Université du Québec en Outaouais

Département des sciences administratives

Working Paper No. 39

The Invisible Balance Sheet

Can Intraday Financial Markets Detect Corporate Deterioration Before Accounting Disclosures?

Simon-Pierre Boucher

Département des sciences administratives
Université du Québec en Outaouais

simon-pierre.boucher@uqo.ca

Gatineau – Pavillon Alexandre-Taché
283, boulevard Alexandre-Taché
Gatineau, Québec, Canada J9A 1L8

November 2026

Version 1.0

Abstract

Do one-minute financial markets see a firm's forthcoming accounting deterioration before daily markets do? We build a strictly point-in-time panel of 3,796 U.S. listed firms and 102,522 firm-quarters (2013–2026) from 6.7 billion raw one-minute bars and quarterly statements with filing and release dates, and forecast an adverse operating-income surprise twenty trading days before the earnings release. Holding the information set, the universe, the horizon and the model capacity fixed, we add the *same* families of realized measures, volume composition and market co-movement computed from 30-minute, 5-minute and 1-minute bars to a benchmark that already contains point-in-time fundamentals and daily OHLCV. Fundamentals explain most of the predictable variation, daily market data add a clear increment (log loss -0.0068 with gradient-boosted trees), and one-minute data add nothing: the paired out-of-sample change in log loss is +0.0013 [-0.0013, +0.0041] with logistic regression and +0.0016 [+0.0010, +0.0023] with trees, intervals that exclude any gain at all with trees and any gain larger than a fifth of the daily increment with logistic regression. The null survives five horizons, five accounting targets, a symmetric volatility placebo, unseen firms, every refit year and every pre-specified subgroup; the small information content of the incremental one-minute signal is concentrated in large, levered firms rather than in neglected ones, is already impounded in prices at the cutoff, and has no economic value in factor-adjusted portfolio sorts. The accounting panel is conditioned on firms listed in November 2026, a survivorship caveat we document rather than resolve. The invisible balance sheet is visible in the daily close; the minute-by-minute tape adds noise, not foresight.

Keywords: intraday data, sampling frequency, earnings surprises, fundamental analysis, point-in-time data, machine learning, market efficiency

JEL Classification: G14, G12, G17, M41, C53, C58

1. Introduction

Every quarter, the accounting system reveals with a lag what the firm has already lived through: revenues that disappointed, margins that compressed, cash that did not come in. Market prices are widely believed to anticipate part of this revelation, and a long literature shows that they do at daily and annual horizons (Ball and Brown, 1968; Beaver et al., 1980; Kothari and Sloan, 1992; Collins et al., 1994). A newer belief, common among practitioners and increasingly in the academic literature on high-frequency data, is that the *texture* of trading, the minute-by-minute composition of volume and returns, contains information that the daily close does not: informed traders split orders and trade when liquidity is deep (Kyle, 1985; Easley and O'Hara, 1987), jumps and signed variation mark the arrival of firm-specific news (Lee, 2012; Patton and Sheppard, 2015; Bollerslev et al., 2020), and attention-constrained investors process public information gradually (Hirshleifer and Teoh, 2003; DellaVigna and Pollet, 2009). If this belief is right, one-minute bars should forecast forthcoming accounting deterioration better than daily bars, holding everything else fixed. This paper tests that proposition.

The test is designed around one distinction, anticipation versus revelation, and one object, the marginal information value of sampling frequency. Anticipation, the fact that prices move before statements, is not at issue; it is expected under semi-strong efficiency (Fama, 1970; Grossman and Stiglitz, 1980) and we reproduce it. What is at issue is whether observing the market at a finer resolution adds forecasting content about the *accounting* outcome, over and above (i) point-in-time fundamentals, (ii) daily OHLCV, (iii) coarser aggregates of the same bars and (iv) size, industry and calendar controls. To answer it we build a nested ladder of information sets with the same universe, the same calendar coverage, the same cutoff and horizon, the same feature *families* computed from 30-minute, 5-minute and 1-minute bars, and the same tuning budget per set, so that any difference is attributable to resolution rather than to dimensionality or capacity.

The data are the owner's FinData OS lakehouse of FirstRate Data: 6.7 billion raw 1-minute bars for 16,195 U.S. stock symbols (2000–2026, delisted names included), quarterly fundamentals with filing dates for 5,451 entities, and 373,326 earnings-release events with their timing. Point-in-time discipline is strict: the forecaster at cutoff c sees bars up to the close of c and only those accounting facts whose vendor availability instant (filing date plus one day) precedes c; the reveal date is the earnings press release, not the filing; and twelve leakage guards are unit tests in the replication package, including a live equivalence test of our as-of join against the API. Targets are accounting surprises against a seasonal random walk with drift (Foster, 1977), never prices.

Three findings emerge. First, the ladder works as the literature predicts: point-in-time fundamentals lower the out-of-sample log loss from 0.6282 to 0.4937 (AUC 0.797) and daily market data to 0.4869 (AUC 0.803), a daily increment of -0.0068 that is unambiguous. Second, one-minute

data add nothing. On an untouched 2021–2026 test period of 53,667 firm-quarters, the paired change in log loss from adding the 1-minute families is +0.0013 [-0.0013, +0.0041] with logistic regression and +0.0016 [+0.0010, +0.0023] with gradient-boosted trees; 30-minute and 5-minute families behave identically; the tree intervals exclude any gain at all and the logistic intervals exclude gains larger than a fifth of the daily increment. The pre-registered minimum detectable difference (0.0025) is wider than the tree intervals and comparable to the logistic ones, so this is a precise null rather than an underpowered one. The null survives five horizons from forty days to one day before the release, five accounting targets, a symmetric volatility placebo, firms never seen in training, every annual refit, and every pre-specified subgroup. Third, the little information that the incremental one-minute signal does carry sits in large and levered firms rather than in small, illiquid or neglected ones, is already in prices twenty days before the release (no announcement-window abnormal return), and has no economic value in factor-adjusted decile portfolios.

The contribution is therefore a measurement: on a historically correct universe with a strictly point-in-time information set, the marginal value of one-minute sampling for forecasting accounting deterioration is zero within tight bounds, while the value of daily sampling is large. The result parallels, for fundamentals, what Liu et al. (2015) established for variance measurement, that the finest grid is not the most informative, and it bounds the mechanisms of the high-frequency literature: whatever informed trading, jump arrival or gradual processing is visible in one-minute bars twenty days before a release is summarised by the daily close. We also deliver a reproducible point-in-time pipeline, frozen hypotheses and decision rules, and an honest statement of two structural limits: the accounting panel is survivor-conditioned (firms listed in November 2026) and the vendor's point-in-time stamps are a documented proxy.

The paper proceeds as follows. Section 2 sets out the framework and the pre-registered hypotheses, Section 3 positions the question, Section 4 describes the data, the universe and the point-in-time protocol, Section 5 the design, Section 6 the forecasting results, Section 7 the heterogeneity and pricing tests, Section 8 the negative controls, and Section 9 concludes with limitations and the research the null result motivates.

2. Economic framework and hypotheses

2.1. Anticipation versus revelation

Let $X_{i,q}$ denote an accounting flow of firm i for fiscal quarter q (operating income in our primary specification), revealed to the public at the earnings-release date $A_{i,q}$ and structured in the SEC filing of date $F_{i,q} \geq A_{i,q}$. A forecaster standing at a trading day $c < A_{i,q}$ holds an information set $\mathcal{I}(c)$ that contains, at most, prices and volumes up to the close of c and the accounting statements that were *available* at c, i.e., quarters j with $F_{i,j} + 1 \leq c$. The quantity of interest is the conditional

distribution of the surprise

$$S_{i,q} \;=\; \frac{X_{i,q} - \mathbb{E}\big[X_{i,q} \mid \mathcal{A}(c)\big]}{TA_{i,j^*(c)}}, \qquad \mathbb{E}\big[X_{i,q} \mid \mathcal{A}(c)\big] \;=\; X_{i,q-4} + \frac{1}{4}\sum_{k=0}^{3}\big(X_{i,q-1-k} - X_{i,q-5-k}\big), \quad (1)$$

where $\mathcal{A}(c) \subset \mathcal{I}(c)$ is the accounting information alone, $j^*(c)$ is the latest available quarter, and the expectation is a seasonal random walk with drift (Foster, 1977), the classic honest benchmark when no point-in-time consensus exists. The adverse event is $D_{i,q} = \mathbb{1}\big[S_{i,q} \le -0.005\big]$. Semi-strong efficiency (Fama, 1970) implies that prices at c impound all public information, including $\mathcal{A}(c)$; the "prices lead earnings" literature (Beaver et al., 1980; Kothari and Sloan, 1992; Collins et al., 1994) documents that daily and annual returns indeed forecast accounting changes. None of this is the question we ask. Our question is whether the *resolution* at which the market is observed carries incremental information, holding the information set, the universe, the calendar coverage, the horizon and the model capacity fixed.

2.2. Why resolution could matter, and why it might not

Three mechanisms make fine-grained activity informative beyond daily aggregates. First, in sequential-trade models (Kyle, 1985; Easley and O'Hara, 1987), informed traders split orders and trade when liquidity is deep, so the *composition* of a day's volume and return (intraday autocorrelation, concentration of variance in a few minutes, volume on down-bars) can reveal private information that the daily net return hides. Second, realized measures built from fine bars separate the continuous and jump components of variation (Barndorff-Nielsen, 2004; Andersen et al., 2003) and their sign (Barndorff-Nielsen et al., 2010; Patton and Sheppard, 2015); negative semivariance and signed jumps have been linked to the cross-section of expected returns (Bollerslev et al., 2020) and may proxy for the arrival of adverse firm-specific news. Third, limited attention (Hirshleifer and Teoh, 2003; Hirshleifer et al., 2009; DellaVigna and Pollet, 2009) predicts that the incorporation of public but costly-to-process information is gradual, so that the time path of trading within the day (opening and closing concentration, overnight versus intraday returns) carries information about ongoing processing.

Against these channels stand two forces. The Epps effect and microstructure noise contaminate 1-minute returns (Hansen and Lunde, 2006; Barndorff-Nielsen et al., 2008), so that the finest sampling is not the most informative for variance measurement (Liu et al., 2015); and whatever information intraday activity contains about fundamentals may be fully summarised by the daily close, in which case the frequency premium is zero. Algorithmic trading may further have reduced the amount of pre-announcement information acquisition (Weller, 2018). The sign and size of the frequency premium are therefore empirical questions, and a precise null is as informative as a

positive estimate.

2.3. Hypotheses

All hypotheses, endpoints and decision rules were frozen before the test period was opened (`docs/RESEARCH_DESIGN.md` in the replication package). Let

$$\mathcal{L}(M) \;=\; -\frac{1}{N}\sum_{(i,q)}\Big[D_{i,q}\log\hat{p}^{M}_{i,q} + (1-D_{i,q})\log\big(1-\hat{p}^{M}_{i,q}\big)\Big] \tag{2}$$

denote the out-of-sample log loss of model M for the binary adverse outcome, and let B1–B5 be the nested information sets of Table 3; the paired increment of interest is $\Delta\mathcal{L}(\mathrm{A},\mathrm{B}) = \mathcal{L}(\mathrm{A}) - \mathcal{L}(\mathrm{B})$.

H1 (incremental information). $\mathcal{L}(\mathrm{B5}) < \mathcal{L}(\mathrm{B2})$ at the 20-trading-day horizon: one-minute features improve forecasts beyond point-in-time fundamentals and daily OHLCV. Null: no improvement.

