support rate 0.07 . Foreign exchange illustrates prediction P3: the implied standard deviation of weekend news is only 17 basis points, compared with 70 to 237 basis points for the other classes, so noise of 12 basis points yields the lowest slope despite deep displayed order books.

Within asset classes, the relation has the sign prediction P3 implies but is small. Internet Appendix Table IA.10 interacts the local price change with a displayed-liquidity measure

Table 4: Displayed Liquidity during Weekend Benchmark Closures

Market group	Markets	Episodes	Mean support lower bound	Median shift-cost lower bound	Spearman correlation
			(%)	(\$)	ρ
All weekend markets	103	1,846	44.2	10,965	−0.072
U.S. stocks and ADRs	64	1,098	42.1	10,025	0.034
U.S.-listed ETFs	13	138	26.9	5,027	0.154
International equities	11	217	20.8	2,760	−0.042
Equity indices	4	106	48.5	6,692	−0.277
Commodity futures	4	121	76.7	23,143	0.119
Spot metals	4	104	76.8	34,463	−0.176
Foreign exchange	3	62	74.7	54,984	0.311

Notes: The unit of observation is a reconstructed internal-pricing source episode during a weekend benchmark closure. The rule-neutral common sample contains 1,846 episodes across 103 markets for which all three reported measures are available. Each episode receives equal weight; consequently, the all-market row reflects the observed episode composition rather than equal weighting across asset groups.

The support measure is the fraction of all scheduled internal-pricing minutes in which both a \$10,000 purchase and an equal-sized sale have displayed-book VWAPs within 10 basis points of the same-snapshot midpoint. Depth-censored and invalid minutes remain in the denominator and contribute no confirmed support to this conservative lower bound. Table IA.9 separately reports confirmed support, confirmed failure, and unknown support among valid snapshots, together with invalid coverage.

The shift-cost measure is the within-episode fifth percentile of the lower bound on displayed notional required to move the static displayed midpoint by 25 basis points, summarized by its cross-episode median. It is a twenty-level displayed-book bound, not an estimate of manipulation cost.

Spearman ρ is computed within each row between the absolute official post-resumption endpoint change from the last external anchor and $\log(1 + \text{shift-cost lower bound})$. The correlations are descriptive, not causal. Rows with only three or four markets provide particularly limited cross-market inference. Table IA.9 reports the censoring of both measures.

standardized within asset class. Across pooled and U.S.-equity specifications, and for liquidity measured in the same episode or in the contract’s previous weekend, a one-standard-deviation deeper order book is associated with a slope 0.01 to 0.06 higher; none of these estimates is statistically different from zero, and the week-bootstrap intervals exclude increases larger than about 0.13.

6.3 Losing the benchmark thins the order book

National holidays close one benchmark while others remain open. These closures are known in advance but differ across national calendars, so in the same UTC hour some perpetuals lose their benchmark while others do not. For contract j and UTC hour t , we estimate

$$Y_{jt} = \beta \text{HolidayClosure}_{jt} + \alpha_{j \times \text{how}(t)} + \delta_t + \varepsilon_{jt}, \quad (12)$$

where $\text{HolidayClosure}_{jt}$ equals one when a full-day holiday closes the contract’s benchmark in an hour in which it is normally open. Contract-by-hour-of-week fixed effects $\alpha_{j \times \text{how}(t)}$ absorb each contract’s usual intraday and weekly pattern, and exact UTC-hour fixed effects δ_t absorb news and trading conditions common to all contracts, including the commonality in liquidity documented by Chordia, Roll, and Subrahmanyam (2000). The coefficient therefore compares, within the same hour, contracts whose benchmark is closed for a holiday with contracts whose benchmark is open; the sample requires at least two open control calendars. Depth and order-book trading enter as $\log(1 + y)$ in dollars, and local-price updating as the share of minutes in which the local price changes.

Inference uses a calendar-assignment test. Twelve of the 13 holiday events have alternative dates with enough comparable contract-hours; the exception is U.S. Memorial Day. For each of these 12 events, we move the holiday to other dates with the same weekday, UTC offset, regular opening arrangements, and at least two open control calendars, holding the treated contract-hours fixed; events in the same treatment week move together, leaving 11 independent treatment weeks. We re-estimate equation (12) on the same rows in each of 999 draws. The test relies on exchangeability across eligible dates.

A holiday closure reduces top-of-book depth by 52.5 percent, order-book trading by 78.1 percent, and the share of minutes in which the local price changes by 21.5 percentage points (Table 5, Panel A). Each effect has the minimum attainable assignment p -value of 0.001, and excluding dates near announced changes to the oracle’s parameters gives the same signs and p -values (Internet Appendix Table IA.12). The pattern is consistent with liquidity providers withdrawing when they cannot hedge in the closed benchmark: the closure that makes the last external price stale also thins the order book that produces the local price.

The estimates describe eligible full-day holidays on four benchmark calendars, three of them East Asian. Early closes provide no eligible treatment hours, and holidays are rarer than weekends. Internet Appendix Table IA.12 reports all nine outcomes examined.

The effects appear in the holiday week and not before (Figure 3). Stacking weeks -4 through $+4$ around each holiday, with event-by-contract-by-hour-of-week and event-by-exact-hour fixed effects, all three outcomes fall discretely in the event week: depth by about 54 percent, order-book trading by about 76 percent, and local-price updating by 16 percentage points (Panel B). Joint tests do not reject flat pre-event paths (p -values of 0.27 to 0.34).

Table 5: Holiday Closures and Market Quality

	(1) Top-of-book depth	(2) Order-book trading	(3) Local-price updating
<i>Panel A: Average Effect of a Holiday Closure</i>			
Holiday closure	-0.744*** [0.001]	-1.519*** [0.001]	-0.215*** [0.001]
Economic effect	-52.5%	-78.1%	-21.5 pp
Observations / treatment weeks	6,281 / 11	6,281 / 11	6,281 / 11
Holiday closure, all events	-0.712 (0.127)	-1.622 (0.399)	-0.221 (0.057)
Observations	61,823	62,008	62,008
Contracts	96	96	96
Holiday events	13	13	13
Contract $\times$ hour-of-week FE	Yes	Yes	Yes
Exact UTC-hour FE	Yes	Yes	Yes
<i>Panel B: Event-Week Estimates Relative to Week -1</i>			
Event week ($k = 0$)	-0.780*** (0.168)	-1.409** (0.540)	-0.162*** (0.049)
Pretrend test: weeks -4 to -2 (p -value)	0.271	0.289	0.337
Observations	5,874	5,895	5,895
Contracts	67	67	67
Holiday events	11	11	11
Event $\times$ contract-hour FE	Yes	Yes	Yes
Event $\times$ exact-hour FE	Yes	Yes	Yes

Notes: The unit of observation is a contract-UTC-hour. The holiday-closure indicator equals one when a full-day holiday closes the contract's benchmark in an hour in which it is normally open; the perpetual trades throughout. The sample requires at least two benchmark calendars that are open in the same hour.

Panel A estimates $Y_{jt} = \beta \text{HolidayClosure}_{jt} + \alpha_{j \times \text{how}(t)} + \delta_t + \varepsilon_{jt}$. The first row is the primary estimate. Its p -values, in brackets, come from a calendar-assignment test that moves each holiday to other eligible dates with the same weekday and the same treated contract-hours; events in the same treatment week move together. The economic effect converts this estimate. The row for all events uses every holiday, with standard errors clustered by contract and date in parentheses. Table IA.11 lists the events, and Table IA.12 reports all nine outcomes examined.

Depth and order-book trading enter as $\log(1 + y)$, with y in dollars, and the economic effect is $100[\exp(\hat{\beta}) - 1]$ percent. Depth is the hourly average of dollar quantities at the best bid and ask. Local-price updating is the share of minutes in which the local price changes, and its effect is in percentage points. Order-book trading includes voluntary and forced trades on the order book and excludes liquidations executed outside it.


Panel B stacks event weeks -4 through +4 and omits week -1. Standard errors in parentheses are clustered by contract and event. The pretrend row jointly tests weeks -4 to -2. The set of contributing events varies by relative week; Table IA.13 holds it fixed.

In columns (1), (2), the Panel B clustered covariance matrix is adjusted to be positive semidefinite.

*, **, and *** denote $p < 0.10$, $p < 0.05$, and $p < 0.01$; in Panel A they use the assignment test, with 999 draws over 12 eligible events and a minimum attainable p of 0.001. Three of the four benchmark calendars are East Asian.

After the holiday the estimates are imprecise, and for order-book trading and local-price updating they lie above zero. The set of contributing events varies across relative weeks; holding it fixed leaves five events and smaller effects (Internet Appendix Table IA.13), so the calendar-assignment test remains the primary evidence.

Figure 3: Event-Week Dynamics around Holiday Closures



Notes: Points are coefficients from stacked event-week regressions; week -1 is omitted. All models absorb episode-by-market-by-UTC-hour-of-week and episode-by-exact-UTC-hour fixed effects. Bars are 95% intervals two-way clustered by market and episode and are diagnostic under few nominal events. Episode composition varies across weeks. Panels A-B are in log points; Panel C is in percentage points. The main table reports the audited calendar-assignment test; the appendix also holds the event set constant across all nine relative weeks.

Notes: Each panel plots coefficients on indicators for weeks relative to a holiday closure, with week -1 omitted, for contracts whose benchmark the holiday closes relative to contracts whose benchmark is open in the same UTC hour. Regressions stack weeks -4 through $+4$ around each of the 11 contributing events, with event-by-contract-by-hour-of-week and event-by-exact-hour fixed effects. Depth and order-book trading are in log points; local-price updating is the share of minutes in which the local price changes. Bars are 95 percent confidence intervals clustered by contract and event. Table 5, Panel B, reports the event-week estimates and pretrend tests.

Section 7 asks which participants withdraw when the benchmark closes and whose trades carry the signal in local prices.

7 Participation and Signal Production

The weekend signal enters the local price through the trades of a few large accounts, and programmatic market makers supply most of the liquidity against them. Over weekends, the local price moves 2.9 basis points in the direction of takers' trades, and the 1 percent of takers with the largest notional account for 95 percent of this information. When a holiday closes the benchmark, the same accounts become 0.76 percentage points less likely to trade the affected contract than contracts whose benchmarks remain open, about a fifth of their usual 3.9 percent rate.

7.1 Who trades on the weekend signal

Section 5 shows that weekend local prices carry information about where prices settle once trading resumes. To see whose trades bring that information, we split the markout of each voluntary weekend fill into two parts (Table 6). The execution cost is the distance between the fill price and the pre-trade local price. The information is the subsequent move of the local price in the direction of the trade, measured 30 minutes after the benchmark reopens, that is, the part of the trade's price impact that persists to that point. Both are notional-weighted and measured in basis points. We classify accounts by the share of their fills over the preceding 28 days that carry a client order identifier, an optional order label that marks programmatic trading: programmatic accounts have a share of at least 50 percent, manual accounts a share of zero, and mixed accounts lie in between.

Taker orders carry the signal. Across all takers, the local price moves 2.88 basis points in the direction of the trade (week-bootstrap interval 1.10 to 4.71), or \$14.8 million over the sample.

Programmatic and manual takers bring the same information but pay different prices to trade. Both move the price by about 2.8 basis points, but manual takers pay 5.74 basis points to execute, more than three times the 1.73 basis points paid by programmatic takers.

Programmatic market makers supply most of the liquidity and are adversely selected. They account for 76 percent of maker notional, and the price moves 3.01 basis points against their fills (interval -4.62 to -1.43). The 3.82 basis points they earn from the spread more than offset this loss before fees, as in models in which the spread compensates liquidity providers for trading with better-informed counterparties (Glosten and Milgrom, 1985).

Information is concentrated in a few accounts. The 1 percent of taker accounts with the largest weekend notional account for 81 percent of taker notional and 95 percent of taker information, and their trades carry 3.38 basis points of information, compared with 0.76 for the remaining accounts. Because a trader who moves the local price by d moves a rule-based reference price by λd (Section 3), these few accounts also move the reference price, in proportion to the weight the rule places on the local price.

7.2 Who withdraws when the benchmark closes

To see who withdraws when the benchmark closes, we follow accounts around the holiday events of Section 6. For each event, we freeze a cohort of accounts that traded at least two contracts during the preceding 28 days, including at least one whose benchmark the holiday

Table 6: Who Trades on the Weekend Signal

Account type	Accounts	Share of notional (%)	Execution cost (bp)	Information (bp)	95% interval	Information (\$ million)
<i>Panel A: Taker trades</i>						
Programmatic	33,569	50.7	1.73	2.81	[0.58, 5.19]	7.3
Mixed	4,679	2.6	3.70	4.42	[-3.11, 13.35]	0.6
Manual	104,382	43.3	5.74	2.82	[-1.36, 6.63]	6.2
No prior history	90,265	3.4	7.42	3.57	[-16.65, 28.87]	0.6
All	196,034	100.0	3.71	2.88	[1.10, 4.71]	14.8
Programmatic – manual				-0.01	[-4.81, 5.02]	
<i>Panel B: Maker trades</i>						
Programmatic	11,897	76.0	-3.82	-3.01	[-4.62, -1.43]	-11.8
Mixed	2,177	1.6	-5.36	-0.27	[-16.96, 17.70]	0.0
Manual	42,993	20.2	-4.07	2.00	[-2.93, 7.13]	2.1
No prior history	20,451	2.1	-5.08	13.62	[-8.04, 41.17]	1.4
All	64,782	100.0	-3.92	-1.61	[-3.29, 0.19]	-8.3
Programmatic – manual				-5.00	[-10.71, 0.28]	
<i>Panel C: Concentration of taker information</i>						
Taker accounts						196,034
Top 1% of taker accounts by notional: share of taker notional (%)						81.0
Top 1% of taker accounts by notional: share of taker information (%)						95.0
Information of top 1% / remaining 99% of taker accounts (bp)						3.38 / 0.76

Notes: The sample contains voluntary fills in weekend benchmark-closure windows, aggregated by account, market, window, and role. For a fill with sign s (positive for purchases), price p , pre-trade midquote m , and notional q , execution cost is $10^4 s(p-m)/m$ and information is $10^4 s(M-m)/m$, where M is the local midquote 30 minutes after the benchmark reopens; both are notional-weighted averages in basis points. Information measures how far the post-reopening price moves in the direction of the trade, net of the spread; information in dollars is $\sum s(M-m)q/m$. Account types use the share of the account's fills with a client order identifier over the preceding 28 days: programmatic accounts have a share of at least 50 percent, manual accounts a share of zero, and mixed accounts lie in between; accounts without prior history are shown separately. Intervals are percentiles from 999 resamples of the 37 reopening weeks. Taker and maker totals need not offset, because forced fills and fills without valid prices are excluded.

will close and one whose benchmark will stay open. Account–contract–hours without a trade enter as zeros, so the design captures the decision to trade as well as the amount traded. We follow the cohort at the same weekday and UTC hours from four weeks before to four weeks after each event. For account i , contract j , event e , and UTC hour h , we estimate

$$Y_{ijeh} = \beta \text{HolidayClosure}_{jeh} + \alpha_{eij} + \theta_{eih} + \mu_{j \times \text{how}(h)} + \varepsilon_{ijeh}. \quad (13)$$

Event–account–contract fixed effects α_{eij} absorb each account’s usual propensity to trade a given contract, event–account–hour fixed effects θ_{eih} absorb changes in the account’s overall activity at that moment, and contract-by-hour-of-week effects $\mu_{j \times \text{how}(h)}$ absorb seasonality. The coefficient therefore compares, for the same account in the same hour, a contract whose benchmark is closed with contracts whose benchmarks are open. This design needs no alternative dates, so it uses all 13 events, which fall in 12 treatment weeks; p -values come from flipping the signs of the 12 treatment-week score contributions in all $2^{12} = 4,096$ combinations.

A holiday closure makes an account 0.76 percentage points less likely to trade the affected contract, about a fifth of the 3.9 percent rate at which these accounts traded affected contracts the week before (Table 7, Panel A; $p = 0.006$). The probability of providing liquidity that is filled falls by 0.24 percentage points, and order-book trading and maker notional fall as well (p -values of 0.005 to 0.009). The decline is relative: in the raw levels of Panel B, activity also falls in contracts whose benchmarks stay open, but by less, from 2.96 to 2.54 percent of accounts, against 3.92 to 2.07 percent in affected contracts. Internet Appendix Figure IA.2 shows these shares in every week from four weeks before to four weeks after the holiday. Grouping events by treatment month into seven broader blocks leaves the coefficients unchanged and raises the p -values to 0.03–0.09 (Internet Appendix Table IA.14).

Two forces thin the order book behind local prices. Adverse selection is present every weekend: market makers lose to informed takers and must recover the loss from the spread. When the benchmark closes, the additional withdrawal documented here and in Section 6 is consistent with liquidity providers that cannot hedge in the closed benchmark (Ho and Stoll, 1981; Grossman and Miller, 1988). Section 8 asks how a reference-price rule should weigh the local price that this order book produces.

8 Reference-Price Design and Policy Implications

A weight on the local price that rises with episode length improves on the frozen rule by 43 basis points and on the venue’s oracle by 4.2 basis points. The stakes lie mainly in avoiding a stale last external price: when external pricing resumes, the frozen rule would have left \$1.1 billion more revaluation of open positions than the oracle, whereas the selected rule and the oracle differ by less than \$0.1 billion. This section evaluates simple updating rules on the realized price paths, compares them with the frozen rule, full pass-through, and the oracle, and translates their errors into dollars.

Table 7: Account Reallocation around Holiday Closures

	(1) Active <i>indicator</i>	(2) Maker active <i>indicator</i>	(3) Ordinary notional <i>asinh USD/1,000</i>	(4) Maker notional <i>asinh USD/1,000</i>
<i>Panel A: Within-Account Reallocation during Extra Closures</i>				
Holiday closure	−0.0076*** (0.0019)	−0.0024*** (0.0007)	−0.0179*** (0.0046)	−0.0090*** (0.0024)
Block-score sign-test p -value	[0.0059]	[0.0068]	[0.0093]	[0.0049]
Standardized-score p -value	0.0059	0.0049	0.0029	0.0024
Observations	9,943,428	9,943,428	9,943,428	9,943,428
Accounts	43,161	43,161	43,161	43,161
Within-hour treatment contrasts	347,141	347,141	347,141	347,141
Markets	80	80	80	80
Nominal episodes	13	13	13	13
Treatment-week inference blocks	12	12	12	12
Three-way fixed effects	Yes	Yes	Yes	Yes
<i>Panel B: Episode-Equal Raw Participation and Notional</i>				
	Active (%)	Maker active (%)	Ordinary (\$)	Maker (\$)
Week −1, affected markets	3.92	2.14	1,770	875
Week −1, contemporaneously open markets	2.96	1.44	1,874	918
Event week, affected markets	2.07	1.17	550	259
Event week, contemporaneously open markets	2.54	1.29	1,392	636
Week +1, affected markets	2.95	1.76	1,340	671
Week +1, contemporaneously open markets	2.75	1.37	2,379	1,162

Notes: Panels A and B use a cohort frozen from the 28 days before each eligible full-day holiday. An account enters only if it traded at least two study markets and had both an affected market and a contemporaneously open control market in its pre-event portfolio. Eligible account-market-hour cells with no trade are coded as zero.

Panel A estimates $Y_{ijeh} = \beta \text{HolidayClosure}_{jeh} + \alpha_{eij} + \theta_{eih} + \mu_{j \times \text{how}(h)} + \varepsilon_{ijeh}$. Fixed effects absorb episode-account-market propensities, episode-account-exact-hour availability, and market-hour-of-week seasonality. Parentheses are diagnostic episode-clustered standard errors. Scores for episodes sharing a treatment week are aggregated before jointly flipping block signs.

Brackets and stars use all $2^{12} = 4,096$ sign assignments under joint block-score sign symmetry. The standardized test rescales each block by its within-block component dispersion. Table IA.14 reports the coarser treatment-month sensitivity. Within-hour treatment contrasts count distinct episode-account-hours with both treated and control markets on the fitted rows. The regression uses available matched hours in weeks −4 through +4; treatment equals one only in affected markets in week 0. Panel B displays weeks −1, 0, and +1. Asinh outcomes use dollar notional divided by 1,000; coefficients are in transformed units.

Panel B uses the fitted regression rows, first averaging within episode-week-market-role cells and then weighting episodes equally. It is a descriptive levels display. The control benchmarks remain open in the same UTC hour; event composition varies across weeks.

*, **, and *** denote sign-test $p < 0.10$, $p < 0.05$, and $p < 0.01$. The design measures relative account activity across markets in the same hour.

8.1 Rules and evaluation design

For contract j and episode w , let $F_{jw} = O_{jw}(c)$ be the last external price, the oracle’s value when internal pricing begins. We consider the rule family of Section 3,

$$\log P_{\lambda,jw}(u) = \log F_{jw} + \lambda [\log M_{jw}(u) - \log M_{jw}(c)], \quad \lambda \in \{0, 0.25, 0.5, 0.75, 1\}, \quad (14)$$

where $\lambda = 0$ is the frozen rule and $\lambda = 1$ is full pass-through of the local price change. Rules are evaluated one minute before scheduled resumption against the post-resumption oracle, with absolute loss

$$\mathcal{L}_{jw}(P) = 10,000 \left| \log P - \log \overline{O}_{jw}^{+30} \right|. \quad (15)$$

At the start of each calendar month, we select λ using only episodes whose target had been observable for at least 24 hours, drawing first on the same asset class and closure type, then on the same closure type, and finally on all episodes when more specific cells lack sufficient history. The selected weight is then held fixed for the month. For each benchmark rule B_{jw} , the paired loss difference is

$$\Delta \mathcal{L}_{jw} = \mathcal{L}_{jw}(P_{\hat{\lambda},jw}) - \mathcal{L}_{jw}(B_{jw}), \quad (16)$$

so negative values favor the selected rule. We report episode means and means that weight the 32 resumption weeks equally, with week-bootstrap intervals from 1,999 resamples and a one-sided sign randomization test. The primary sample contains 3,749 episodes in 102 contracts.

8.2 The best weight rises with episode length

The selected weights follow prediction P4 (Table 8, Panel D). The selected λ averages 0.10 for maintenance breaks and 0.46 for lunch breaks, against 0.79 to 0.87 for overnight, weekend, and holiday episodes. The realized errors of the two polar rules show the same pattern: the frozen rule is more accurate than full pass-through in maintenance and lunch breaks, and less accurate in the three long closure types; over weekends its mean absolute error is 140 basis points, against 88 for full pass-through. Little news accumulates during short episodes, so local noise dominates; over long episodes, staleness dominates.

8.3 The selected rule improves on the frozen rule and the oracle

The selected rule improves most on the frozen rule. The frozen rule’s mean absolute error is 125 basis points, and the selected rule lowers week-weighted loss by 43 basis points (95 percent interval -53 to -33 ; Table 8, Panel A). It also improves on the venue’s oracle, from a mean absolute error of 89.5 to 83.6 basis points, a week-weighted gain of 4.2 basis points (interval -6.3 to -2.0). The gain over full pass-through is 1.9 basis points; the one-sided sign test favors the selected rule, but the interval $(-3.3$ to $0.2)$ includes zero. The data thus show that local updating beats the frozen rule, and more weakly that shrinking the local change improves on full pass-through. Excluding the seven foreign-exchange and spot-metal contracts, whose benchmark clocks follow TradeXYZ’s published reference hours (Section 2), leaves 3,186 episodes and a gain over the oracle of 4.3 basis points (Panel B).