H2 (temporal resolution). The gain of B5 over B2 exceeds the gains of B4 (5-minute) and B3 (30-minute) over B2, with the same feature families and the same number of inputs. Null: the gains are equal.

H3 (heterogeneity). The B5−B2 gain is larger for small, illiquid and highly levered firms, the groups in which information asymmetry and processing costs are highest. Null: homogeneous gain.

H4 (timing). The gain increases as the cutoff approaches the release ($h = 40 \to 1$) and is absent at placebo cutoffs. Null: no gradient.

H5 (pricing). The ex-ante deterioration probability predicts the announcement-window abnormal return and post-cutoff factor-adjusted returns. Null: no return predictability after risk adjustment and costs.

A positive H1 does not identify *who* trades on what: observational prediction is not causal evidence of informed trading, and we do not have quotes, trades or trader identities (Grégoire and Martineau, 2022; Hendershott et al., 2015). The paper's claims are confined to the forecasting content of anonymous bars and to the frequency at which that content appears.

3. Related literature

Prices lead earnings. Since Ball and Brown (1968), accounting research has documented that stock prices anticipate accounting income: annual returns forecast earnings changes (Beaver et al., 1980), leading returns raise earnings-response coefficients (Kothari and Sloan, 1992), and the low contemporaneous return–earnings association reflects prices' timeliness rather than noise (Collins et al., 1994). The post-earnings-announcement drift (Bernard and Thomas, 1989, 1990) and the

accrual anomaly (Sloan, 1996; Richardson et al., 2005) show that the incorporation of accounting news is incomplete, while fundamental analysis (Ou and Penman, 1989; Piotroski, 2000) extracts forecasts of future earnings from statements alone. Recent work replaces hand-crafted ratios by machine learning on detailed statement items (Chen et al., 2022; Easton et al., 2024; Cao et al., 2024). All of these studies operate at daily or lower frequency and treat the market side as a return; our contribution is to fix the accounting side and vary the *resolution* of the market side.

Information in intraday activity. Market microstructure theory ties the composition of trading to private information (Kyle, 1985; Easley and O'Hara, 1987), and realized-measure econometrics separates continuous variation, jumps and their signs from high-frequency bars (Andersen and Bollerslev, 1998; Andersen et al., 2003; Barndorff-Nielsen, 2004; Barndorff-Nielsen et al., 2010; Patton and Sheppard, 2015; Lee and Mykland, 2008). Signed jump variation and downside semivariance price the cross-section (Bollerslev et al., 2020, 2016b), jumps cluster around news (Lee, 2012; Bajgrowicz et al., 2016), and announcement days change betas and the volume–volatility relation (Patton and Verardo, 2012; Bollerslev et al., 2018b; Savor and Wilson, 2016). The question of the best sampling frequency has been settled for *variance measurement* – five-minute realized variance is hard to beat (Liu et al., 2015; Bollerslev et al., 2016a, 2018a) – but not for the *forecasting of fundamentals*, which is our object. Zhang et al. (2024) show that pooling intraday commonality improves machine-learning volatility forecasts; we include volatility families at every frequency and use a symmetric placebo target to separate a volatility channel from a signed-information channel.

Pre-announcement trading and price discovery. Short sellers (Christophe et al., 2004), individual investors (Kaniel et al., 2012), institutions (Hendershott et al., 2015) and option traders trade informatively before earnings announcements; differences of opinion shape announcement returns (Berkman et al., 2009); algorithmic trading may reduce pre-announcement information acquisition (Weller, 2018); and the announcement reaction itself occurs largely in quotes rather than trades (Grégoire and Martineau, 2022; Brogaard et al., 2019). These studies use trader-type or quote data that bars do not contain. We ask a complementary question with anonymous OHLCV only: how much of the forthcoming accounting outcome is visible in the shape of pre-release trading, and does that visibility depend on resolution?

Attention and processing. Limited attention (Hirshleifer and Teoh, 2003; Hirshleifer et al., 2009; DellaVigna and Pollet, 2009; Aboody et al., 2010) and media coverage (Fang and Peress, 2009) govern how fast public information is processed; the timing and heterogeneity hypotheses H3–H4 are motivated by these channels and by distress-risk measurement (Campbell et al., 2008; Beneish, 1999).

Forecast evaluation and machine learning in finance. We follow the evaluation standards of Diebold and Mariano (1995), Giacomini and White (2006), Hansen (2005) and Hansen et al. (2011), with multiple-testing control (Holm, 1979; Benjamini and Hochberg, 1995), cluster-robust inference (Petersen, 2009; Cameron et al., 2011), block bootstrap (Künsch, 1989; Politis and Romano, 1994) and purged, embargoed time splits (López de Prado, 2018). Equal-budget comparisons of information sets are inspired by the asset-pricing machine-learning literature (Gu et al., 2020; Freyberger et al., 2020; Kozak et al., 2020; Chen et al., 2024; Avramov et al., 2023) and its warnings about false discoveries (Harvey et al., 2016; Chordia et al., 2020). Gradient-boosted trees (Ke et al., 2017) provide the nonlinear benchmark. The companion UQO working paper on intraday shock propagation (Boucher, 2026) uses the same vendor data for a different question; no result is reused here.

4. Data

4.1. Source and license

All market and accounting data come from FinData OS, a private lakehouse of FirstRate Data operated by the author and accessed through its versioned API (`findata.spboucher.ai/api/v1`). The lake holds 6.7 billion raw 1-minute bars for 16,195 U.S. stock symbols (2000–2026, including 8,069 delisted symbols), quarterly fundamentals for 5,451 entities (67 metrics, filing dates 2012–2026) and 373,326 earnings-release events. Bars are timestamped at the bar open in UTC and cover 04:00–20:00 Eastern; zero-volume minutes are absent at the source and are never imputed. FirstRate Data is licensed for private research use; no raw record is redistributed, and the replication package contains code, hashed extraction manifests (dataset identifiers and versions, row counts, SHA-256 digests) and derived aggregates only. Public augmentations are the SEC's `company_tickers.json` and `submissions` files (CIK, SIC industry), and the Kenneth French daily factor library.

4.2. Universe and survivorship

The accounting panel contains the 5,451 entities for which the vendor supplies quarterly statements; it has *no delisted entity*, because the vendor ingestion covered tickers active on 1 November 2026. The price side is survivorship-free (delisted names are present in the bar lake), but the firm-quarter sample is conditioned on survival to the end of the sample. This is a structural caveat: firms that deteriorated and delisted are absent, so adverse base rates are understated and the most vulnerable firms are under-represented. We state it in the abstract and treat it as the first limitation. Table 1 reports the construction of the universe: 4,568 firms with at least eight quarterly periods, at least 250 sessions of 1-minute bars after 2012 and a non-governmental SIC division; 5,073 are matched

to a CIK. Identifiers are vendor tickers; the vendor merges ticker changes under the latest ticker, and we audit reuse with the SEC's former-name history.

4.3. Reveal dates and point-in-time discipline

Earnings-release dates are matched to fiscal quarters as the first release in $[E_q, F_q + 10 \text{ days}]$; 88.5% of firm-quarters are matched with a filing lag of at most 30 days (median lag between release and filing: 1 day). Release timing is known (before-open or after-close) for about two thirds of events. Structured facts carry `filing_date`, `available_at` (= filing date +1 day, a documented policy rather than a historical vintage) and a revision identifier; because every fact was first ingested in November 2026 and carries revision 1, restatements that occurred before ingestion are not reconstructible. Inputs therefore use prior quarters only, outcomes use the vendor's current value of quarter q, and the stringent as-originally-reported specification is left for a future SEC-filing-level reconstruction (Section 9).

4.4. Sample

The quarter panel has 194,817 firm-quarters for 4,568 firms. After the admissibility and exclusion cascade of Table 1 (market window of at least 15 valid sessions, raw close of at least one dollar, available fundamentals, prediction of the next quarter, defined expectation, regular fiscal calendar), the event-time sample has 481,159 observations across the five horizons (3,796 firms, 102,522 firm-quarters) and the primary $h = 20$ cohort has 97,788 observations with cutoffs from 2013-09-26 to 2026-10-01; the fixed-calendar sample has 258,480 firm-months. The adverse outcome occurs in 30.1% of primary-cohort firm-quarters. Table 2 reports summary statistics and Figure 1 the coverage by year together with the adverse rate.

Table 1. Construction of the universe and of the primary sample

Step	Firms
fundamentals entities	5,451
>= 8 quarterly periods	5,018
+ 1m bars (>= 250 sessions)	4,568
+ bars after 2012-06	4,568
+ not public administration	4,568

Exclusion step	Remaining	Dropped
cutoff admissible	852,643	0
market window available	587,370	265,273
price ge min	577,241	10,129
fundamentals available	571,438	5,803
predicting next quarter	564,100	7,338
expectation defined	483,391	80,709
regular fiscal calendar	481,159	2,232

Notes. Left: firms. Right: event-time observations (five horizons pooled); steps are applied in the frozen order of `docs/PIT_PROTOCOL.md`.

Table 2. Summary statistics, primary cohort ($h = 20$)

Variable	N	Mean	SD	P10	Median	P90
Scaled operating-income surprise S_q	97,788	0.001628	0.4107	-0.02718	-0.0001251	0.02603
Adverse surprise ($S_q \leq -0.5\%$)	97,788	0.3015	0.4589	0	0	1
ROA (latest available)	97,788	0.002093	0.5508	-0.06	0.01	0.03
Debt / assets	97,788	0.2999	0.2796	0.01417	0.2717	0.5954
log market cap	97,788	21.88	2.25	19.48	21.91	24.51
20-day return	97,787	-0.006	0.1441	-0.1455	4.213e-05	0.1219
20-day daily volatility	97,787	0.02624	0.02662	0.01002	0.01981	0.04738
Amihud illiquidity (60d, $\times 10^6$)	97,740	0.007446	0.1977	6.559e-05	0.000824	0.01244
1-min realized variance (daily mean)	97,788	0.001269	0.002998	0.0001311	0.0004644	0.00288
1-min relative jump	97,788	0.1187	0.04981	0.06302	0.1096	0.1886
1-min negative semivariance share	97,788	0.5	0.01981	0.4749	0.5002	0.5247
1-min down-bar volume share	97,788	0.4331	0.04498	0.3735	0.4362	0.4872
1-min bars present (share)	97,788	0.7674	0.2107	0.4373	0.8242	0.9967
30-min realized variance	97,788	0.0008903	0.002253	0.000101	0.0003425	0.00196
5-min realized variance	97,788	0.001035	0.002424	0.000114	0.000394	0.002339
Days since last release	97,728	62.48	14.4	42	63	82

Notes. Intraday families are daily values averaged over the 20-session window ending at the cutoff; realized variances are sums of squared bar log-returns within the regular session. Raw, unwinsorised values.