Table 8: A Policy Application: Out-of-Sample Reference-Price Comparisons before External Inputs Resume

Benchmark	N	Candidate MAE	Benchmark MAE	Episode-mean Δ loss	Week-mean Δ loss	Week-block 95% CI	Sign p -value
<i>Panel A: Primary Comparison at One Minute before Resumption</i>							
Actual within-closure oracle	3,749	83.64	89.51	-5.87	-4.20	[-6.32, -2.00]	< 0.001
Frozen last external price	3,749	83.64	124.81	-41.17	-42.75	[-53.08, -33.39]	< 0.001
Full path-rebased mid	3,749	83.64	86.40	-2.76	-1.86	[-3.32, 0.19]	0.019
<i>Panel B: Sensitivity to Self-Reported Benchmark Clocks</i>							
Sample definition	N	Candidate MAE	Actual oracle MAE	Episode-mean Δ loss	Week-mean Δ loss	Week-block 95% CI	Sign p -value
Full primary sample	3,749	83.64	89.51	-5.87	-4.20	[-6.32, -2.00]	< 0.001
Exclude self-reported-clock markets	3,186	93.84	100.23	-6.39	-4.29	[-6.62, -1.87]	0.001
Price object	N	Spearman ρ	Market-cluster 95% CI		Reopening-week-cluster 95% CI		
<i>Panel C: Displayed-Shift Lower Bounds and Endpoint Errors</i>							
Signed full-mid gain	4,003	-0.039	[-0.091, 0.022]			[-0.087, 0.005]	
Absolute full-mid gain	4,003	-0.365	[-0.473, -0.220]			[-0.421, -0.317]	
Frozen-price endpoint error	4,003	-0.366	[-0.492, -0.199]			[-0.416, -0.318]	
Full-mid endpoint error	4,003	-0.393	[-0.509, -0.250]			[-0.436, -0.354]	
Closure type	N	Mean selected λ	Frozen MAE	Full-mid MAE	Selected-rule MAE	Actual oracle MAE	
<i>Panel D: Selected Weights and Polar-Rule Errors by Closure Type</i>							
Maintenance	766	0.10	21.47	25.35	21.51	21.23	
Lunch break	181	0.46	101.11	119.39	104.59	132.01	
Overnight	1,080	0.79	178.85	118.20	118.19	126.15	
Weekend	1,496	0.85	140.31	88.12	85.47	89.71	
Holiday	226	0.87	133.19	103.46	100.24	110.44	

Notes: Panel A evaluates a rolling, strictly out-of-sample candidate at the last observed minute before scheduled external-input resumption. At each calendar-month origin, the candidate selects $\lambda \in \{0, 0.25, 0.5, 0.75, 1\}$ using only endpoint targets available at least 24 hours earlier and sets $\log P_\lambda = \log F + \lambda \log(\text{mid}_u/\text{mid}_c)$. The endpoint is the mean official oracle over the first 30 minutes after resumption, not an independently observed true value.

MAE and paired loss differences are in basis points; the reported MAEs are episode averages. Δ loss equals candidate absolute error minus benchmark absolute error, so negative values favor the candidate. The week statistic first averages paired losses within reopening week and then weights weeks equally. Confidence intervals resample the 32 reopening weeks 1,999 times. The reported sign-test p -value is one-sided for a lower candidate loss and uses 19,999 sign randomizations. No significance stars are used because the full-mid row's one-sided sign test and two-sided week-block interval give different boundary assessments.

Panel B repeats the actual-oracle comparison, holding the selected rules fixed, after excluding the 7 markets whose benchmark clocks follow TradeXYZ's published OTC reference hours (xyz:EUR, xyz:GBP, xyz:GOLD, xyz:JPY, xyz:PALLADIUM, xyz:PLATINUM, xyz:SILVER). The full sample is repeated to make the sample change explicit.

Panel C uses the 4,003-episode censoring-aware lower-bound sample and the fixed full path-rebased midquote P_1 ($\lambda = 1$), not the selected rule in Panels A–B. Full-mid gain is $G = \mathcal{L}(F) - \mathcal{L}(P_1)$ in basis points; the absolute-gain row uses $|G|$. These gains are functions of the two component errors. Correlations are with the log lower bound on displayed notional required to move the book-implied mid by 25 bp. Confidence intervals separately resample markets and reopening weeks 999 times. The panel is contemporaneous and descriptive; the displayed-book bound abstracts from replenishment, strategic response, and impact persistence.

Panel D partitions the Panel A sample by closure type. It reports the mean weight selected from past outcomes and the episode-average absolute errors, in basis points, of the frozen anchor ($\lambda = 0$), the full path-rebased midquote ($\lambda = 1$), the selected rule, and the actual within-closure oracle.

The candidate in Panels A–B is unconstrained by the oracle's complete historical Discovery Bound and reset rules. Historical bound methodology is not reconstructed for 3,600 of 3,749 primary observations (96.03%). The remaining records carry CL methodology-period classifications: 120 v2, 28 archived v1, and 1 unverified transition. These classifications identify rule periods, not complete historical bound and reset paths. Table IA.16 imposes the documented bounds on the candidate as a scenario. The table establishes an unconstrained predictive comparison on realized paths; implementability, equilibrium response, and welfare remain separate questions. Resumption times are reconstructed from official rules and schedules.

The gain over the oracle survives the venue’s price bounds. The oracle may not move beyond a documented bound around the last external price, whereas the rules above are unconstrained. Internet Appendix Table IA.16 imposes each contract’s documented bound on every candidate, in training as well as evaluation, and re-selects the monthly weight. The bounds rarely bind: the selected candidate is projected onto a bound in 0.48 percent of episodes, and the oracle itself lies outside the documented bound in 0.61 percent. The gain over the oracle becomes 3.5 basis points (interval -5.6 to -1.6), or 3.7 when the documented reset budget is allowed. It remains with bounds tightened to 75 percent of their documented width (2.9 basis points) and disappears at 50 percent (1.2 basis points, interval -4.3 to 2.0).

Episodes whose order books are cheaper to move are harder to price by either polar rule (Panel C). On 4,003 episodes with shift-cost bounds, the rank correlation of the shift cost with the error is -0.37 for the frozen rule and -0.39 for full pass-through, whereas the advantage of full pass-through is unrelated to the shift cost (-0.04 , interval including zero).

8.4 The stakes at resumption

Resumption is usually orderly but has a long tail. Across 4,640 scheduled resumptions, the mark price typically converges to the oracle within two minutes, but it has not converged after two hours in 12 percent of resumptions, and forced order-book trading follows 21 percent (Internet Appendix Table IA.15).

Basis-point errors translate into dollars through open interest. When external pricing resumes, the oracle catches up with the external price, and every open position is revalued by the rule’s remaining error. Table 9 multiplies each rule’s error one minute before resumption by open interest in dollars. Over the 260-day sample, the oracle leaves \$1.52 billion of revaluation at resumption. The frozen rule would have left \$2.60 billion, or \$1.08 billion more (95 percent interval \$0.68 to \$1.53 billion). The selected rule would have left \$86 million less than the oracle, an improvement whose interval ($-\$170$ to $+\$3$ million) narrowly includes zero.

Large revaluations are followed by forced trading. Sorting episodes by the oracle’s error (Panel B), forced order-book trading, backstop liquidation, or auto-deleveraging⁴ follows 6 percent of resumptions in the lowest decile and 73 percent in the highest, where forced trading reaches 70 basis points of open interest. The top decile alone accounts for \$116 million of the \$178 million of forced trading after resumptions in the sample. Applying these decile rates to each rule’s error, the frozen rule is associated with \$145 million more forced trading than the oracle (interval \$72 to \$295 million), and the selected rule with \$7 million less (interval $-\$35$ to $+\$16$ million). These are associations on realized paths, and they locate the stakes in avoiding a stale last external price.

The dollar size of forced trading rests on a few large resumptions, but the ranking of the rules does not. Panel C of Table 9 repeats the comparisons with a contract fixed-effects slope, which gives \$240 million for the frozen rule in the full sample. Dropping the 38 episodes in the top 1 percent by forced trading, which account for \$130 million of the \$178 million, lowers this figure to \$17 million (interval \$11 to \$26 million). The revaluation differences

⁴Backstop liquidation transfers distressed positions to a dedicated liquidator when order-book liquidation is insufficient. Auto-deleveraging reduces counterparties’ positions to resolve an insolvent account.

Table 9: The Stakes of Reference-Price Errors at Source Resumption

Rule	Revaluation at resumption		Forced trading after resumption			
	Total (\$m)	Difference from actual oracle (\$m)	Decile-rate (\$m)	Difference (\$m)	Market-FE slope (\$m)	Difference (\$m)
<i>Panel A: Revaluation and forced trading implied by each rule</i>						
Frozen last external price	2,596.4	1,080.2 [684.3, 1,529.9]	322.9	145.3 [71.6, 294.8]	418.1	240.4 [62.1, 520.4]
Actual within-closure oracle	1,516.2	—	177.6	—	177.6	—
Full path-rebased mid	1,441.7	-74.5 [-157.8, 6.9]	169.5	-8.1 [-33.0, 11.8]	161.0	-16.6 [-46.0, 1.9]
Selected rule	1,430.4	-85.8 [-170.3, 2.6]	170.6	-7.0 [-34.7, 16.0]	158.5	-19.1 [-51.2, 0.5]
<i>Panel B: Forced trading by decile of the implemented oracle’s error</i>						
Decile	Median error (bp)	Episodes	Share with forced trading (%)	Forced trading (\$m)	Forced trading (bp of open interest)	
1	2.2	375	6.1	1.9	0.45	
2	7.0	375	7.5	1.2	0.28	
3	13.6	375	5.9	1.2	0.38	
4	21.9	375	10.1	1.7	0.66	
5	33.3	374	12.3	8.0	3.54	
6	50.1	375	16.0	6.6	3.53	
7	73.6	375	27.2	7.5	4.34	
8	109.5	375	30.9	8.4	5.19	
9	171.8	375	44.0	25.5	14.11	
10	315.3	375	72.5	115.6	70.05	
<i>Panel C: Differences from the actual oracle after excluding large events (\$m)</i>						
Sample	Episodes	Forced trading (\$m)	Frozen last external price		Selected rule	
			Revaluation	Forced trading (market FE)	Revaluation	Forced trading (market FE)
Full sample	3,749	177.6	1,080.2 [684.3, 1,529.9]	240.4 [62.1, 520.4]	-85.8 [-170.3, 2.6]	-19.1 [-51.2, 0.5]
Excl. WTI crude, 8 Mar 2026	3,748	145.5	1,072.1 [670.4, 1,517.9]	117.7 [50.2, 205.7]	-87.6 [-171.6, 0.8]	-9.6 [-23.6, 0.1]
Excl. largest reopening week	3,676	133.4	1,063.4 [676.2, 1,527.8]	115.6 [48.6, 203.9]	-89.0 [-167.2, -4.5]	-9.7 [-22.0, -0.5]
Excl. top 1% by forced trading	3,711	47.4	1,050.6 [657.7, 1,489.7]	17.0 [10.6, 25.6]	-71.9 [-139.4, -5.7]	-1.2 [-2.7, -0.1]
Excl. SK Hynix markets	3,497	150.9	478.0 [302.3, 686.8]	280.9 [60.2, 564.6]	-58.8 [-141.8, 31.1]	-34.6 [-95.6, 16.7]

Notes: The sample is the 3,749 primary episodes of Table 8 in 102 markets, evaluated one minute before scheduled resumption. Open interest is converted to dollars at the oracle price in that minute (mean \$66.9 million). A rule's revaluation at resumption is open interest times the rule's absolute error against the mean official oracle over the first 30 minutes after resumption; it is the mark-to-market transfer between long and short positions that the rule leaves to occur when external inputs return. Forced trading is the absolute net forced order-book flow plus backstop and auto-deleveraging notional in the 30 minutes after resumption (\$177.6 million in total). The decile-rate column assigns each episode the ratio of forced trading to open interest observed in its decile of the implemented oracle's error (Panel B) and applies it to the error of each rule; the market-FE column uses the slope of forced trading on the implemented oracle's revaluation with market fixed effects (0.223). Both are associations on realized paths. Brackets are 95% intervals from 999 resamples of the 32 reopening weeks, with deciles and slopes re-estimated in each draw. Totals cover the 260-day sample. Panel C repeats the Panel A differences after removing, in turn, the episode with the most forced trading (WTI crude oil, reopening on 8 March 2026; Figure 4), the reopening week with the most forced trading, the top 1% of episodes by forced trading, and every episode in the two SK Hynix markets (xyz:SKHX and xyz:SKHY); each exclusion re-estimates the market-FE slope and the bootstrap. The 28 July 2026 SK Hynix event in Figure 4 lies in an unverified source interval and is not in any row.

change little under these exclusions, except that dropping the two SK Hynix contracts, which carry 18 percent of open interest, lowers the frozen rule’s excess revaluation to \$478 million. In every subsample the frozen rule is worse than the oracle on both measures, with intervals that exclude zero, and the selected rule’s point estimates are below the oracle’s.