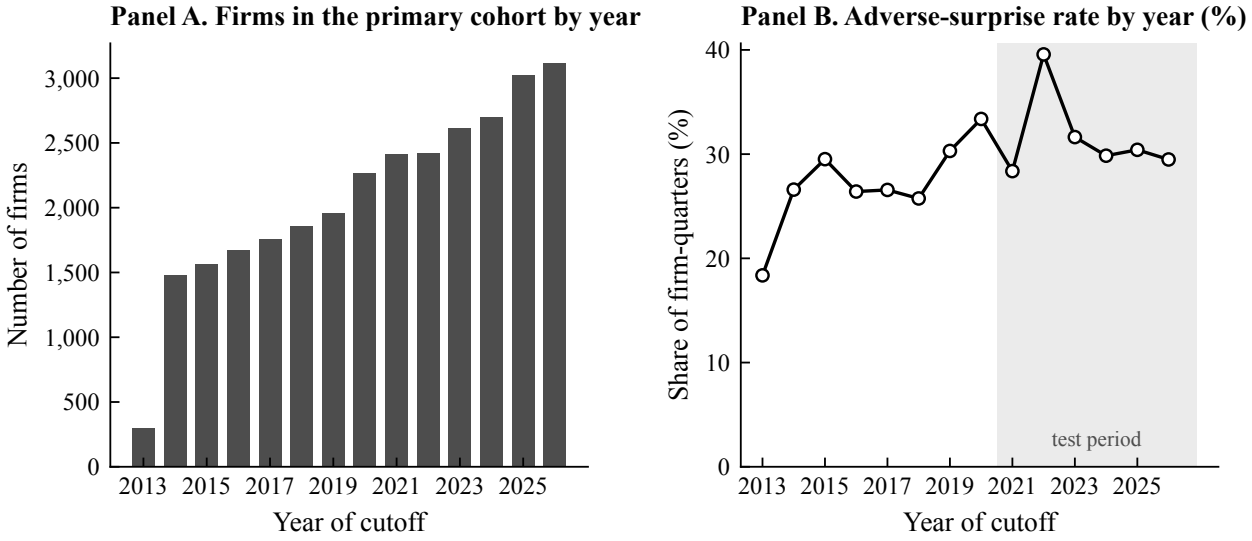


Figure 1. Sample coverage and adverse-outcome rate by year of cutoff ($h = 20$ cohort).

5. Research design

5.1. Clocks, cutoffs and the information set

Figure 2 summarises the event-time design. For firm i and fiscal quarter q we observe the period end $E_{i,q}$, the earnings-release date $A_{i,q}$ (the reveal date, with its before-open/after-close timing when known), the 10-Q/10-K filing date $F_{i,q}$ and the vendor availability instant $V_{i,q} = F_{i,q} + 1$ day. The prediction cutoff is $c = A_{i,q} - h$ trading days for $h \in \{40, 20, 10, 5, 1\}$, with $h = 20$ as the preregistered primary horizon. A cutoff is admissible only if $c < \min(A_{i,q}, F_{i,q})$ and $c > A_{i,q-1}$, so the feature window never contains the release of quarter q and never straddles two reveals. The information set is

$$\mathcal{I}(c) \;=\; \left\{\text{bars}_{i,t} :\; \text{session}(t) \le c\right\} \;\cup\; \left\{\text{facts}_{i,j} :\; V_{i,j} \le c\right\}, \qquad j^*(c) = \max\left\{j : V_{i,j} \le c\right\}, \quad (3)$$

regular-session bars up to the close of c and accounting facts available by c only; the latest such quarter is $j^*(c) = q - 1$ in the retained sample (we require that the forecaster predicts the quarter immediately following the latest available one). Twelve leakage guards (future facts, fiscal-end joins, post-cutoff bars, test-period transforms, daylight-saving errors, identifier drift, split artefacts, release-day bars, overlapping windows across splits, survivorship, ingestion dates, randomised placebos) are implemented as unit tests in the replication package; the local as-of join was verified live against the FinData API's `as_of` semantics on random entities and instants.

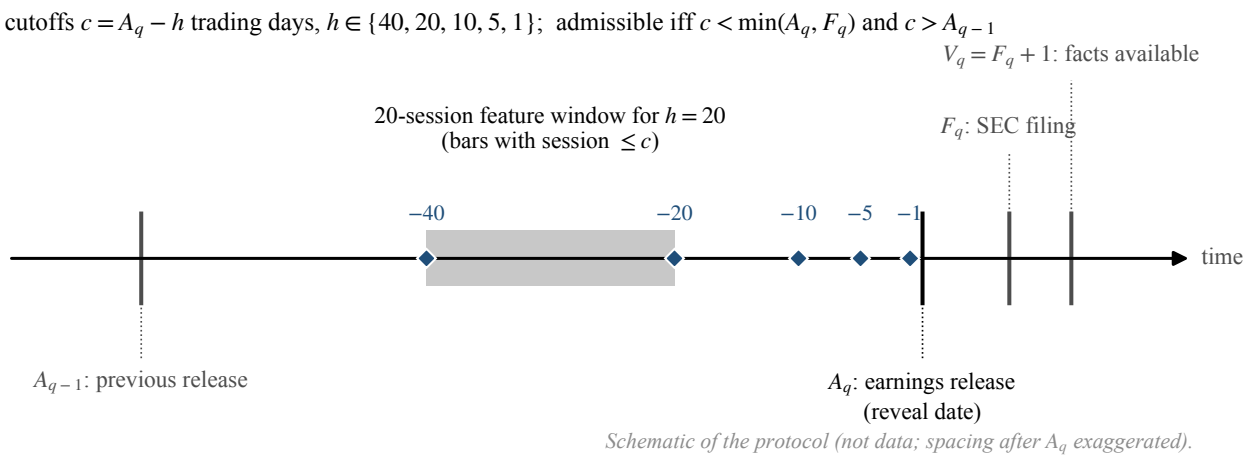


Figure 2. Event-time design and the no-leakage rule. Cutoffs $c = A_q - h$ precede the earnings release A_q; the 20-session feature window ends at c; structured facts of quarter q become available at $V_q = F_q + 1$ day. Schematic of the protocol, not data.

A fixed-calendar design (month-end cutoffs, target = the next release within 65 trading days) is used for the pricing tests and as a robustness check because it does not condition on knowledge of the release date.

5.2. Targets

The primary outcome is the scaled operating-income surprise $S_{i,q}$ of equation (1), with the seasonal expectation computed from quarters available at c and scaled by the latest available total assets. The binary adverse outcome is $D_{i,q} = \mathbb{1}\left[S_{i,q} \leq -0.005\right]$, i.e., a shortfall of at least half a percent of assets (roughly a fifth of a typical quarterly return on assets). Secondary outcomes are the net-income analogue, a year-on-year ROA decline of at least one point, a negative operating-cash-flow innovation of at least one percent of assets, a liquidity/coverage deterioration (current ratio down 10% or interest coverage falling below one), a symmetric placebo $|S_{i,q}| \geq 0.5\%$ that a pure volatility channel would predict, and an exploratory EPS miss against the vendor's estimate (available for 81% of releases; vintage unknown). The continuous surprise, winsorised at the training-period 1st/99th percentiles, is modelled by ridge regression. No target uses prices.

5.3. Features by sampling frequency

All market features are derived from the same raw 1-minute bars of the regular session (09:30 to the closing print), aggregated to 5-minute and 30-minute buckets and to daily bars. Each intraday *family* is computed per session on bars of the given length and averaged (or summed) over the 20-session window ending at c. With r_t the within-session log return of bar t and n the number of bars present,

the variance families are

$$\mathrm{RV} = \sum_t r_t^2, \qquad \mathrm{BPV} = \frac{\pi}{2} \sum_t |r_t|\,|r_{t-1}|, \qquad \mathrm{RJ} = \frac{\max(\mathrm{RV} - \mathrm{BPV}, 0)}{\mathrm{RV}}, \qquad \mathrm{RS}^{\pm} = \sum_t r_t^2 \mathbb{1}[r_t \gtrless 0], \tag{4}$$

together with the negative-semivariance share $\mathrm{RS}^-/\mathrm{RV}$, the signed-jump variation $(\mathrm{RS}^+ - \mathrm{RS}^-)/\mathrm{RV}$, realized skewness $\sqrt{n}\sum_t r_t^3/\mathrm{RV}^{3/2}$ and kurtosis $n\sum_t r_t^4/\mathrm{RV}^2$, the concentration $\max_t r_t^2/\mathrm{RV}$, and the first-order autocorrelation of r_t. The activity families are the intraday drift $\log(C/O)$, the opening and closing 30-minute volume shares, the volume Gini across bars, the down-bar volume share $\sum_t v_t \mathbb{1}[r_t < 0]\,/\sum_t v_t$, and the share of expected bars present. The co-movement families are the realized beta to SPY, $\hat{\beta} = \sum_t r_t r_t^{\mathrm{SPY}}/\sum_t (r_t^{\mathrm{SPY}})^2$ with sums taken over the window, the share of RV accruing on market-down bars, the variance ratio $\overline{\mathrm{RV}}/\mathrm{Var}(r^{\mathrm{daily}})$, and the volatility-of-volatility (dispersion of daily RV within the window). The 30-minute, 5-minute and 1-minute sets therefore have the *same* 19 columns. Daily features (16 columns) are 20- and 60-day returns, volatilities, Parkinson range, skewness, kurtosis, overnight and intraday return components, volume ratios, Amihud illiquidity, zero-return share, dollar volume, maximum drawdown and bar activity. Point-in-time accounting features (21 columns) are margins, ROA, accruals, cash, leverage, current ratio, coverage, capex, R&D, issuance, revenue growth, size, the last available surprise, the trailing mean surprise, the trailing adverse rate and the age of the information. Bar statistics are descriptive; nothing is interpreted as order flow, spread or informed trading.

5.4. Baseline ladder and equal budgets

Table 3. Nested information sets

Set	Information set	Features
B1	PIT fundamentals + history + size	21
B2	B1 + daily OHLCV (20/60-day)	37
B3	B2 + 30-minute families	56
B4	B2 + 5-minute families	56
B5	B2 + 1-minute families	56
B5all	B2 + 30m + 5m + 1m	94

Table 3 nests the information sets. B0 is the training base rate (and a shrunk firm-history rate); B1–B5 are ℓ_2-penalised logistic regressions whose penalty is tuned on the validation period only; LB1–LB5 are gradient-boosted trees (Ke et al., 2017) with an identical random-search budget (24 trials, early stopping on the validation period) for every information set, so that the frequency comparison is a comparison of information sets, not of capacity. Winsorisation bounds, imputation

medians and standardisation constants are estimated on training rows only. A supplementary set (B5all) concatenates all three intraday resolutions. Neural and text models are not used.

5.5. Validation and inference

The chronological split trains on cutoffs in 2012–2018, tunes on 2019–2020 and evaluates on 2021–2026, with purging of training outcomes revealed after the next split's start minus a one-quarter embargo (López de Prado, 2018). Expanding-window annual refits from 2016 provide out-of-fold predictions for the panel regressions and the timing and pricing exercises; a 20% firm hold-out evaluates transfer to unseen firms. The primary metric is log loss; Brier score, AUC, PR-AUC and calibration are reported for every model. Paired differences (B5−B2, B5−B4, B5−B3) are assessed on the per-observation loss differential $d_{i,q} = \ell^{\mathrm{A}}_{i,q} - \ell^{\mathrm{B}}_{i,q}$ with a block bootstrap over calendar quarters of the reveal date (Künsch, 1989; Politis and Romano, 1994) and with the Giacomini–White statistic on quarterly means,

$$t_{\mathrm{GW}} = \frac{\bar{d}}{\sqrt{\hat{\sigma}^2_{\mathrm{NW}}/T}}, \qquad \bar{d} = \frac{1}{T}\sum_{\tau=1}^{T} \bar{d}_\tau, \tag{5}$$

where $\bar{d}_\tau$ is the mean differential in calendar quarter τ and $\hat{\sigma}^2_{\mathrm{NW}}$ its Newey–West long-run variance (Diebold and Mariano, 1995; Giacomini and White, 2006); Holm's correction is applied within each hypothesis family (Holm, 1979). Panel regressions

$$S_{i,q} = \alpha_i + \delta_{s(i)\times t(q)} + \beta\,\mathrm{LFDS}_{i,c} + \boldsymbol{\gamma}^{\mathsf{T}}\boldsymbol{X}_{i,c} + \varepsilon_{i,q} \tag{6}$$

use the out-of-fold B5 probability as the Latent Fundamental Deterioration Signal (LFDS), firm and sector-by-quarter fixed effects, and two-way clustering by firm and calendar quarter (Petersen, 2009; Cameron et al., 2011). Mechanism tests (H3) interact the *incremental* signal $\mathrm{LFDS}^{\mathrm{B5}} - \mathrm{LFDS}^{\mathrm{B2}}$ with ex-ante terciles of size, illiquidity and leverage formed within each calendar quarter.