8.5 Two failures at resumption

Large errors at resumption can come from the local price or from the external price itself. Figure 4 shows one of each.


The weekend of March 6–8, 2026, in WTI crude oil is the resumption with the most forced trading in the evaluation sample (Panel A). While CME was closed, the local price rose 5.0 percent from the last external price, and the oracle followed it. The post-resumption oracle was 14.5 percent above the last external price, so the local price captured about a third of the move. In the 30 minutes after resumption, \$28.1 million of net forced order-book flow and \$4.1 million of auto-deleveraging followed, and open interest fell from \$168 million to \$147 million. The error here is the local price’s: it moved in the right direction, but by too little.

SK Hynix on July 28, 2026, shows the opposite failure (Panel B). Until late July, the venue’s documented external coverage for Korean equities began with the pre-market session at 08:00 Korea time. At that minute, as external pricing returned, the oracle fell 23 percent within one minute and was back near its pre-resumption level by the seventh minute; its 30-minute mean differed from the last pre-resumption value by 0.1 percent. The order book carried no signal of a move of this size. Within four minutes, \$84.7 million of forced order-book flow, backstop liquidation, and auto-deleveraging were executed, and open interest fell from \$483 million to \$314 million, almost half of the \$178 million of forced trading after all 3,749 resumptions in Table 9. A parameter notice dated the same day moved the start of pre-market coverage from 08:00 to 08:10. Because the notice gives no effective time, the documented schedule does not determine which window applied on July 28, so the episode lies in an unverified source state and is excluded from Tables 8 and 9. The error here is the external price’s.

8.6 Policy implications

The comparison measures forecast accuracy on realized paths, holding quotes and trading fixed; a different rule could change how participants quote and trade. Within this scope, the evidence has three implications for venues and regulators. First, the relevant state for reference-price rules is whether the oracle is in external or internal pricing, not the trading hours of the benchmark; the two differ in most benchmark-closed minutes of U.S. equities. Second, a universal frozen rule is costly in long episodes, whereas full reliance on local prices imports noise that is larger where order books are thinner, and the order book thins further when the benchmark itself is closed. The appropriate weight is therefore state- and product-specific, rising with episode length and news intensity and falling with local noise. Third, because the right weight depends on these states, rules should be auditable. Venues should record and disclose when external pricing stops and resumes, which inputs form the external price, and the history of price bounds and resets, so that regulators can evaluate rules against

Figure 4: Two Resumptions with Large Forced Trading



Notes: Panel A: WTI crude oil (xyz:CL), March 6, 20:00 UTC, to March 8, 23:00 UTC, 2026; the shaded area is the CME closure, and the dashed line is the scheduled resumption at 22:00 UTC. Panel B: SK Hynix (xyz:SKHX), July 27, 22:30 UTC, to July 27, 23:40 UTC, 2026; triangles mark the lowest oracle and mark prices observed within a minute when they lie more than one dollar below the minute's closing value. Top row: local price, oracle, and mark price at one-minute frequency. Middle row: net forced order-book flow (absolute value of net forced taker trading), backstop liquidation, and auto-deleveraging notional per minute. Bottom row: open interest in dollars at the oracle price. Middle and bottom rows cover 10 minutes before to 30 minutes after resumption. The Panel A episode is in the evaluation sample of Table 9; the Panel B episode lies in an unverified source state and is not.

their actual constraints and distinguish failures of the external price from failures of the local price that replaces it.

9 Conclusion

Derivatives that trade around the clock need a reference price when their benchmark is closed. We study this problem in 103 perpetual futures contracts on traditional assets. Loss of the external price is a distinct event from a benchmark closure: 47 percent of benchmark-closed minutes are in internal pricing, while the rest remain in external pricing. During internal pricing, the venue must decide how much weight to place on the local price.

A simple signal-extraction framework frames that decision. The optimal weight on local prices is the signal share of the local price change, rising with accumulated news and falling with local noise. The evidence is consistent with this structure. Weekend local price changes map into the post-resumption oracle with slopes of 0.68 to 0.93 and lower out-of-sample forecast errors by 17 to 73 percent in the four asset classes with enough history. The signal strengthens as the episode proceeds, and the slopes below one, post-resumption reversals, and lower slopes at benchmark reopening are traces of local noise. That noise is lower in asset classes whose order books are costlier to move, although displayed liquidity does not measurably change the weight within classes.

The order book that supplies the signal is thin. A \$10,000 trade can be verified to execute within 10 basis points on both sides in only 44 percent of weekend minutes, and a holiday closure reduces top-of-book depth by about half and order-book trading by about 78 percent, as the same accounts trade less in affected contracts than in contracts whose benchmarks remain open. The signal enters the local price through the aggressive trades of a small number of large accounts, and programmatic market makers absorb it at an informational loss covered by the spread. The closure that makes the last external price stale thus also thins the order book that produces the local price. When external pricing resumes, the mark price usually converges to the oracle within minutes, with a tail of slower adjustment.

Weights selected from past outcomes are near 0.1 for brief maintenance breaks and near 0.85 for weekends and holidays, and the frozen rule beats full pass-through only in short episodes. The selected rule improves on the oracle by 4.2 basis points, or 3.5 basis points under the venue’s documented price bounds, a predictive comparison that holds behavior fixed. The stakes lie mainly in avoiding a stale last external price: the frozen rule would have left \$1.1 billion more revaluation at resumption, and forced trading after resumption is concentrated in episodes with large reference-price errors. Those errors can come from a local price that moved too little, as in the WTI crude oil weekend of March 2026, or from the external price itself, as in the SK Hynix resumption of July 2026.

For regulators considering 24/7 trading and listed perpetuals, the evidence favors rules for the reference price that depend on the observable state: whether external pricing is on, how long it has been off, and how deep the order book is. Such rules require venues to record and disclose source states and price constraints. Extending the evidence to other venues, longer samples, and the equilibrium effects of alternative rules remains for future work.

References

- Rosa M. Abrantes-Metz, Michael Kraten, Albert D. Metz, and Gim S. Seow. LIBOR manipulation? *Journal of Banking and Finance*, 36(1):136–150, 2012.
- Damien Ackerer, Julien Hugonnier, and Urban Jermann. Perpetual futures pricing. *Mathematical Finance*, 36(3):481–499, 2026.
- Carol Alexander, Jaehyuk Choi, Heungju Park, and Sungbin Sohn. BitMEX Bitcoin derivatives: Price discovery, informational efficiency, and hedging effectiveness. *Journal of Futures Markets*, 40(1):23–43, 2020.
- Guillermo Angeris, Tarun Chitra, Alex Evans, and Matthew Lorig. A primer on perpetuals. *SIAM Journal on Financial Mathematics*, 14(1):SC17–SC30, 2023.
- Michael J. Barclay and Terrence Hendershott. Price discovery and trading after hours. *Review of Financial Studies*, 16(4):1041–1073, 2003.
- Michael J. Barclay and Terrence Hendershott. Liquidity externalities and adverse selection: Evidence from trading after hours. *Journal of Finance*, 59(2):681–710, 2004.
- Michael J. Barclay and Terrence Hendershott. A comparison of trading and non-trading mechanisms for price discovery. *Journal of Empirical Finance*, 15(5):839–849, 2008.
- Bruno Biais, Pierre Hillion, and Chester Spatt. Price discovery and learning during the preopening period in the Paris Bourse. *Journal of Political Economy*, 107(6):1218–1248, 1999.
- Nina Boyarchenko, Lars C. Larsen, and Paul Whelan. The overnight drift. *Review of Financial Studies*, 36(9):3502–3547, 2023.
- William A. Brock and Allan W. Kleidon. Periodic market closure and trading volume: A model of intraday bids and asks. *Journal of Economic Dynamics and Control*, 16(3–4):451–489, 1992.
- James Brugler, Marta Khomyn, and Tālis J. Putniņš. Benchmarking benchmarks. *Journal of Financial Economics*, 168:104018, 2025.
- Charles Cao, Eric Ghysels, and Frank Hatheway. Price discovery without trading: Evidence from the Nasdaq preopening. *Journal of Finance*, 55(3):1339–1365, 2000.
- Tarun Chordia, Richard Roll, and Avanidhar Subrahmanyam. Commonality in liquidity. *Journal of Financial Economics*, 56(1):3–28, 2000.
- Todd E. Clark and Kenneth D. West. Approximately normal tests for equal predictive accuracy in nested models. *Journal of Econometrics*, 138(1):291–311, 2007.
- Commodity Futures Trading Commission. CFTC staff seek public comment on 24/7 trading. Press Release No. 9068-25, April 21. <https://www.cftc.gov/PressRoom/PressReleases/9068-25>, 2025a.

- Commodity Futures Trading Commission. CFTC staff seek public comment regarding perpetual contracts in derivatives markets. Press Release No. 9069-25, April 21. <https://www.cftc.gov/PressRoom/PressReleases/9069-25>, 2025b.
- Commodity Futures Trading Commission. Policy statement concerning the listing of perpetual contracts. Federal Register 91, 33160, June 3, 2026a.
- Commodity Futures Trading Commission. Request for comment on the extension of standard futures contracts to 24/7 trading and on perpetual contracts referencing physically delivered or storable energy commodities. Press Release No. 9259-26, June 22. <https://www.cftc.gov/PressRoom/PressReleases/9259-26>, 2026b.
- Commodity Futures Trading Commission. CFTC to stay self-certified contract on 24/7 trading for crude oil futures. Press Release No. 9265-26, July 9. <https://www.cftc.gov/PressRoom/PressReleases/9265-26>, 2026c.
- Lin William Cong, Liam Fox, Siguang Li, and Luofeng Zhou. A primer on oracle economics. *Journal of Corporate Finance*, 94:102800, 2025.
- Shane A. Corwin and Marc L. Lipson. Order flow and liquidity around NYSE trading halts. *Journal of Finance*, 55(4):1771–1801, 2000.
- Darrell Duffie and Piotr Dworczak. Robust benchmark design. *Journal of Financial Economics*, 142(2):775–802, 2021.
- Darrell Duffie and Jeremy C. Stein. Reforming LIBOR and other financial market benchmarks. *Journal of Economic Perspectives*, 29(2):191–212, 2015.
- Financial Stability Board. Reforming major interest rate benchmarks. Report, July 22. https://www.fsb.org/2014/07/r_140722/, 2014.
- Kenneth R. French and Richard Roll. Stock return variances: The arrival of information and the reaction of traders. *Journal of Financial Economics*, 17(1):5–26, 1986.
- Mason S. Gerety and J. Harold Mulherin. Trading halts and market activity: An analysis of volume at the open and the close. *Journal of Finance*, 47(5):1765–1784, 1992.
- Lawrence R. Glosten and Paul R. Milgrom. Bid, ask and transaction prices in a specialist market with heterogeneously informed traders. *Journal of Financial Economics*, 14(1):71–100, 1985.
- John M. Griffin and Amin Shams. Manipulation in the VIX? *Review of Financial Studies*, 31(4):1377–1417, 2018.
- Sanford J. Grossman and Merton H. Miller. Liquidity and market structure. *Journal of Finance*, 43(3):617–633, 1988.
- Joel Hasbrouck. Measuring the information content of stock trades. *Journal of Finance*, 46(1):179–207, 1991.