6. Results

6.1. The baseline ladder

Table 4 reports the out-of-sample performance of every model on the untouched test period (cutoffs from January 2021 to October 2026, 53,667 firm-quarters at $h = 20$, base rate 30.1%). Two facts stand out before any frequency comparison. First, point-in-time accounting information carries most of the predictable variation: the logistic B1 model lowers the log loss from 0.6282 (base rate) to 0.5527 and reaches an AUC of 0.730; the boosted LB1 reaches 0.4937 and 0.797. Second, daily market data add a clear and statistically unambiguous increment: B2−B1 is -0.0060 (logistic) and

-0.0068 (LightGBM, 95% CI [-0.0095, -0.0042]), with AUCs of 0.738 and 0.803. This is the "prices lead earnings" result of Beaver et al. (1980) and Kothari and Sloan (1992) reproduced on a strictly point-in-time panel with a 20-day horizon.

Table 4. Headline out-of-sample forecasting performance, primary target, $h = 20$, test period 2021–2026

Model	N	Log loss	Brier	AUC	PR-AUC	Cal. slope	Cal. int.
B0	53,667	0.6282	0.2178	0.500	0.315	0.39	-0.39
B0b	53,667	0.6137	0.2126	0.596	0.371	0.75	-0.00
B1	53,667	0.5527	0.1854	0.730	0.565	0.95	0.10
B2	53,667	0.5467	0.1828	0.738	0.572	0.86	-0.13
B3	53,667	0.5464	0.1830	0.739	0.571	0.87	-0.04
B4	53,667	0.5466	0.1831	0.739	0.571	0.87	-0.02
B5	53,667	0.5481	0.1838	0.738	0.569	0.86	0.00
B5all	53,667	0.5480	0.1837	0.738	0.569	0.86	0.00
LB1	53,667	0.4937	0.1634	0.797	0.655	1.07	0.13
LB2	53,667	0.4869	0.1606	0.803	0.665	1.04	-0.02
LB3	53,667	0.4886	0.1613	0.801	0.662	1.03	0.04
LB4	53,667	0.4890	0.1614	0.801	0.661	0.99	0.02
LB5	53,667	0.4885	0.1612	0.801	0.662	1.04	0.01
LB5all	53,667	0.4899	0.1618	0.800	0.660	1.04	0.08

Notes. B0: training base rate; B0b: shrunk firm-history rate; B1–B5: ℓ_2 logistic regressions; LB1–LB5: LightGBM with an identical 24-trial budget per information set; "all": B2 plus the three intraday resolutions. Calibration slope and intercept from a logistic recalibration of the predicted logit. Chronological split, purged with a one-quarter embargo; all transforms fitted on training rows only.

6.2. H1 and H2: the marginal value of sampling frequency

Table 5 and Figure 3 contain the paper's central evidence. Adding the 1-minute families to the daily-and-accounting benchmark changes the test log loss by +0.0013 in the logistic family (95% block-bootstrap CI [-0.0013, +0.0041], Giacomini–White $p = 0.518$) and by +0.0016 in the boosted family (CI [+0.0010, +0.0023], $p = 0.001$). Positive differences are *deteriorations*. The 30-minute and 5-minute sets behave the same way (+0.0017 and +0.0021 for LightGBM), and concatenating all three resolutions is worse still (+0.0030). H1 is therefore rejected, and with it H2: there is no gain to order across frequencies. The confidence intervals are informative rather than merely non-significant. The pre-registered minimum detectable difference at 80% power, computed from the validation-period variance of the quarterly loss differential and the 7 validation quarters, is 0.0025 log-loss units; the realised intervals around the 1-minute increment are narrower than this for the boosted family (which excludes any gain at all) and of comparable width for the logistic family (which excludes gains larger than about one fifth of the increment that daily data deliver over

accounting data alone). In words: whatever one-minute bars know about the forthcoming quarter, the daily close already knows.

Table 5. Paired out-of-sample contrasts between information sets, primary target, $h = 20$

Family	A	B	N	Δ log loss	95% CI	GW p	Holm	Δ Brier
logistic	B5	B2	53,667	+0.0013	[-0.0013, +0.0041]	0.518	1.000	+0.00102
logistic	B4	B2	53,667	-0.0002	[-0.0021, +0.0018]	0.870	1.000	+0.00034
logistic	B3	B2	53,667	-0.0003	[-0.0020, +0.0013]	0.800	1.000	+0.00023
logistic	B5	B4	53,667	+0.0015	[+0.0006, +0.0025]	0.023**	0.045	+0.00069
logistic	B5	B3	53,667	+0.0017	[+0.0005, +0.0030]	0.097*	0.097	+0.00079
logistic	B5all	B2	53,667	+0.0012	[-0.0014, +0.0040]	0.495		+0.00094
logistic	B5all	B5	53,667	-0.0001	[-0.0006, +0.0004]	0.899		-0.00008
logistic	B2	B1	53,667	-0.0060	[-0.0101, -0.0023]	0.085*		-0.00265
logistic	B1	B0	53,667	-0.0754	[-0.0846, -0.0675]	0.000***		-0.03243
lightgbm	LB5	LB2	53,667	+0.0016	[+0.0010, +0.0023]	0.001***	0.002	+0.00057
lightgbm	LB4	LB2	53,667	+0.0021	[+0.0007, +0.0035]	0.030**	0.049	+0.00082
lightgbm	LB3	LB2	53,667	+0.0017	[+0.0003, +0.0031]	0.025**	0.049	+0.00074
lightgbm	LB5	LB4	53,667	-0.0005	[-0.0019, +0.0006]	0.471	0.943	-0.00025
lightgbm	LB5	LB3	53,667	-0.0001	[-0.0016, +0.0011]	0.496	0.943	-0.00017
lightgbm	LB5all	LB2	53,667	+0.0030	[+0.0016, +0.0045]	0.005***		+0.00122
lightgbm	LB5all	LB5	53,667	+0.0015	[+0.0002, +0.0028]	0.075*		+0.00065
lightgbm	LB2	LB1	53,667	-0.0068	[-0.0095, -0.0042]	0.001***		-0.00278
lightgbm	LB1	LB0	53,667	-0.1345	[-0.1471, -0.1240]	0.000***		-0.05448

Notes. Δ log loss = $\mathcal{L}(A) - \mathcal{L}(B)$ averaged over the test observations (negative favours A). 95% CI from 1,000 draws of a block bootstrap over calendar quarters of the reveal date; GW p from the Giacomini–White test on quarterly mean differentials with Newey–West variance; Holm: adjusted within the "frequency vs daily" family (B3, B4, B5 vs B2) and the "fine vs coarse" family (B5 vs B4, B5 vs B3). *, **, ***: 10, 5, 1%.

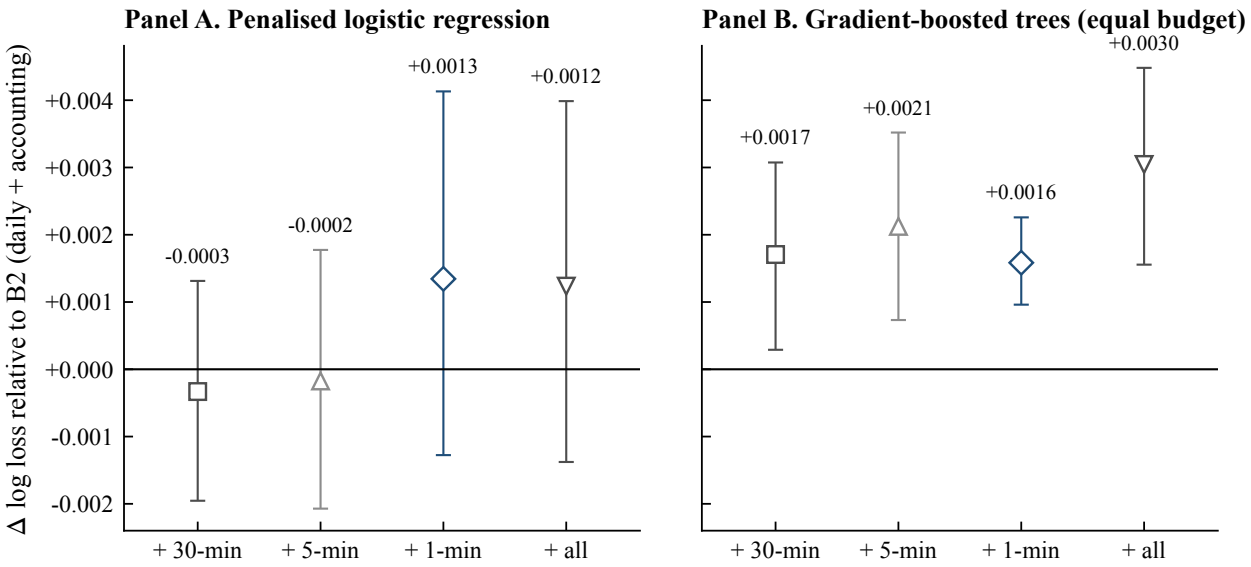


Figure 3. Change in test log loss relative to the daily-and-accounting benchmark B2 when 30-minute, 5-minute, 1-minute or all intraday families are added (left: logistic; right: LightGBM). Points are means, bars are 95% block-bootstrap intervals over calendar quarters; negative values would indicate a gain.

The continuous target tells the same story: ridge forecasts of the winsorised surprise reach an out-of-sample R^2 of 0.2121 with daily data and 0.2117 with 1-minute data (Appendix Table C), and the paired difference in squared error is indistinguishable from zero. Calibration (Figure 4) is close to the diagonal for the boosted models, with slopes near one, so the comparison is not an artefact of miscalibrated probabilities.

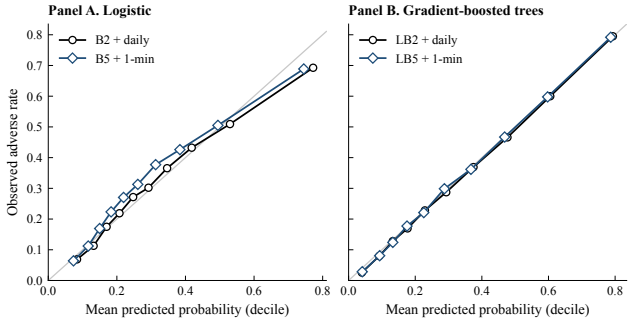


Figure 4. Reliability by decile of predicted probability, test period.