- Songrun He, Asaf Manela, Omri Ross, and Victor von Wachter. Fundamentals of perpetual futures. Working paper, SSRN 4301150, 2024.
- Yan He, Hai Lin, Jing Wang, and Chunchi Wu. Price discovery in the round-the-clock U.S. Treasury market. *Journal of Financial Intermediation*, 18(3):464–490, 2009.
- Terrence Hendershott, Dmitry Livdan, and Dominik Rösch. Asset pricing: A tale of night and day. *Journal of Financial Economics*, 138(3):635–662, 2020.
- Thomas Ho and Hans R. Stoll. Optimal dealer pricing under transactions and return uncertainty. *Journal of Financial Economics*, 9(1):47–73, 1981.
- International Organization of Securities Commissions. Principles for financial benchmarks. Final Report FR07/13, July. <https://www.iosco.org/library/pubdocs/pdf/IOSCO415.pdf>, 2013.
- Praveen Kumar and Duane J. Seppi. Futures manipulation with “cash settlement”. *Journal of Finance*, 47(4):1485–1502, 1992.
- Charles M. C. Lee, Mark J. Ready, and Paul J. Seguin. Volume, volatility, and New York Stock Exchange trading halts. *Journal of Finance*, 49(1):183–214, 1994.
- Andrew W. Lo and A. Craig MacKinlay. Stock market prices do not follow random walks: Evidence from a simple specification test. *Review of Financial Studies*, 1(1):41–66, 1988.
- Dong Lou, Christopher Polk, and Spyros Skouras. A tug of war: Overnight versus intraday expected returns. *Journal of Financial Economics*, 134(1):192–213, 2019.
- Mathis Mourey, Mohamad H. Shahrour, and Florentina Şoiman. A crypto-stock weekend effect: Predicting Monday stock returns using weekend cryptocurrency returns. *Finance Research Letters*, 86:108661, 2025.
- Richard Roll. A simple implicit measure of the effective bid-ask spread in an efficient market. *Journal of Finance*, 39(4):1127–1139, 1984.
- Donald J. Schuirmann. A comparison of the two one-sided tests procedure and the power approach for assessing the equivalence of average bioavailability. *Journal of Pharmacokinetics and Biopharmaceutics*, 15(6):657–680, 1987.
- Robert J. Shiller. Measuring asset values for cash settlement in derivative markets: Hedonic repeated measures indices and perpetual futures. *Journal of Finance*, 48(3):911–931, 1993.

A Internet Appendix

The Internet Appendix follows the order of the paper. Section A.1 documents the price objects of Section 2; Section A.2 derives and calibrates the framework of Section 3; Sections A.4–A.8 report the samples, robustness checks, and additional results for Sections 5–8.

A.1 Institutional Details

Table IA.1 defines the four price objects used in the paper and summarizes how each updates in the four benchmark and source states.

The mark price combines the oracle with order-book information through a separate smoothing mechanism. Let P_t^{mid} denote the local price and define the local basis relative to the oracle as $b_t = P_t^{mid} - S_t$. One documented mark-price component is

$$C_t^{basis} = S_t + \text{EMA}_{150s}(b_t), \quad (17)$$

which smooths the local premium or discount with a 150-second time constant. At the Hyperliquid infrastructure layer, the local mark component is

$$L_t = \text{median}(P_t^{bid}, P_t^{ask}, P_t^{last}), \quad (18)$$

and the reported mark price is the robust median of L_t and the available deployer-supplied mark-price components, including C_t^{basis} . Its adjustment therefore depends on the available components and applicable price bounds, as well as on basis smoothing. The oracle is the reference for funding, whereas the mark price is used for margin, liquidation, unrealized profit and loss, and related risk controls. The empirical analysis uses the mark price reported by the protocol.

The main text omits two details of the oracle rule. For equity contracts, an April 30, 2026 announcement shortened the smoothing time constant in equation (1) from one hour to 30 minutes, which speeds the transmission of local prices into the oracle. The external price that returns at resumption comes from the venue’s external data sources for U.S. equities, from the resumed futures session for commodity futures, and from the reopened local session after a lunch break for Asian equities.

Table IA.1: Price Objects and Their Updating Rules

Panel A: The Four Price Objects

Price object	General definition	Economic role
Local price	Midquote of the best bid and offer in the venue's order book.	Measures the center of locally executable prices and traces the local price path.
External Price field	Relayer field carrying the external price.	Records the external price during external pricing and the last external price during internal pricing.
Oracle	Contract reference price published by the Relayer.	Enters contract-level reference-price mechanisms, including funding and price constraints.
Mark price	Risk-management price formed as a robust median of oracle-based and local order-book components.	Used for margin, liquidation, and account-level risk measurement.

Panel B: Price Formation across the Four States

State	Benchmark	Source	Local price	External Price field	Oracle	Mark price
Open / external pricing	Open	External	Live order-book midquote.	Updates with external fair price.	Equals external fair price.	Robust median including an external-oracle-plus-basis component and the local mark component.
Closed / external pricing	Closed	External	Live order-book midquote.	Continues to update externally.	Remains externally derived.	Robust median including an external-oracle-plus-basis component and the local mark component.
Open / internal pricing	Open	Internal	Live order-book midquote.	Fixed at last external value.	Smoothed impact-price updates with a dead zone.	Robust median including an internal-oracle-plus-basis component and the local mark component.
Closed / internal pricing	Closed	Internal	Live order-book midquote.	Fixed at last external value.	Smoothed impact-price updates with a dead zone.	Robust median including an internal-oracle-plus-basis component and the local mark component.

Notes: Panel A defines the four price objects; Panel B crosses benchmark trading status with the scheduled oracle source, as in Table 1. The local price requires a valid two-sided order book. Impact prices are average execution prices for a configured notional on each side of the order book. The open/internal-pricing state arises at contract-specific clock boundaries. Contract-level price bounds (Discovery Bounds) limit how far the oracle and mark price can move from the last external price and do not change the source state. Source states and transition times are reconstructed from documented schedules.

A.2 Framework Derivations and Calibration

This section derives the reversal slope in prediction P1 of Section 3 and the difference between the source-resumption and benchmark-reopening slopes, and calibrates the framework by asset class, using the notation of equations (2)–(4).

Reversal after resumption. Suppose the reversal after resumption is measured as $m_{o+h} - m_o$ and regressed on a local change that ends k minutes before resumption, as in Table IA.8. News increments over the two non-overlapping intervals are independent, so only the noise terms covary. Abstracting from controls and from any change in the noise process at resumption, and for closures long enough that $\gamma(\tau_o) \approx \gamma(\tau_o + h) \approx 0$, the population slope is

$$\frac{\text{Cov}(m_{o+h} - m_o, m_{o-k} - m_c)}{\text{Var}(m_{o-k} - m_c)} \approx -\frac{\gamma(k) - \gamma(h+k)}{\sigma^2\tau_{o-k} + \psi_{o-k}^2}. \quad (19)$$

The slope is negative when the autocovariance of local deviations declines between lags k and $h+k$, and it is zero when deviations are serially uncorrelated at lag k , as with bid–ask bounce. A negative reversal slope measured after a k -minute gap therefore indicates local deviations that persist beyond k minutes. If resumption also compresses local deviations, the reversal is larger.

Two endpoints. In the baseline framework, the source-resumption slope equals $\lambda^*(u)$. Let the benchmark-reopening target be $Y_b = v_b - v_c + \zeta_b$ for $b \geq o$. Because $v_b - v_o$ is independent of the local path through $u \leq o$, $\text{Cov}(Y_b, X(u)) = \sigma^2\tau_u$, and the slope is unchanged: later news lowers the fit but not the slope. Now relax the baseline and suppose the resumed external source reports $v_o + \xi_o$, where ξ_o is a transitory component that the reopened benchmark removes. Then

$$\beta_b(u) - \beta_o(u) = -\frac{\text{Cov}(\xi_o, \eta_u - \eta_c)}{\sigma^2\tau_u + \psi_u^2}. \quad (20)$$

A negative difference indicates that the source-resumption endpoint shares a transitory component with the local path that is undone when the benchmark reopens. Panel A of Table IA.5 estimates this difference directly.

Calibration. In the framework, the population slope of the endpoint regression is the share of residual local-price variation that is news. Table IA.2 uses this identity to split the residual variance of the weekend local price change into an implied news component and an implied noise component for each asset class, and compares the noise component with the displayed-liquidity measures of Table 4, computed on the same episodes.

Table IA.2: Implied News and Noise in Weekend Local Prices

Asset class	N	β	Implied decomposition				Displayed book	
			News s.d. (bp)	Noise s.d. (bp)	95% interval noise s.d.	Signal-to- noise ratio	Support (%)	Median shift cost (\$)
Equity indices	65	0.934	70.2	18.7	[0.0, 43.8]	14.1	78.3	115,603
U.S. equities and ADRs	715	0.832	155.4	69.9	[42.8, 83.4]	4.9	59.1	17,958
Spot metals	90	0.813	82.1	39.4	[0.0, 63.7]	4.4	81.6	39,364
U.S. ETFs	57	0.807	174.2	85.3	[0.0, 115.4]	4.2	55.6	14,058
Commodity futures	103	0.792	237.4	121.6	[38.4, 186.1]	3.8	81.1	24,340
International equities	40	0.735	222.7	133.9	[0.0, 172.3]	2.8	3.1	468
Foreign exchange	44	0.675	17.3	12.0	[3.1, 14.9]	2.1	79.8	59,259

Notes: The sample and specification are those of Table 3, Panel A: weekend internal-pricing episodes, the local price change one minute before scheduled resumption, the mean oracle over the first 30 minutes after resumption, contract fixed effects, and predetermined controls, including the initial basis between the local price and the oracle. Let $\tilde{x}$ and $\tilde{y}$ denote the local change and the target after removing contract means and projecting out the controls. Under the framework in Section 3, the news variance is $\widehat{\text{Cov}}(\tilde{x}, \tilde{y}) = \hat{\beta} \widehat{\text{Var}}(\tilde{x})$ and the noise variance is $(1 - \hat{\beta}) \widehat{\text{Var}}(\tilde{x})$, both with residual degrees of freedom; the columns report their square roots in basis points over the internal-pricing interval. Because the controls include the initial basis, the noise term corresponds to $\text{Var}(\eta_u \mid \eta_c)$, approximately the variance of the local price's deviation from value at the evaluation minute. The signal-to-noise ratio equals $\hat{\beta}/(1 - \hat{\beta})$. Intervals are 2.5th and 97.5th percentiles from 999 resamples of whole resumption weeks, with contract effects and controls re-estimated in each draw. Displayed-book measures use the definitions of Table 4, computed on the calibration episodes: the fraction of scheduled minutes with confirmed two-sided support for a \$10,000 trade within 10 basis points of the local price, and the median across episodes of the within-episode fifth percentile of the lower bound on displayed notional required to move the local price by 25 basis points. Rank correlations across seven classes are descriptive.

Spearman correlations across the seven classes: noise s.d. with support -0.46 and with shift cost -0.86 ; signal-to-noise ratio with support 0.07 and with shift cost 0.29 .