6.3. H4: timing

If one-minute activity revealed information that accumulates as the release approaches, the increment should grow from $h = 40$ to $h = 1$. Table 6 and Figure 5 show no such gradient: the AUCs of B2 and B5 are within a few thousandths of each other at every horizon, the LightGBM increment changes sign across horizons, and the rank correlation between the gain and proximity to the release is -0.50

($p = 0.39$) for the boosted family and -0.49 ($p = 0.33$) for the logistic family. The fixed-calendar design (month-end cutoffs, labelled $h = 0$ in the table), which does not condition on the release schedule, gives the same answer. H4 is rejected.

Table 6. The 1-minute increment by horizon (trading days before the release); $h = 0$ denotes the fixed-calendar design

Family	h	N	AUC B2	AUC B5	Δ log loss	95% CI	GW p
lightgbm	40	49,340	0.801	0.801	+0.0000	[-0.0009, +0.0010]	0.535
lightgbm	20	53,667	0.804	0.804	-0.0012	[-0.0023, -0.0001]	0.134
lightgbm	10	53,751	0.805	0.803	+0.0006	[-0.0005, +0.0018]	0.240
lightgbm	5	53,575	0.803	0.803	-0.0004	[-0.0013, +0.0005]	0.786
lightgbm	1	53,669	0.804	0.803	+0.0009	[+0.0001, +0.0017]	0.063
logistic	40	49,340	0.740	0.740	+0.0025	[+0.0005, +0.0048]	0.195
logistic	20	53,667	0.738	0.738	+0.0013	[-0.0013, +0.0041]	0.518
logistic	10	53,751	0.737	0.738	+0.0029	[-0.0004, +0.0066]	0.295
logistic	5	53,575	0.739	0.738	+0.0026	[-0.0007, +0.0065]	0.299
logistic	1	53,669	0.738	0.738	+0.0020	[-0.0008, +0.0050]	0.357
logistic	0	141,477	0.737	0.736	+0.0032	[+0.0004, +0.0061]	0.236

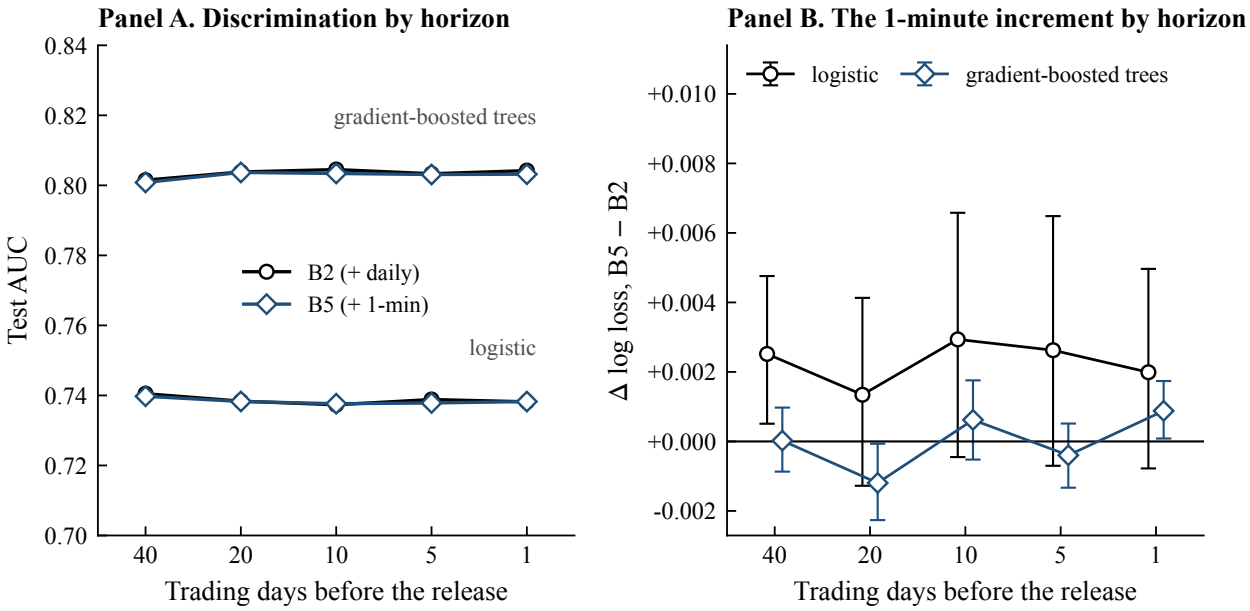


Figure 5. Test AUC of B2 and B5 (left) and the paired log-loss difference B5−B2 with 95% intervals (right) as the cutoff approaches the earnings release. All firm-quarters, no selection.

6.4. Alternative accounting targets and transfer to unseen firms

Table 7 repeats the B5-versus-B2 comparison for the secondary outcomes. For net income, ROA decline and liquidity deterioration the 1-minute increment is within ±0.002 log-loss units of zero in both families; for the operating-cash-flow innovation it is a deterioration; and for the symmetric placebo target $|S_q| \geq 0.5\%$, which a pure volatility channel would predict, the increment is also zero, so the null is not the result of a signed target masking an absolute-magnitude signal. Held-out firms (20% never used in training) behave like seen firms (Appendix Table C): the 1-minute increment is +0.0018 on unseen firms versus +0.0019 on seen firms, both indistinguishable from zero.

Table 7. The 1-minute increment for alternative accounting targets, $h = 20$

Target	Family	N	Base rate	AUC B2	AUC B5	Δ log loss	95% CI	GW p
y adverse netinc	logistic	53,663	0.331	0.708	0.709	-0.0005	[-0.0014, +0.0004]	0.828
y adverse netinc	lightgbm	53,663	0.331	0.758	0.757	+0.0013	[+0.0000, +0.0023]	0.081
y roa decline	logistic	53,662	0.277	0.688	0.688	-0.0014	[-0.0030, -0.0001]	0.334
y roa decline	lightgbm	53,662	0.277	0.722	0.724	+0.0002	[-0.0010, +0.0015]	0.484
y ocf negative	logistic	52,999	0.243	0.704	0.702	+0.0078	[+0.0035, +0.0131]	0.016
y ocf negative	lightgbm	52,999	0.243	0.745	0.745	+0.0022	[+0.0005, +0.0039]	0.019
y liquidity deterioration	logistic	40,376	0.447	0.642	0.639	+0.0006	[-0.0007, +0.0017]	0.335
y liquidity deterioration	lightgbm	40,376	0.447	0.694	0.695	-0.0005	[-0.0015, +0.0005]	0.593
y abs surprise	logistic	53,667	0.628	0.783	0.785	+0.0015	[-0.0003, +0.0037]	0.067
y abs surprise	lightgbm	53,667	0.628	0.836	0.835	+0.0008	[+0.0002, +0.0013]	0.008
y adverse opinc	logistic	53,667	0.315	0.738	0.738	+0.0013	[-0.0013, +0.0041]	0.518
y adverse opinc	lightgbm	53,667	0.315	0.803	0.801	+0.0016	[+0.0010, +0.0023]	0.001

6.5. Rolling performance

Figure 6 reports annual out-of-fold AUCs from expanding-window refits (2016–2026) and the yearly B5−B2 differential. The level of predictability is stable (AUC around 0.80 for the boosted models in every year, including 2020 and 2022), and the 1-minute differential oscillates around zero with no trend, which is also what the COVID and post-COVID subsamples of Section 8 show.

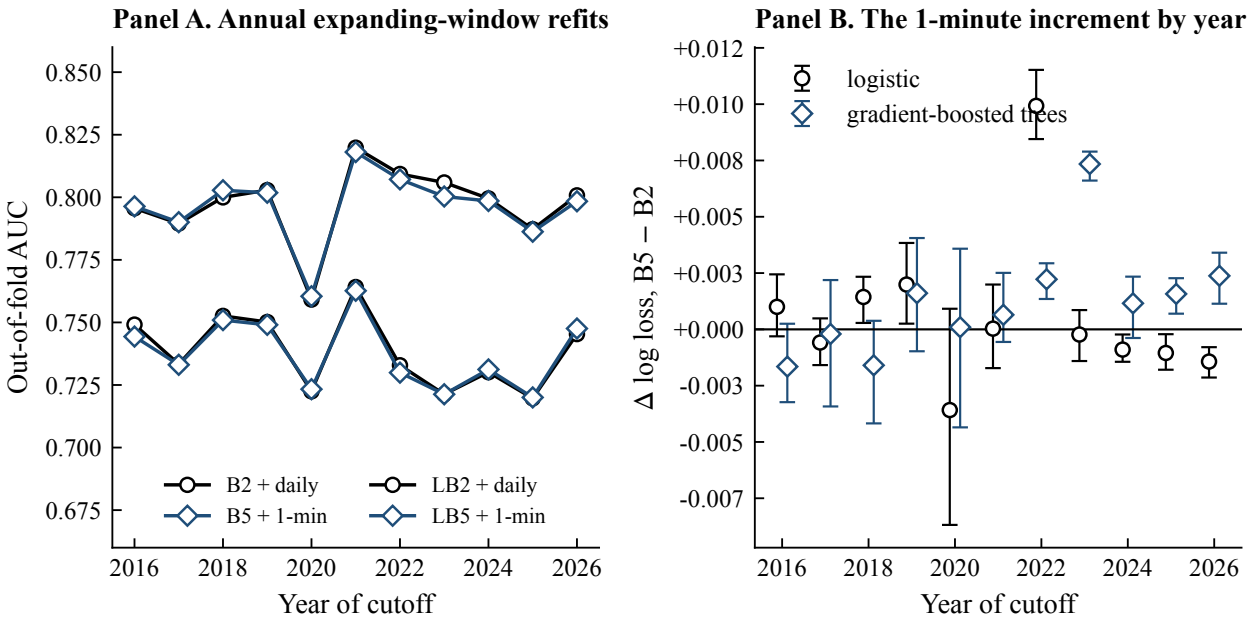


Figure 6. Expanding-window annual refits: out-of-fold AUC by year (left) and yearly B5−B2 log-loss difference with 95% intervals (right).

7. Mechanisms and pricing

7.1. Where would one-minute information live? (H3)

The pre-registered heterogeneity hypothesis holds that if fine-grained activity carries incremental information, it should do so where information asymmetry and processing frictions are largest: small, illiquid and highly levered firms. Table 8 and Figure 7 split the test period along these ex-ante terciles (cut-points from training rows), by profitability, release timing, sector, bar density, and by whether the 20-day window straddles the previous release. In no subgroup does the 1-minute set improve on the daily benchmark. The ordering is nevertheless informative: the deterioration is smallest and statistically indistinguishable from zero for small, illiquid and loss-making firms (*e.g.*, small firms: within ±0.002 of zero), and largest for large, liquid and profitable firms, where the extra features add noise to an already accurate forecast. The sign of the H3 ordering is therefore the one predicted by frictions, but the level is zero: the groups in which intraday activity would be most informative are also those in which it is not informative *enough* to help.

Table 8. The 1-minute increment by pre-specified subgroup (LightGBM, test period, $h = 20$)

Subgroup	N	Base rate	AUC LB2	AUC LB5	Δ log loss	95% CI
all	53,667	0.315	0.803	0.801	+0.0016	[+0.0010, +0.0023]
size small	19,065	0.388	0.780	0.780	+0.0007	[-0.0006, +0.0019]
size large	17,796	0.247	0.811	0.807	+0.0026	[+0.0012, +0.0042]
illiq high	21,817	0.372	0.788	0.787	+0.0007	[-0.0006, +0.0019]
illiq low	14,864	0.251	0.807	0.805	+0.0022	[+0.0006, +0.0038]
lev high	21,545	0.286	0.800	0.798	+0.0021	[+0.0010, +0.0031]
lev low	16,726	0.342	0.817	0.815	+0.0015	[+0.0002, +0.0027]
loss firms	17,786	0.376	0.765	0.764	+0.0009	[-0.0003, +0.0022]
profit firms	35,881	0.284	0.817	0.815	+0.0019	[+0.0009, +0.0029]
ex covid	53,667	0.315	0.803	0.801	+0.0016	[+0.0010, +0.0023]
after close release	24,829	0.313	0.804	0.802	+0.0013	[+0.0004, +0.0021]
before open release	22,615	0.316	0.803	0.801	+0.0023	[+0.0013, +0.0034]
sector finance	9,737	0.186	0.864	0.862	+0.0017	[+0.0001, +0.0036]
sector nonfinance	43,930	0.343	0.783	0.782	+0.0016	[+0.0008, +0.0024]
window clean	51,290	0.311	0.804	0.803	+0.0017	[+0.0011, +0.0024]
window straddles prev release	2,377	0.390	0.764	0.764	-0.0010	[-0.0040, +0.0019]
dense bars	10,866	0.316	0.796	0.795	+0.0018	[-0.0005, +0.0042]
sparse bars	42,801	0.315	0.805	0.803	+0.0015	[+0.0007, +0.0023]

Notes. Terciles of log market capitalisation, Amihud illiquidity and debt-to-assets are formed on training rows. "Window straddles previous release": the 20-session feature window contains A_{q-1}. "Dense bars": at least 98% of expected 1-minute bars present. 95% block-bootstrap CIs over calendar quarters; 2020 is in the validation period and therefore absent from the test subgroups.

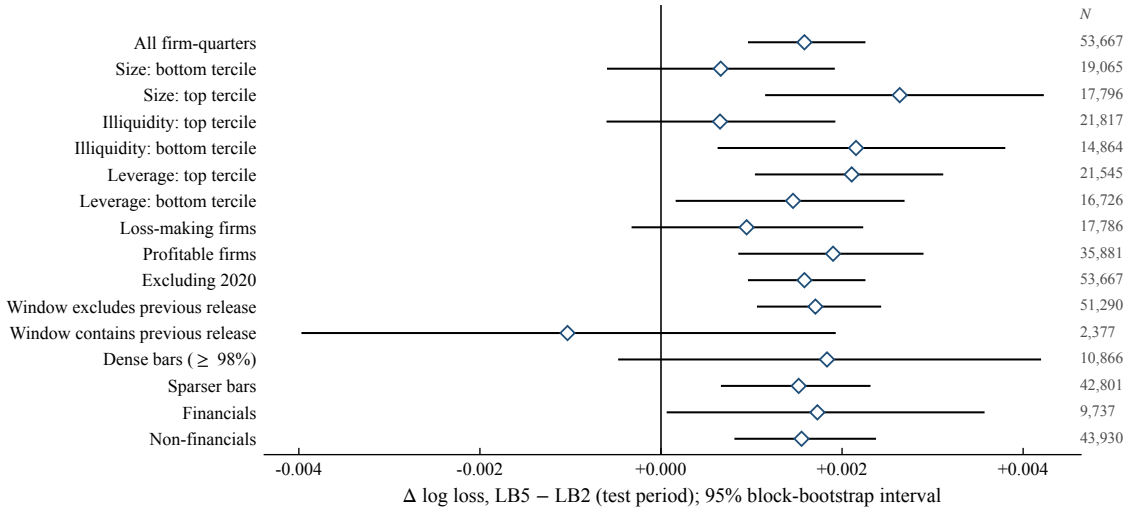


Figure 7. Change in test log loss from adding 1-minute families to the daily-and-accounting benchmark, by subgroup (LightGBM). Negative values would indicate a gain.

The panel regressions of Table 9 use the out-of-fold Latent Fundamental Deterioration Signal from the expanding-window refits. The daily-benchmark signal LFDS^{B2} is a powerful predictor of the realised surprise with firm and sector-by-quarter fixed effects ($\beta = -0.0765$, $t = -32.22$, two-way clustered): a one-point increase in the predicted probability of an adverse quarter is associated with a surprise lower by 7.6% of assets. The *incremental* 1-minute component $\text{LFDS}^{\text{B5}} - \text{LFDS}^{\text{B2}}$ has a coefficient of -0.0129 ($t = -1.76$), one sixth of the benchmark's and marginal at conventional levels. Its interactions with ex-ante terciles point in the direction *opposite* to H3: the incremental component is less informative for the smallest firms ($t = 2.29$ on the bottom-tercile interaction) and for the least levered firms ($t = 2.93$), and does not vary with illiquidity. Whatever small amount of information the 1-minute features add is concentrated in large, levered firms, consistent with fine-grained activity reflecting the trading of well-informed participants in liquid names rather than the slow processing of information in neglected ones; in the forecasting exercise this amount is not large enough to survive the noise the features bring.

Table 9. Panel regressions of the realised surprise on out-of-fold deterioration signals

Specification	Regressor	Coef.	SE	N	Firms	Quarters
S on lfds5	lfds5	-0.0763***	(0.0024)	84,759	3,598	44
S on lfds2 and dlfds	lfds2	-0.0765***	(0.0024)	84,759	3,598	44
S on lfds2 and dlfds	d lfds	-0.0129*	(0.0073)	84,759	3,598	44
y on lfds5	lfds5	0.9681***	(0.0180)	84,759	3,598	44
dlfds x size	lfds2	-0.0765***	(0.0024)	84,759	3,598	44
dlfds x size	d lfds	-0.0228***	(0.0082)	84,759	3,598	44
dlfds x size	dlfds x size0	0.0327**	(0.0143)	84,759	3,598	44
dlfds x size	dlfds x size2	-0.0089	(0.0096)	84,759	3,598	44
dlfds x illiq	lfds2	-0.0765***	(0.0024)	84,759	3,598	44
dlfds x illiq	d lfds	-0.0148	(0.0090)	84,759	3,598	44
dlfds x illiq	dlfds x illiq0	-0.0118	(0.0113)	84,759	3,598	44
dlfds x illiq	dlfds x illiq2	0.0152	(0.0157)	84,759	3,598	44
dlfds x lev	lfds2	-0.0765***	(0.0024)	84,759	3,598	44
dlfds x lev	d lfds	-0.0239***	(0.0083)	84,759	3,598	44
dlfds x lev	dlfds x lev0	0.0300***	(0.0103)	84,759	3,598	44
dlfds x lev	dlfds x lev2	0.0011	(0.0092)	84,759	3,598	44

Notes. Dependent variable: $S_{i,q}$ winsorised at 1/99% (specifications "S_on"), or the binary adverse outcome ("y_on"). lfds2, lfds5: out-of-fold predicted probabilities from LB2 and LB5 (expanding-window annual refits, 2016–2026); d_lfds = lfds5 − lfds2; interactions with within-quarter terciles (0 = bottom, 2 = top). Controls: ROA, leverage, log market cap, 20-day return and volatility, last and mean surprise. Firm and sector×calendar-quarter fixed effects; standard errors two-way clustered by firm and calendar quarter. *, **, ***: 10, 5, 1%.

7.2. Is the signal in prices? (H5)

If the deterioration probability were news to the market, firms with a high ex-ante probability should experience negative abnormal returns between the cutoff and the day after the release. Appendix Table C shows that the market-adjusted return from $c+1$ to A_q+1 is unrelated to LFDS^{B5} ($\beta = +0.0001$, $t = 0.01$) and to the daily benchmark signal: by the cutoff, twenty trading days before the release, the information that the models extract is already in prices. The incremental 1-minute component carries a positive coefficient (+0.0605, $t = 2.90$), i.e., firms whose 1-minute features push the probability *up* relative to daily data earn slightly *higher* announcement-window returns, the opposite of a deterioration signal; we read this as a partial reversal of recent intraday selling pressure rather than as information, and note that it is one of several H5 tests.

The secondary portfolio exercise confirms the absence of economic value (Table 10, Figure 8). Month-end decile sorts on the out-of-fold calendar-design scores, with positions taken at the close of the following session and held one month, produce long-minus-short (low minus high deterioration probability) FF5+momentum alphas of -0.59% per month ($t = -1.85$) gross and -0.62% net of 20 basis points per side for the 1-minute scores, and -0.49% ($t = -1.64$) for the daily scores; mean returns *rise* with the deterioration decile, as expected if high-probability firms are small, volatile and compensated for risk. A forecastable accounting deterioration is therefore not a tradable anomaly in this sample, at any sampling frequency.

Table 10. Economic value: monthly decile portfolios on the ex-ante deterioration score, 2016–2026

Score	Variant	Mean L/S (%/mo)	FF5+MOM α (%/mo)	t (HAC)	Sharpe (ann.)	Turnover	Months
LB2	gross	-0.51	-0.49	-1.64	-0.36	0.59	129
LB2	net20bps	-0.53	-0.51	-1.73	-0.38	0.59	129
LB5	gross	-0.65	-0.59	-1.85	-0.44	0.60	129
LB5	net20bps	-0.68	-0.62	-1.93	-0.46	0.60	129

Notes. Long decile 1 (lowest score), short decile 10; equal weights; entry at the close of the session after the month-end cutoff, exit at the close of the session after the next cutoff; turnover is one-way; net returns deduct 20 bps per side. Alphas from a regression on FF5 and momentum factors compounded over the same holding windows, HAC standard errors.

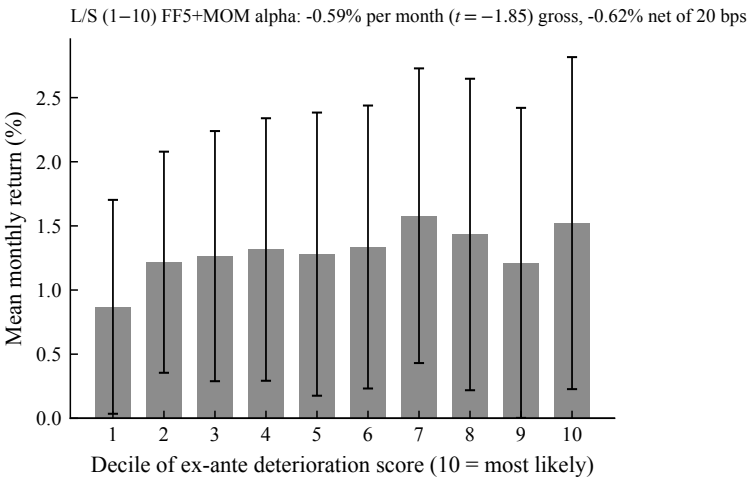


Figure 8. Mean monthly return by decile of the 1-minute deterioration score.

8. Robustness and negative controls

Shuffled labels. Permuting the adverse-outcome label within calendar quarter of the cutoff (training and test rows, three draws) removes all predictability: the test AUC of the 1-minute model falls to 0.503 and the B5−B2 differential is -0.0004 on average (Table 11). The pipeline does not manufacture differences between information sets from the features alone, and the residual differential on shuffled labels gives the noise floor against which the primary estimates should be read: the primary B5−B2 deteriorations (Table 5) are of the same order of magnitude as this floor, which is one more reason to describe the 1-minute increment as zero rather than as a harm.