A.3 Data Construction

Contract coverage follows the venue’s documented listing information where it is available and the observed trading record otherwise; when a listing date is unavailable, coverage begins on the first observed trading date. Each contract keeps one identifier through listing changes, so its prices, trades, and institutional attributes can be followed throughout the sample. Benchmark schedules are taken from published trading hours and exchange holiday calendars in each benchmark’s local time zone.

Table 2 reports scheduled closure duration, the time from the last pre-closure boundary to scheduled benchmark reopening. The contract–closure-window panel also records observed duration, the minutes for which the trading record is available, and the local price at the quarter, half, three-quarter, and final pre-reopening points of each window.

The minute counts differ across exhibits because they classify minutes differently. The 15,226,677 contract–minutes assigned to complete benchmark-closure windows include 9,950 reopening minutes whose minute-end benchmark state is open. Figure 1 instead classifies all 15,379,020 contract–minutes whose minute-end benchmark state is closed, including 162,293 minutes outside complete windows. The closure calendar contains 15,291,300 scheduled minutes; the 64,623-minute difference from the complete-window count reflects minutes without an observed record and the treatment of window boundaries.

A.4 Information in Local Prices

Table IA.3 traces the weekend sample from eligible internal-pricing episodes to the paired and out-of-sample samples of Table 3, and Table IA.4 compares retained with excluded episodes. The sample also excludes episodes whose transition between external and internal pricing cannot be verified or whose benchmark and oracle schedules disagree. Table IA.6 reports the forecasting specification that pools contracts within asset class. Table IA.7 compares local price dynamics across source states, and Table IA.8 reports the post-resumption reversals discussed under prediction P1.

For equity indices, international equities, U.S. equities, and U.S. ETFs, the benchmark reopens after source resumption. Panel A of Table IA.5 moves the target to the oracle after benchmark reopening while keeping the same episodes and predictor. The slopes fall to 0.605 for equity indices, 0.633 for international equities, 0.552 for U.S. equities, and 0.254 for U.S. ETFs; the ETF estimate is imprecise. The same panel tests the difference between the two slopes on the same observations and bootstrap draws. The differences are negative in all four classes: -0.329 for equity indices (wild $p = 0.007$), with p -values of 0.066, 0.106, and 0.334 for international equities, U.S. equities, and ETFs. In the framework, news arriving after resumption lowers the fit at the later endpoint but not the slope (equation (20)). A lower slope instead indicates a transitory component shared by local prices and the resumed external price that the reopened benchmark removes.

Panel B of Table IA.5 reports forecast gains against the later target: 59.0 percent for commodity futures, 18.2 percent for equity indices, 5.9 percent for international equities, and 66.1 percent for spot metals. The smaller gains for equity indices and international equities are consistent with news arriving after resumption, which raises forecast errors without changing the predictor. Figure IA.1 repeats Figure 2 for both endpoints.

Table IA.3: Information-Sample Availability and Forecasting Gates

Panel A. Sequential weekend sample flow					
Requirement	Retained	Excluded	Contracts		
Weekend internal-pricing episodes	1,882	0	103		
Eligible source resumption and rule regime	1,882	0	103		
Scheduled external-source anchor	1,863	19	100		
Fresh, positive anchor prices	1,749	114	99		
Valid final internal quote	1,599	150	99		
Valid quotes at all four evaluation points	1,214	385	94		
Complete source-resumption target	1,211	3	94		
Benchmark target begins at or after resumption	1,158	53	94		
Complete benchmark-reopening target	1,119	39	91		
Available predetermined controls: common sample	1,114	5	91		
Market and closure type represented in training	959	155	77		
At least 30 training observations	866	93	77		
At least 4 training weeks	866	0	77		
Local path adds an identified regressor	866	0	77		
At least 20 residual degrees of freedom	866	0	77		
At least 10 training rows per parameter: forecasting sample	181	685	16		
Panel B. Asset composition and endpoint requirements					
Market group	Base	Source point only	Paired point only	Common full path	Forecasting sample
All markets	1,882	1,566	1,454	1,114	181
U.S. stocks and ADRs	1,117	945	945	715	0
U.S.-listed ETFs	148	107	107	57	0
International equities	218	160	48	40	19
Equity indices	106	82	82	65	32
Commodity futures	121	117	117	103	68
Spot metals	104	103	103	90	62
Foreign exchange	68	52	52	44	0

Notes: Panel A applies requirements sequentially, so each excluded episode appears once, at its first failed requirement. The common sample is the paired full-path sample in Table 3. Forecasting counts refer to the baseline model with contract indicators at the final evaluation point and count each episode once across the two endpoints.

In Panel B, the two point-only columns require only the final internal quote, the indicated target(s), and predetermined controls. The common sample additionally requires quotes at all four evaluation points.

Quotes must be at most 180 seconds old. Training uses all eligible closure types within the asset class and only targets realized before the forecast month. Sequential counts depend on the order shown.

Table IA.4: Characteristics of Retained and Excluded Weekend Episodes

Panel A. Weekend base to paired full-path sample					
Characteristic	Retained <i>N</i>	Excluded <i>N</i>	Retained median	Excluded median	Std. mean difference
Duration (hours)	1,114	768	48.02	48.02	0.232
Quote-age p95 (seconds)	1,114	768	57.38	917.69	-0.627
Fresh quotes (%)	1,114	768	99.48	69.30	1.583
Order-book trading per hour (\$)	1,114	768	75,629	9,058	-0.077
Mean top-of-book depth (\$)	1,114	768	5,394	2,796	-0.013
Mean spread (bp)	1,114	768	3.11	17.50	-0.905
Panel B. Paired full-path sample to forecasting sample					
Characteristic	Retained <i>N</i>	Excluded <i>N</i>	Retained median	Excluded median	Std. mean difference
Duration (hours)	181	933	49.00	48.02	-0.112
Quote-age p95 (seconds)	181	933	34.58	58.73	-0.286
Fresh quotes (%)	181	933	99.93	99.41	0.259
Order-book trading per hour (\$)	181	933	221,570	69,495	0.474
Mean top-of-book depth (\$)	181	933	7,094	5,228	0.226
Mean spread (bp)	181	933	1.03	3.62	-0.410

Notes: Each observation is a weekend internal-pricing episode. Panel A compares the common information sample with excluded base episodes. Panel B compares forecasting-sample episodes with common-sample episodes that fail the forecasting gates. The *N* columns report nonmissing observations for each characteristic.

Characteristics are calculated over the same scheduled internal interval for both groups. Quote age is the age of the best local quote in the minute panel. Freshness uses the 180-second limit of the main sample. Trading is order-book trading per scheduled hour; depth is the sum of dollar bid and ask depth at the best quotes.

The standardized mean difference is the retained-minus-excluded mean divided by the square root of the average of the two sample variances. It accompanies medians to make skewness visible. Asset composition is reported in Table IA.3.

Table IA.5: The Benchmark-Reopening Endpoint

Panel A: Mapping into the Oracle after Benchmark Reopening

	(1) Equity indices	(2) International equities	(3) U.S. equities and ADRs	(4) U.S. ETFs
Local price change	0.605*** (0.169)	0.633* (0.159)	0.552*** (0.158)	0.254 (0.532)
Wild p : $\beta = 1$	0.053	0.118	0.020	0.226
Wild equivalence-test p	0.851	0.858	0.925	0.837
Bootstrap 95% interval	[0.111, 1.035]	[0.155, 1.145]	[0.284, 0.865]	[-0.739, 1.518]
Difference from resumption	-0.329	-0.101	-0.279	-0.553
Wild p : difference = 0	0.007	0.066	0.106	0.334
95% interval, difference	[-0.5938, -0.0005]	[-0.3047, 0.1100]	[-0.5667, 0.0098]	[-1.5417, 0.7134]
Bootstrap draws used	999	999	999	999
Contract fixed effects	Yes	Yes	Yes	Yes
Predetermined controls	Yes	Yes	Yes	Yes
Episodes	65	40	715	57
Contracts	4	5	60	11
Resumption weeks	33	11	31	21
R^2	0.512	0.548	0.162	0.136

Notes: The dependent variable is the change in the oracle from the last external price to the mean oracle over the 30 minutes after benchmark reopening, in basis points. The regressor is the local price change from the start of internal pricing to one minute before resumption. All regressions include contract fixed effects and predetermined controls: episode length, the initial basis between the local price and the oracle, and oracle returns over the preceding 30 minutes and 24 hours. Standard errors, clustered by resumption week, are in parentheses. Stars test $\beta = 0$ at the 10, 5, and 1 percent levels with a null-imposed wild cluster bootstrap- t over resumption weeks; the same procedure tests $\beta = 1$ and equivalence to one within $[0.8, 1.2]$, using 9,999 weights, or all sign weights with at most 10 weeks. Bootstrap intervals resample resumption weeks (999 draws) and re-estimate the fixed effects in each draw. Observations, regressor, and controls match Table 3. The difference rows test the benchmark-reopening slope minus the resumption slope on the same episodes and bootstrap draws. Commodity futures, foreign exchange, and spot metals are omitted because their benchmark reopens at resumption.


Table IA.5: The Benchmark-Reopening Endpoint

Panel B: Forecast Gain against the Oracle after Benchmark Reopening

	(1) Commodity futures	(2) Equity indices	(3) International equities
Forecast gain (OOS R^2)	0.590	0.182	0.059
Clark–West p	0.009	< 0.001	0.003
Leave-one-week range	[0.560, 0.848]	[0.153, 0.241]	[-0.008, 0.100]
Weeks with lower forecast error	16 / 18	9 / 12	3 / 6
Episodes	68	32	19
Test weeks	18	12	6
Forecast months	4	3	1
Contract indicators	Yes	Yes	Yes
Closure-type indicators	Yes	Yes	Yes
Predetermined controls	Yes	Yes	Yes
	(4) Spot metals	(5) U.S. equities and ADRs	(6) U.S. ETFs
Forecast gain (OOS R^2)	0.661	–	–
Clark–West p	0.003	–	–
Leave-one-week range	[0.605, 0.725]	–	–
Weeks with lower forecast error	13 / 18	–	–
Episodes	62	–	–
Test weeks	18	–	–
Forecast months	5	–	–
Contract indicators	Yes	Yes	Yes
Closure-type indicators	Yes	Yes	Yes
Predetermined controls	Yes	Yes	Yes

Notes: The baseline model uses predetermined controls and contract and closure-type indicators; the enhanced model adds the local price change, with a separate coefficient for each closure type. The forecast gain is the incremental out-of-sample R^2 of the enhanced model over the baseline. Both models are re-estimated at the start of each month on episodes whose target windows have ended. Training requires four weeks, 30 observations, 10 observations per effective parameter, and 20 residual degrees of freedom. Forecast gains and Clark–West statistics weight episodes equally; Clark–West p -values are asymptotic, with weekly HAC standard errors of at most four lags. Leave-one-week ranges drop one test week at a time with forecasts held fixed. Dashes indicate cells without enough training history.

Figure IA.1: Information Accumulation at Resumption and at Benchmark Reopening



Notes: The left column repeats Figure 2, which targets the mean oracle over the 30 minutes after resumption. The right column targets the mean oracle over the 30 minutes after benchmark reopening, using the same episodes and local price paths. Upper panels report mapping coefficients with 95 percent week-bootstrap intervals; lower panels report forecast gains. Commodity futures, foreign exchange, and spot metals coincide across the two columns because their benchmark reopens at resumption.