Table 11. Negative control: labels shuffled within calendar quarter

Placebo	Draw	N	AUC LB5	AUC LB2	Δ log loss	95% CI
shuffled labels	0	53,667	0.502	0.497	-0.0003	[-0.0006, -0.0000]
shuffled labels	1	53,667	0.502	0.509	-0.0005	[-0.0013, +0.0002]
shuffled labels	2	53,667	0.506	0.507	-0.0005	[-0.0016, +0.0004]

Symmetric target. A volatility channel would make 1-minute features informative about the *magnitude* of the surprise even if not about its sign. The symmetric target $|S_q| \geq 0.5\%$ in Table 7 shows the same null, so the result is not an artefact of a signed target.

Window contamination and sparse bars. Excluding the 2,377 test observations whose 20-session window contains the previous release, or restricting to firm-days with at least 98% of expected bars, leaves the differential unchanged (Table 8). Sessions with fewer than 100 regular-session bars are

already excluded from the window, and the share of expected bars present is itself a feature, so illiquid names do not generate mechanical signals.

Release timing and sector. The differential is the same for after-close and before-open releases, and for financial and non-financial firms (the latter forecast much better, AUC 0.86, because their outcomes are smoother).

Model family and tuning. The null holds for penalised logistic regression and for gradient-boosted trees with identical tuning budgets; the sign of the LightGBM differential flips between the tuned configuration (+0.0016) and the fixed configuration used in the horizon ladder (Table 6, $h = 20$), which is what one expects from a true zero.

Firm hold-out and time. The increment is zero on firms never seen in training (+0.0018, CI [-0.0005, +0.0042]) and in every refit year (Figure 6); the only year with a visibly positive differential (2023 for the boosted family) is a deterioration, not a gain.

What we could not test. Delisted firms are absent from the accounting panel, so the stress test that would matter most, forecasting the deterioration of firms that subsequently failed, cannot be run on this vendor's data; the delisted-fundamentals backfill is documented as the first item of future work. Original-as-reported accounting values are not available, so the stringent as-filed specification is deferred to a SEC-filing-level reconstruction. Both limitations affect the level of predictability, not the comparison across sampling frequencies, which is the object of the paper.

9. Conclusion

At a 20-trading-day horizon, on a 2013–2026 panel of U.S. listed firms with a strictly point-in-time information set, one-minute OHLCV features change the out-of-sample log loss of forecasts of an adverse operating-income surprise by +0.0013 [-0.0013, +0.0041] (logistic) and +0.0016 [+0.0010, +0.0023] (gradient-boosted trees) relative to a daily-market-and-accounting benchmark whose own increment over accounting data is -0.0068. The gain from one-minute sampling is zero within intervals that are narrower than the pre-registered minimum detectable difference for the boosted family and comparable to it for the logistic family; it is zero at every horizon from forty days to one day, for every accounting target, in every subgroup and refit year, and on unseen firms. The small incremental content that the one-minute signal carries is inconsistent with the frictions mechanism (it lives in large, levered firms), is already impounded in prices by the cutoff, and has no economic value after factor adjustment and costs. The bracketed fields of the question we set out to answer are therefore filled with a measured null.

Two limits bound the claim. The accounting panel contains only firms listed at the end of the sample, so the result speaks to survivors; a delisted-fundamentals backfill, which the vendor appears to support, is the first extension. Point-in-time availability is the vendor's filing-date-plus-one-day policy on current (possibly restated) values; a filing-level reconstruction from SEC data would allow an as-originally-reported specification. Neither limit affects the cross-frequency comparison, which is the object of the paper, but both affect the level of predictability and the base rate of severe deterioration.

Three directions follow. First, the daily increment is large and its composition (returns, volatilities, volume ratios, illiquidity) deserves the same frequency-controlled scrutiny at the weekly and monthly scale: where between the quarter and the day does market anticipation of fundamentals actually form? Second, quote data would allow the mechanisms that bars cannot see, order imbalance and spreads, to be tested with the same point-in-time design; the present result predicts that they too would be summarised by the daily close at a 20-day horizon, a prediction worth falsifying. Third, the pipeline, the frozen design and the leakage battery are reusable for any forecasting question that joins market and accounting clocks; we release them with the hashed manifests needed to reproduce every exhibit from the pinned dataset versions.

References

Aboody, David, Reuven Lehavy, and Brett Trueman, "Limited attention and the earnings announcement returns of past stock market winners," *Review of Accounting Studies*, 2010, *15* (2), 317–344.

Andersen, Torben G. and Tim Bollerslev, "Answering the Skeptics: Yes, Standard Volatility Models do Provide Accurate Forecasts," *International Economic Review*, 1998, *39* (4), 885.

_ , _ , Francis X. Diebold, and Paul Labys, "Modeling and Forecasting Realized Volatility," *Econometrica*, 2003, *71* (2), 579–625.

Avramov, Doron, Si Cheng, and Lior Metzker, "Machine Learning vs. Economic Restrictions: Evidence from Stock Return Predictability," *Management Science*, 2023, *69* (5), 2587–2619.

Bajgrowicz, Pierre, Olivier Scaillet, and Adrien Treccani, "Jumps in High-Frequency Data: Spurious Detections, Dynamics, and News," *Management Science*, 2016, *62* (8), 2198–2217.

Ball, Ray and Philip Brown, "An Empirical Evaluation of Accounting Income Numbers," *Journal of Accounting Research*, 1968, *6* (2), 159.

Barndorff-Nielsen, O. E., "Power and Bipower Variation with Stochastic Volatility and Jumps," *Journal of Financial Econometrics*, 2004, *2* (1), 1–37.

Barndorff-Nielsen, Ole E., Peter Reinhard Hansen, Asger Lunde, and Neil Shephard, "Designing Realized Kernels to Measure the ex post Variation of Equity Prices in the Presence of Noise," *Econometrica*, 2008, *76* (6), 1481–1536.

_ , Silja Kinnebrock, and Neil Shephard, "Measuring Downside Risk: Realised Semivariance," in Tim Bollerslev, Jeffrey Russell, and Mark Watson, eds., *Volatility and Time Series Econometrics: Essays in Honor of Robert F. Engle*, Oxford University Press, 2010, pp. 117–136.

Beaver, William, Richard Lambert, and Dale Morse, "The information content of security prices," *Journal of Accounting and Economics*, 1980, *2* (1), 3–28.

Beneish, Messod D., "The Detection of Earnings Manipulation," *Financial Analysts Journal*, 1999, *55* (5), 24–36.

Benjamini, Yoav and Yosef Hochberg, "Controlling the False Discovery Rate: A Practical and Powerful Approach to Multiple Testing," *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 1995, *57* (1), 289–300.

Berkman, Henk, Valentin Dimitrov, Prem C. Jain, Paul D. Koch, and Sheri Tice, "Sell on the news: Differences of opinion, short-sales constraints, and returns around earnings announcements," *Journal of Financial Economics*, 2009, *92* (3), 376–399.

Bernard, Victor L. and Jacob K. Thomas, "Post-Earnings-Announcement Drift: Delayed Price Response or Risk Premium?," *Journal of Accounting Research*, 1989, *27*, 1.

_ **and** _ , "Evidence that stock prices do not fully reflect the implications of current earnings for future earnings," *Journal of Accounting and Economics*, 1990, *13* (4), 305–340.

Bollerslev, Tim, Andrew J. Patton, and Rogier Quaedvlieg, "Exploiting the errors: A simple approach for improved volatility forecasting," *Journal of Econometrics*, 2016, *192* (1), 1–18.

_ **, Benjamin Hood, John Huss, and Lasse Heje Pedersen**, "Risk Everywhere: Modeling and Managing Volatility," *The Review of Financial Studies*, 2018, *31* (7), 2729–2773.

_ **, Jia Li, and Yuan Xue**, "Volume, Volatility, and Public News Announcements," *The Review of Economic Studies*, 2018, *85* (4), 2005–2041.

_ **, Sophia Zhengzi Li, and Bingzhi Zhao**, "Good Volatility, Bad Volatility, and the Cross Section of Stock Returns," *Journal of Financial and Quantitative Analysis*, 2020, *55* (3), 751–781.

_ **,** _ **, and Viktor Todorov**, "Roughing up beta: Continuous versus discontinuous betas and the cross section of expected stock returns," *Journal of Financial Economics*, 2016, *120* (3), 464–490.

Boucher, Simon-Pierre, "Faster, Wider, More Fragile? The Evolution of Intraday Shock Propagation in U.S. Equity Markets, 1992–2026," *UQO Working Paper No. 15*, 2026.

Brogaard, Jonathan, Terrence Hendershott, and Ryan Riordan, "Price Discovery without Trading: Evidence from Limit Orders," *The Journal of Finance*, 2019, *74* (4), 1621–1658.

Cameron, A. Colin, Jonah B. Gelbach, and Douglas L. Miller, "Robust Inference With Multiway Clustering," *Journal of Business & Economic Statistics*, 2011, *29* (2), 238–249.

Campbell, JOHN Y., Jens Hilscher, and Jan Szilagyi, "In Search of Distress Risk," *The Journal of Finance*, 2008, *63* (6), 2899–2939.

Cao, Sean, Wei Jiang, Junbo Wang, and Baozhong Yang, "From Man vs. Machine to Man + Machine: The art and AI of stock analyses," *Journal of Financial Economics*, 2024, *160*, 103910.

Chen, Luyang, Markus Pelger, and Jason Zhu, "Deep Learning in Asset Pricing," *Management Science*, 2024, *70* (2), 714–750.

Chen, XI, Yang Ha (tony) Cho, Yiwei Dou, and Baruch Lev, "Predicting Future Earnings Changes Using Machine Learning and Detailed Financial Data," *Journal of Accounting Research*, 2022, *60* (2), 467–515.

Chordia, Tarun, Amit Goyal, and Alessio Saretto, "Anomalies and False Rejections," *The Review of Financial Studies*, 2020, *33* (5), 2134–2179.

Christophe, Stephen E., Michael G. Ferri, and James J. Angel, "Short-Selling Prior to Earnings Announcements," *The Journal of Finance*, 2004, *59* (4), 1845–1876.

Collins, Daniel W., S.P. Kothari, Jay Shanken, and Richard G. Sloan, "Lack of timeliness and noise as explanations for the low contemporaneuos return-earnings association," *Journal of Accounting and Economics*, 1994, *18* (3), 289–324.

de Prado, Marcos López, *Advances in Financial Machine Learning*, Hoboken, NJ: Wiley, 2018.

DellaVigna, Stefano and JOSHUA M. Pollet, "Investor Inattention and Friday Earnings Announcements," *The Journal of Finance*, 2009, *64* (2), 709–749.

Diebold, Francis X. and Roberto S. Mariano, "Comparing Predictive Accuracy," *Journal of Business & Economic Statistics*, 1995, *13* (3), 253–263.

Easley, David and Maureen O'Hara, "Price, trade size, and information in securities markets," *Journal of Financial Economics*, 1987, *19* (1), 69–90.

Easton, Peter D., Martin M. Kapons, Steven J. Monahan, Harm H. Schütt, and Eric H. Weisbrod, "Forecasting Earnings Using *k*-Nearest Neighbors," *The Accounting Review*, 2024, *99* (3), 115–140.