Table IA.6: Pooled Forecast Sensitivity

Sensitivity Specification: Asset-Class Pooling

	(1) Commodity futures	(2) FX	(3) Equity indices	(4) International equities
Resumption: forecast gain	0.748	–	0.739	0.266
Clark–West p	< 0.001	–	0.001	< 0.001
Leave-one-week range	[0.717, 0.764]	–	[0.611, 0.795]	[0.158, 0.403]
Weeks with lower forecast error	20 / 21	–	15 / 17	5 / 6
Benchmark reopening: forecast gain	0.748	–	0.188	0.143
Clark–West p	< 0.001	–	< 0.001	0.002
Leave-one-week range	[0.717, 0.764]	–	[0.129, 0.296]	[0.083, 0.185]
Weeks with lower forecast error	20 / 21	–	9 / 17	5 / 6
Episodes	80	–	44	19
Test weeks	21	–	17	6
Forecast months	5	–	4	1
Contract indicators	No	No	No	No
Closure-type indicators	Yes	Yes	Yes	Yes
Predetermined controls	Yes	Yes	Yes	Yes
	(5) Spot metals		(6) U.S. equities and ADRs	(7) U.S. ETFs
Resumption: forecast gain	0.494		0.765	–
Clark–West p	< 0.001		0.001	–
Leave-one-week range	[0.406, 0.575]		[0.736, 0.779]	–
Weeks with lower forecast error	16 / 23		20 / 23	–
Benchmark reopening: forecast gain	0.494		0.045	–
Clark–West p	< 0.001		0.058	–
Leave-one-week range	[0.406, 0.575]		[0.009, 0.086]	–
Weeks with lower forecast error	16 / 23		15 / 23	–
Episodes	75		558	–
Test weeks	23		23	–
Forecast months	6		6	–
Contract indicators	No		No	No
Closure-type indicators	Yes		Yes	Yes
Predetermined controls	Yes		Yes	Yes

Notes: This sensitivity pools contracts within each asset class by removing the contract indicators; it was specified after the baseline results. The forecast gain is the incremental out-of-sample R^2 of the enhanced model over the baseline. Both models are re-estimated at the start of each month on episodes whose target windows have ended. Training requires four weeks, 30 observations, 10 observations per effective parameter, and 20 residual degrees of freedom. Forecast gains and Clark–West statistics weight episodes equally; Clark–West p -values are asymptotic, with weekly HAC standard errors of at most four lags. Leave-one-week ranges drop one test week at a time with forecasts held fixed. Dashes indicate cells without enough training history.

Table IA.7: Local Price Dynamics under External and Internal Pricing

Panel A: Matched Source-State Comparison

Oracle state	Contracts	Episodes	Median spell hours	$VR(30)$	$AR(1)$
				0.783	0.027
				[0.746, 0.830]	[0.013, 0.041]
External pricing	42	728	13.5	{0.701, 0.864}	{0.009, 0.046}
				0.878	0.010
				[0.784, 0.981]	[-0.014, 0.032]
Internal pricing	42	451	48.0	{0.717, 1.099}	{-0.021, 0.050}
				0.095	-0.016
				[-0.012, 0.209]	[-0.046, 0.010]
Internal minus external pricing	42	—	—	{-0.085, 0.323}	{-0.061, 0.024}

Notes: The table measures the dynamics of the local price, not of the oracle. Each spell is a continuous interval with one scheduled oracle state, external or internal pricing, following the venue’s documented schedule (Section 2). The sample retains contract–closure-type cells observed in both states, excludes one-minute boundary fragments and unverified schedule intervals, and requires a complete six-hour path of local prices no more than 90 seconds old; returns never cross a change of state. $VR(30)$ is the finite-sample Lo–MacKinlay overlapping variance ratio, $\hat{\sigma}^2(30)/(30\hat{\sigma}^2(1))$, and $AR(1)$ is the one-minute autocorrelation of local price changes within a spell. Both are computed within each spell; spells are equally weighted within contract–closure-type–state cells, closure types within contracts, and contracts equally. Square brackets contain 95% contract-block percentile intervals and braces 95% resumption-week-block percentile intervals, each from 999 resamples. The difference is internal minus external pricing on the matched sample. Values below one indicate transitory mean reversion.

Table IA.8: Post-Resumption Reversals by Asset Class and Closure Type

Asset class	Closure	Reversal slope	Clustered SE	Wild p	Bootstrap 95% interval	N	Contracts	Weeks
<i>Panel A: Reversal Slopes by Asset Class and Closure Type</i>								
Commodity futures	Weekend	-0.045	(0.136)	0.758	[-0.305, 0.223]	115	4	32
Foreign exchange	Weekend	-0.292	(0.114)	0.257	[-0.452, 0.089]	50	3	27
Equity indices	Weekend	-0.157	(0.134)	0.309	[-0.407, 0.111]	79	4	33
International equities	Weekend	-0.203**	(0.063)	0.015	[-0.354, -0.091]	160	8	28
Spot metals	Weekend	-0.060	(0.199)	0.775	[-0.385, 0.385]	100	4	28
U.S. equities and ADRs	Weekend	-0.204***	(0.046)	< 0.001	[-0.290, -0.100]	922	64	31
U.S. ETFs	Weekend	-0.357***	(0.083)	0.002	[-0.532, -0.185]	100	12	25
International equities	Lunch break	-0.776**	(0.195)	0.012	[-1.071, -0.276]	178	6	12
Spot metals	Maintenance	-0.560***	(0.149)	< 0.001	[-0.842, -0.245]	461	4	37
Commodity futures	Maintenance	-0.189	(0.216)	0.417	[-0.579, 0.223]	360	3	35

Notes: Each row regresses the reversal, the change in the local price from resumption o to $o + 30$ minutes, on the local price change from the start of internal pricing c to $o - 5$ minutes, both in basis points, with contract fixed effects and controls for episode length and the initial basis between the local price and the oracle. The five-minute gap separates the two changes, and the quote at o must be newer than the quote at $o - 5$; prices are at most 180 seconds old. A negative slope indicates that the preceding local move partly reverses. Standard errors are clustered by resumption week. Wild p -values and stars use the null-imposed week-cluster wild bootstrap- t at the 10, 5, and 1 percent levels. Bootstrap intervals resample resumption weeks 999 times and are reported when at least 499 draws identify the slope. Episodes with unverified transitions are excluded.

A.5 Displayed Liquidity

Table IA.9 reports the censoring of the two displayed-liquidity measures in Table 4, and Table IA.10 tests prediction P3 within asset classes by interacting the local price change with standardized displayed-liquidity measures in the weekend mapping regression.

Table IA.9: Censoring of the Displayed-Liquidity Measures

Panel A. Support rate						
Group	Episodes	Supported valid (%)	Failed valid (%)	Unknown valid (%)	Invalid all (%)	Lower bound all (%)
All contracts	1,846	44.25	54.70	1.06	0.18	44.16
U.S. stocks and ADRs	1,098	42.22	56.40	1.39	0.20	42.12
U.S.-listed ETFs	138	26.94	72.84	0.22	0.23	26.87
International equities	217	20.84	78.12	1.04	0.12	20.82
Equity indices	106	48.58	51.31	0.11	0.14	48.50
Commodity futures	121	76.80	22.18	1.03	0.14	76.67
Spot metals	104	76.91	22.74	0.35	0.16	76.78
Foreign exchange	62	74.83	25.11	0.06	0.18	74.70

Panel B. Shift cost: measurement coverage						
Group	Episodes	Exact-p05 episodes	Exact valid (%)	Interval valid (%)	Right-cens. valid (%)	Invalid all (%)
All contracts	1,846	1,299	37.48	6.17	56.35	0.18
U.S. stocks and ADRs	1,098	826	35.62	6.96	57.41	0.20
U.S.-listed ETFs	138	115	55.35	7.69	36.96	0.23
International equities	217	169	59.08	5.88	35.04	0.12
Equity indices	106	54	42.90	3.21	53.89	0.14
Commodity futures	121	54	13.18	2.00	84.81	0.14
Spot metals	104	42	19.85	4.40	75.75	0.16
Foreign exchange	62	39	22.59	6.01	71.39	0.18

Panel C. Shift cost: medians of episode-level fifth percentiles (\$)			
Group	Exact minutes exact-available episodes	All valid minutes same episodes	All valid minutes all episodes
All contracts	6,056	5,355	10,965
U.S. stocks and ADRs	7,383	5,836	10,025
U.S.-listed ETFs	2,183	2,164	5,027
International equities	2,504	1,039	2,760
Equity indices	643	494	6,692
Commodity futures	21,323	12,359	23,143
Spot metals	9,580	19,404	34,463
Foreign exchange	70,870	50,014	54,984

Notes: The sample is the weekend sample of Table 4, and all percentages are episode-equal means. In Panel A, the trade is a \$10,000 purchase and an equal-sized sale, each with a volume-weighted execution price within 10 basis points of the local price in the same snapshot. Supported, failed, and unknown are mutually exclusive states among usable snapshots and sum to 100 percent; unknown means that the twenty recorded levels do not determine support. The invalid share uses all scheduled minutes. The final column divides confirmed support by all scheduled minutes and equals the support rate in Table 4.

In Panels B and C, the shift cost is the displayed notional that moves the local price by 25 basis points. Exact, interval-censored, and right-censored observations partition the usable measurements and sum to 100 percent; invalid measurements use all scheduled minutes. An exact-p05 episode contains at least one exact minute. Panel C compares the fifth percentile over exact minutes with the fifth percentile of lower bounds over all usable minutes, first in the same episodes and then in all episodes. All dollar entries are medians across episodes.

Table IA.10: Displayed Liquidity and the Weight on Local Prices

	(1)	(2)	(3)	(4)	(5)	(6)
Sample	Pooled	Pooled	Pooled	Pooled	U.S. equities	U.S. equities
Liquidity timing	Same	Same	Lagged	Lagged	Same	Same
Liquidity measure	Support	Log shift	Support	Log shift	Support	Log shift
$X \times L$ (δ)	0.028 (0.044)	0.038 (0.037)	0.059 (0.038)	0.041 (0.035)	0.008 (0.037)	0.015 (0.034)
Wild p : $\delta = 0$	0.623	0.458	0.218	0.328	0.883	0.848
Wild p : $\delta > 0$	0.307	0.225	0.116	0.162	0.434	0.422
Bootstrap 95% interval	[-0.06, 0.11]	[-0.05, 0.10]	[-0.02, 0.13]	[-0.04, 0.12]	[-0.07, 0.09]	[-0.06, 0.07]
Slope at mean liquidity	—	—	—	—	0.829	0.831
Observations	1,114	1,114	1,066	1,066	715	715
Contracts	91	91	88	88	60	60
Resumption weeks	38	38	37	37	31	31

Notes: The sample, target, local price change, controls, and contract fixed effects are those of Table 3, Panel A (weekend internal-pricing episodes; local change one minute before scheduled resumption; mean oracle over the first 30 minutes after resumption). Each column adds the interaction of the local change with a displayed-book measure L , standardized to mean zero and unit standard deviation within asset class, together with the main effect of L . Pooled columns allow the slope on the local change, the main effect of L , and all controls to differ by asset class; δ is common. δ is the change in the mapping slope per standard deviation of L ; prediction P3 implies $\delta > 0$. Support is the fraction of scheduled minutes with confirmed two-sided support for a \$10,000 trade within 10 basis points of the local price; log shift cost is $\log(1 + \cdot)$ of the within-episode fifth percentile of the displayed notional required to move the local price by 25 basis points, both as in Table 4. Same-episode measures are contemporaneous; lagged measures use the same contract's preceding weekend internal-pricing episode and are predetermined. Parentheses contain week-clustered standard errors. Two-sided and one-sided p -values use the null-imposed week-cluster wild bootstrap- t of Table 3; intervals resample resumption weeks 999 times, with fixed effects re-estimated in each draw. The slope at mean liquidity is reported for single-class columns.

A.6 Holiday Design and Account Reallocation

Table IA.11 lists the holiday events and their inference blocks. Table IA.12 reports all nine market-quality outcomes examined, together with a sensitivity that excludes dates near announced changes to the oracle’s parameters, and Table IA.13 holds the event set fixed across relative weeks. Table IA.14 reports account-level inference under alternative block definitions, and Figure IA.2 plots account activity from four weeks before to four weeks after each holiday.