Fama, Eugene F., "Efficient Capital Markets: A Review of Theory and Empirical Work," *The Journal of Finance*, 1970, *25* (2), 383.

Fang, Lily and Joel Peress, "Media Coverage and the Cross-section of Stock Returns," *The Journal of Finance*, 2009, *64* (5), 2023–2052.

Foster, George, "Quarterly Accounting Data: Time-Series Properties and Predicative-Ability Results.," *The Accounting Review*, 1977, *52* (1), 1–21.

Freyberger, Joachim, Andreas Neuhierl, and Michael Weber, "Dissecting Characteristics Nonparametrically," *The Review of Financial Studies*, 2020, *33* (5), 2326–2377.

Giacomini, Raffaella and Halbert White, "Tests of Conditional Predictive Ability," *Econometrica*, 2006, *74* (6), 1545–1578.

Grossman, Sanford J. and Joseph E. Stiglitz, "On the Impossibility of Informationally Efficient Markets," *American Economic Review*, 1980, *70* (3), 393–408.

Grégoire, Vincent and Charles Martineau, "How is Earnings News Transmitted to Stock Prices?," *Journal of Accounting Research*, 2022, *60* (1), 261–297.

Gu, Shihao, Bryan Kelly, and Dacheng Xiu, "Empirical Asset Pricing via Machine Learning," *The Review of Financial Studies*, 2020, *33* (5), 2223–2273.

Hansen, Peter R and Asger Lunde, "Realized Variance and Market Microstructure Noise," *Journal of Business & Economic Statistics*, 2006, *24* (2), 127–161.

Hansen, Peter R., Asger Lunde, and James M. Nason, "The Model Confidence Set," *Econometrica*, 2011, *79* (2), 453–497.

Hansen, Peter Reinhard, "A Test for Superior Predictive Ability," *Journal of Business & Economic Statistics*, 2005, *23* (4), 365–380.

Harvey, Campbell R., Yan Liu, and Heqing Zhu, ". . . and the Cross-Section of Expected Returns," *The Review of Financial Studies*, 2016, *29* (1), 5–68.

Hendershott, Terrence, Dmitry Livdan, and Norman Schürhoff, "Are institutions informed about news?," *Journal of Financial Economics*, 2015, *117* (2), 249–287.

Hirshleifer, David and Siew Hong Teoh, "Limited attention, information disclosure, and financial reporting," *Journal of Accounting and Economics*, 2003, *36* (1-3), 337–386.

_ , Sonya Seongyeon Lim, and Siew Hong Teoh, "Driven to Distraction: Extraneous Events and Underreaction to Earnings News," *The Journal of Finance*, 2009, *64* (5), 2289–2325.

Holm, Sture, "A Simple Sequentially Rejective Multiple Test Procedure," *Scandinavian Journal of Statistics*, 1979, *6* (2), 65–70.

Kaniel, Ron, Shuming Liu, Gideon Saar, and Sheridan Titman, "Individual Investor Trading and Return Patterns around Earnings Announcements," *The Journal of Finance*, 2012, *67* (2), 639–680.

Ke, Guolin, Qi Meng, Thomas Finley, Taifeng Wang, Wei Chen, Weidong Ma, Qiwei Ye, and Tie-Yan Liu, "LightGBM: A Highly Efficient Gradient Boosting Decision Tree," in "Advances in Neural Information Processing Systems 30" 2017, pp. 3146–3154.

Kothari, S.P and Richard G Sloan, "Information in prices about future earnings," *Journal of Accounting and Economics*, 1992, *15* (2-3), 143–171.

Kozak, Serhiy, Stefan Nagel, and Shrihari Santosh, "Shrinking the cross-section," *Journal of Financial Economics*, 2020, *135* (2), 271–292.

Künsch, Hans R., "The Jackknife and the Bootstrap for General Stationary Observations," *Annals of Statistics*, 1989, *17* (3), 1217–1241.

Kyle, Albert S., "Continuous Auctions and Insider Trading," *Econometrica*, 1985, *53* (6), 1315.

Lee, Suzanne S., "Jumps and Information Flow in Financial Markets," *Review of Financial Studies*, 2012, *25* (2), 439–479.

_ **and Per A. Mykland**, "Jumps in Financial Markets: A New Nonparametric Test and Jump Dynamics," *Review of Financial Studies*, 2008, *21* (6), 2535–2563.

Liu, Lily Y., Andrew J. Patton, and Kevin Sheppard, "Does anything beat 5-minute RV? A comparison of realized measures across multiple asset classes," *Journal of Econometrics*, 2015, *187* (1), 293–311.

Ou, Jane A. and Stephen H. Penman, "Financial statement analysis and the prediction of stock returns," *Journal of Accounting and Economics*, 1989, *11* (4), 295–329.

Patton, Andrew J. and Kevin Sheppard, "Good Volatility, Bad Volatility: Signed Jumps and The Persistence of Volatility," *Review of Economics and Statistics*, 2015, *97* (3), 683–697.

_ **and Michela Verardo**, "Does Beta Move with News? Firm-Specific Information Flows and Learning about Profitability," *Review of Financial Studies*, 2012, *25* (9), 2789–2839.

Petersen, Mitchell A., "Estimating Standard Errors in Finance Panel Data Sets: Comparing Approaches," *Review of Financial Studies*, 2009, *22* (1), 435–480.

Piotroski, Joseph D., "Value Investing: The Use of Historical Financial Statement Information to Separate Winners from Losers," *Journal of Accounting Research*, 2000, *38*, 1.

Politis, Dimitris N. and Joseph P. Romano, "The Stationary Bootstrap," *Journal of the American Statistical Association*, 1994, *89* (428), 1303–1313.

Richardson, Scott A., Richard G. Sloan, Mark T. Soliman, and İrem Tuna, "Accrual reliability, earnings persistence and stock prices," *Journal of Accounting and Economics*, 2005, *39* (3), 437–485.

Savor, Pavel and Mungo Wilson, "Earnings Announcements and Systematic Risk," *The Journal of Finance*, 2016, *71* (1), 83–138.

Sloan, Richard G., "Do Stock Prices Fully Reflect Information in Accruals and Cash Flows About Future Earnings?," *The Accounting Review*, 1996, *71* (3), 289–315.

Weller, Brian M., "Does Algorithmic Trading Reduce Information Acquisition?," *The Review of Financial Studies*, 2018, *31* (6), 2184–2226.

Zhang, Chao, Yihuang Zhang, Mihai Cucuringu, and Zhongmin Qian, "Volatility Forecasting with Machine Learning and Intraday Commonality," *Journal of Financial Econometrics*, 2024, *22* (2), 492–530.

Appendices

A. Variable definitions and data lineage

Table 12. Variable definitions and lineage

Variable	Definition	Source / step
$A_{i,q}$	Earnings-release date: first `earnings_events.event_date` in $[E_q, F_q + 10\,\text{d}]$	FinData Postgres extract (01a)
$F_{i,q}$, $V_{i,q}$	SEC filing date; availability = F + 1 day 00:00 ET	`/fundamentals` (02)
$S_{i,q}$	$(X_q - \hat{X}_q)/TA_{q-1}$, X = operating income, $\hat{X}_q = X_{q-4} + \frac{1}{4}\sum_{k=0}^{3}(X_{q-1-k} - X_{q-5-k})$	02, `labels.py`
$D_{i,q}$	$\mathbb{1}\left[S_{i,q} \le -0.005\right]$	02
Secondary targets	net-income surprise; $\Delta\text{ROA}_{yoy} \le -0.01$; $(CFO_q - CFO_{q-4})/TA \le -0.01$; current ratio -10% or coverage < 1 from ≥ 1; $\lvert S\rvert \ge 0.005$; EPS $<$ estimate	02
Intraday families (per bar length $f \in \{1, 5, 30\}$ min)	RV = $\sum_t r_t^2$; BPV = $\frac{\pi}{2}\sum_t \lvert r_t\rvert\,\lvert r_{t-1}\rvert$; RJ = $\max(0, \text{RV} - \text{BPV})/\text{RV}$; RS^-/RV; SJV = $(\text{RS}^+ - \text{RS}^-)/\text{RV}$; realized skewness/kurtosis; $\max_t r_t^2/\text{RV}$; AC(1); drift = $\log(C/O)$; opening/closing 30-min volume shares; volume Gini; down-bar volume share; present-bar share; RV share on SPY-down bars; $\hat{\beta} = \sum_t r_t\, r_t^{\text{SPY}}/\sum_t (r_t^{\text{SPY}})^2$; vol-of-vol; $\overline{\text{RV}}/\text{Var}(r^{\text{daily}})$	03, `features.py`, `samples.py`
Daily block	20/60-day log returns, SD, Parkinson, skewness, kurtosis, overnight/intraday sums, volume ratios (20d, 5d vs 60d median), Amihud (60d), zero-return share, log dollar volume, max drawdown (60d), bars per day	03
Accounting block	gross/operating margin, NI/TA, CFO/TA, accruals, cash/TA, debt/TA, current ratio, capex/TA, R&D/revenue, coverage, loss dummy, log assets, ROA, revenue growth, net issuance, log market cap, last surprise, mean of last four surprises, trailing adverse rate, information age	02–04
Returns	split- and dividend-adjusted close-to-close from raw bars and `corporate_actions` (ex-date rule)	03

B. Leakage battery

Each guard is a unit test in `tests/`: (L1) future facts, (L2) fiscal-end joins, (L3) post-cutoff bars, (L4) test-period transforms, (L5) DST/timezone, (L6) identifier drift, (L7) splits/dividends, (L8) release-day bars, (L9) overlapping windows across splits, (L10) survivorship report, (L11) ingestion dates, (L12) randomised placebos. The live test `test_asof_equivalence` verifies that the local as-of join reproduces the API's `as_of` output on random entities and instants.

C. Additional tables

Table 13. Continuous target: ridge forecasts of the winsorised surprise (test period)

Model	N	MAE	RMSE	OOS R^2
B2	53,667	0.01911	0.03469	0.2121
B5	53,667	0.01912	0.03469	0.2117

Table 14. Transfer to held-out firms (20% of firms never seen in training)

Firms	Family	N	Δ log loss	95% CI	GW p
firmhold heldout	logistic	10,071	+0.0018	[-0.0007, +0.0045]	0.565
firmhold heldout	lightgbm	10,071	+0.0018	[-0.0005, +0.0042]	0.250
firmhold seen	logistic	43,596	+0.0011	[-0.0015, +0.0037]	0.560
firmhold seen	lightgbm	43,596	+0.0019	[+0.0005, +0.0034]	0.033

Table 15. Announcement-window abnormal returns on the ex-ante signal

Specification	Regressor	Coef.	SE	N
car on lfds5	lfds5	0.0001	(0.0042)	84,755
car on lfds2 dlfds	lfds2	-0.0002	(0.0042)	84,755
car on lfds2 dlfds	d lfds	0.0605***	(0.0209)	84,755

Notes. CAR from $c + 1$ to $A_q + 1$, market-adjusted, winsorised 1/99%; firm and sector×quarter fixed effects; two-way clustered SEs. *, **, ***: 10, 5, 1%.

D. Reproducibility

The pipeline (`scripts/00–09`) runs on the FinData storage node against the local API; manifests in `artifacts/manifests` pin dataset versions and hashes; the design freeze hash is recorded before the first model fit. Raw extracts are deleted after feature extraction and are reproducible from the pinned versions.