Table IA.11: Holiday Events and Inference Blocks

2026 dates	Calendar	Contracts	Contract-hours	UTC hours	Week block	Month block
01-19	US	21	84	4	1	1
02-16	US	25	100	4	2	2
03-02	KR	3	18	6	3	3
04-29–05-06	JP/KR	5	64	24	4	4
05-25	US	48	192	4	5	4
06-03	KR	4	24	6	6	5
06-19	HK/US	60	182	8	7	5
07-01	HK	2	10	5	8	6
07-03	US	64	192	3	8	6
07-17	KR	4	24	6	9	6
07-20	JP	3	12	4	10	6
08-11	JP	3	12	4	11	7
08-17	KR	4	24	6	12	7

Notes: The design has 13 holiday events, 12 treatment-week blocks, and 7 treatment-month blocks. Dates summarize the eligible treatment hours in UTC. US, JP, KR, and HK denote the U.S., Japanese, Korean, and Hong Kong equity calendars. Golden Week and simultaneous holidays are treated as single events. Blocks join events that share a treatment week or month. The calendar-assignment test requires common treated contract-hours across eligible dates, preserving at least 80% of actual treatment cells. Same-weekday dates retain local UTC offsets, ordinary opening arrangements, and at least two other open control calendars. Actual dates belong to the assignment support.

Table IA.12: Complete Holiday Outcomes and Calendar-Assignment Sensitivity

Outcome	All events $\hat{\beta}$	Assignment $\hat{\beta}$	p	Notice $\hat{\beta}$	p
Log spread	0.265	0.145	0.094	0.110	0.223
Log dollar depth	-0.712	-0.744	0.001	-0.715	0.001
Log order-book trading	-1.622	-1.519	0.001	-1.344	0.001
Log active makers	-0.273	-0.315	0.001	-0.294	0.012
No-trade minute share	0.148	0.125	0.003	0.129	0.003
Local-price update share	-0.221	-0.215	0.001	-0.198	0.001
Signed local price-oracle basis (bp)	-7.705	-6.970	0.111	-8.452	0.133
Absolute basis (bp)	8.734	9.604	0.001	8.229	0.001
Basis dispersion (bp)	4.125	4.872	0.006	6.247	0.005
Common observations	Varies	6,281		4,083	
Holiday events / treatment weeks		12 / 11		11 / 10	

Notes: The table reports all nine outcomes examined; the main text reports top-of-book depth, order-book trading, and local-price updating. All-event coefficients use each outcome's available observations. The calendar-assignment test and its notice-restricted version hold the regression rows fixed across observed and reassigned dates, using rows common to all nine outcomes. Each test uses 999 feasible joint assignments and the plus-one Monte Carlo correction. The notice-restricted version also excludes candidate dates near announced changes to the oracle's parameters. Basis measures also reflect changes in the oracle's rules over the sample.

Table IA.13: Event-Week Dynamics on a Common Event Set

Outcome	All events β_0	SE	Pre p	Common set β_0	SE	Pre p
Log spread	0.088	(0.097)	< 0.001	-0.043	(0.129)	0.157
Log dollar depth	-0.780	(0.168)	0.271	-0.520	(0.118)	0.525
Log order-book trading	-1.409	(0.540)	0.289	-0.843	(0.508)	0.273
Log active makers	-0.196	(0.105)	0.004	-0.095	(0.125)	0.135
No-trade minute share	0.057	(0.054)	< 0.001	0.004	(0.074)	0.019
Local-price update share	-0.162	(0.049)	0.337	-0.122	(0.054)	0.161
Signed local price-oracle basis (bp)	-7.184	(11.026)	0.574	-13.197	(17.213)	0.876
Absolute basis (bp)	7.802	(4.473)	0.754	10.181	(7.545)	0.857
Basis dispersion (bp)	3.471	(2.860)	0.444	5.674	(4.513)	0.440
Holiday events	11			5		


Notes: Both specifications omit week -1 and absorb event-contract-hour-of-week and event-exact-hour fixed effects. The common set keeps only events observed in every relative week from -4 to $+4$, so event membership is constant across weeks. SE denotes standard errors clustered by contract and event; Pre p jointly tests weeks -4 , -3 , and -2 . Covariance matrices are adjusted to be positive semidefinite when required.

Table IA.14: Account Inference under Alternative Event Blocks

Outcome	$\hat{\beta}$	Event p	Week p	Month p	Month std. p
Active	-0.0076	0.0029	0.0059	0.0781	0.0781
Maker active	-0.0024	0.0107	0.0068	0.0938	0.0625
Asinh order-book trading	-0.0179	0.0046	0.0093	0.0938	0.0469
Asinh maker notional	-0.0090	0.0027	0.0049	0.0312	0.0312
Inference blocks		13	12	7	7
Sign assignments		8,192	4,096	128	128

Notes: The coefficient, observations, and three fixed effects are identical in every column. Scores are summed within blocks defined by the event, the shared treatment week (the main specification), or the shared treatment month, and each column enumerates all sign assignments of the block scores. The final column rescales month-block scores by their within-block dispersion. Month blocks allow wider shared shocks and give larger p -values.

Figure IA.2: Account Reallocation around Holiday Closures



Notes: The figure reports event-equal mean shares on the fitted account-regression rows from four weeks before through four weeks after each holiday event. The cohort is frozen before each event; event composition varies across relative weeks. Affected contracts have benchmarks closed by the holiday during the event hour; control contracts have benchmarks open in the same UTC hour. Panel A reports order-book activity and Panel B maker activity. The week -1 , 0 , and $+1$ levels coincide with Table 7.

A.7 Resumption

This section describes price adjustment and forced trading at all scheduled resumptions, the setting for the dollar stakes in Section 8. Table IA.15 summarizes 4,640 scheduled resumptions across 103 contracts and 214 dates.

Table IA.15: Operational Outcomes around Scheduled Resumption

Closure type	Events	Contracts	Resumption dates	Mark adjustment (bp)		Post-resumption pressure	
				5-minute absolute	30-minute maximum	Detected forced-flow events	Backstop/ADL events
Panel A: Event Census and Operational Adjustments							
Holiday	368	77	44	40.70	87.55	56 (15.2%)	1 (0.27%)
Lunch break	186	6	51	59.88	181.77	52 (28.0%)	0 (0.00%)
Maintenance	865	7	137	12.13	29.00	91 (10.5%)	1 (0.12%)
Overnight	1,339	19	172	36.08	107.90	353 (26.4%)	5 (0.37%)
Weekend	1,882	103	91	41.41	89.07	420 (22.3%)	5 (0.27%)
All events	4,640	103	214	30.32	78.25	972 (20.9%)	12 (0.26%)
Panel B: Mark–Oracle Convergence							
Median first stable-convergence minute				2.0			
Right-censored at 120 minutes				565 / 4,640 (12.18%)			

Notes: Each observation is a scheduled resumption in one contract, the transition from internal to external pricing under the documented schedule. Events on the same date can share common shocks.

In Panel A, the five-minute adjustment is $10,000|\log(\text{Mark}_{t+5}/\text{Mark}_{t-1})|$. The 30-minute measure is the maximum absolute cumulative mark movement relative to Mark_{t-1} over minutes 1–30. Medians use available mark paths.

Forced-flow events have $\sum_{k=1}^{30} |F_{t+k}| > 0$, where F is dollar forced-taker purchases minus sales within a contract-minute. Absolute values precede aggregation across minutes; opposing executions within a minute can offset. The statistic is detected net forced flow. Backstop/ADL events have positive separately recorded backstop or auto-deleveraging notional in the same interval.

Percentages use all events in each row. Of the 4,640 events, 4,639 have all 30 flow minutes; 1 has 14 minutes and no detected pressure. On complete windows, detected forced flow and backstop/ADL occur in 20.95% and 0.26%, respectively.

Panel B defines stable convergence as the first entry of the absolute mark-oracle log gap into a 10 bp band, maintained for 5 consecutive minutes. The median conditions on observed convergence; right-censoring records no completed qualifying spell within 120 minutes.

The median absolute adjustment of the mark price is 30.32 basis points over the first five minutes, and the median maximum adjustment is 78.25 basis points over the first 30 minutes. Adjustments are largest after lunch breaks (medians of 59.88 and 181.77 basis points), which occur only in six international-equity contracts, and smallest after maintenance breaks (12.13 and 29.00 basis points), in three commodity-futures and four spot-metal contracts. These differences reflect products and trading clocks jointly.

Forced order-book flow, detected when at least one minute has nonzero net forced taker trading, appears in the first 30 minutes of 972 resumptions, or 20.9 percent; because purchases and sales within a minute can offset, this is a detection rate rather than a count of forced trades. Backstop liquidation or auto-deleveraging appears in 12 resumptions, or 0.26 percent. The mark price converges quickly to the oracle: conditional on convergence, defined as the gap staying within 10 basis points for five consecutive minutes, the median starting minute is two. In 565 events, or 12.18 percent, convergence is not observed within two hours. Resumption is thus usually smooth, with a tail of slow adjustment.

A.8 Reference-Price Rules

Table IA.16 repeats the primary comparison of Table 8 with each candidate constrained by the venue’s documented price bound. Imposing the documented bounds in training and evaluation lowers the improvement over the oracle from 4.20 to 3.51 basis points, and allowing the documented reset budget gives 3.71. The oracle lies outside the documented bound in 0.61 percent of episodes and outside the reset budget in 0.24 percent, so the documented parameters are consistent with the oracle’s realized position in nearly all episodes.

Table IA.16: Reference-Price Rules under Documented Discovery Bounds

<i>Panel A: Selected rule versus the implemented oracle, one minute before resumption</i>							
Constraint on candidate	Mean selected λ	Binding (%)	Candidate MAE	Oracle MAE	Episode-mean Δ loss	Week-mean Δ loss	Sign p -value
Unconstrained (Table 8)	0.66	–	83.64	89.51	-5.87	-4.20 [-6.32, -2.00]	< 0.001
Documented bound	0.65	0.48	84.53	89.51	-4.98	-3.51 [-5.55, -1.57]	0.001
Bound and reset budget	0.65	0.27	84.19	89.51	-5.32	-3.71 [-5.66, -1.62]	< 0.001
75% of documented bound	0.64	1.44	85.06	89.51	-4.45	-2.90 [-5.03, -0.69]	0.007
50% of documented bound	0.64	3.36	86.67	89.51	-2.84	-1.21 [-4.26, 2.04]	0.238
<i>Panel B: Prices relative to the documented bound, one minute before resumption</i>							
Measure	Value						
Implemented oracle outside the documented bound (%)	0.61						
Implemented oracle outside the documented reset budget (%)	0.24						
Post-resumption official price outside the documented bound (%)	1.07						
Unconstrained candidate outside the documented bound (episodes)	18						
Share of the unconstrained episode-sum gain from these episodes (%)	3.0						

Notes: The sample is the 3,749 primary episodes of Table 8 (102 contracts, 32 resumption weeks). Each candidate’s cumulative log change from the last external price is projected into the documented Discovery Bound around that price, in training as well as in evaluation, and the monthly weight is re-selected under the constraint with the grid, training gates, and 24-hour availability lag of Table 8. The unconstrained row equals the comparison in Table 8. Bounds are the per-market values in the specification snapshot of September 7, 2026 (median 10%, range 2–20%), back-dated with two dated notices: before June 12, 2026, the bound is 10% rather than 5% for five large U.S. equities and resets are zero for the 41 contracts whose resets were raised on that date; before March 13, 2026, crude oil has no resets. On a notice date the tighter setting applies. 1,090 episodes use a back-dated value. The reset budget allows up to $1 + R$ successive bounds, the farthest range reachable with R resets. The stress rows scale the documented bound without resets. Binding is the share of episodes in which the selected candidate is projected. Δ loss is candidate absolute error minus implemented-oracle absolute error in basis points; negative values favor the candidate. Intervals resample resumption weeks 1,999 times; p -values are one-sided week sign-randomization tests.