Heavyweight Earnings Nights and Brazilian Index Volatility: A Registered Forecasting Test That Failed

Fernando Varani Espoladore
Independent researcher
This version: 24 November 2026

Abstract

When Brazil's largest listed companies file their quarterly and annual results, is the next-day volatility of the Ibovespa, and of the WIN mini-index future, larger than a standard day-ahead forecast expects? Can a calendar of those filings improve the forecast? For every B3 session from 2012 to September 2026 we build an "earnings dose". Each filing in the securities regulator's (CVM) open-data index is weighted by the company's share of cash-market traded value and by the company's own typical past two-session move around earlier filings. We add the dose of the session being forecast as one extra regressor to a HAR forecast of log volatility. The Ibovespa target is ln(ln(H/L)), the log of the daily log high-low range, with 3,405 evaluated sessions, 1,358 of them with a positive dose. The WIN target is the log of five-minute realised volatility, with an implied-volatility term in the baseline, over 971 evaluated sessions, 355 of them dosed. Both models are refitted each day on an expanding window over the same rows. The test, its six pass marks and the predicted range of every number were registered before any outcome was computed. The registration records two tries: one counting run that computed only sample sizes and spreads, and one look at outcomes. **The registered test KILLED the claim**, meaning that two of the six pass marks failed. On the Ibovespa the dose is statistically significantly associated with a higher range than the forecast expected. The slope is +0.160 log per point of dose (95% CI 0.084 to 0.235), with a Newey-West t of 4.13 ($p \approx 0.00004$). The t is 3.66 after controlling for the same calendar shifted five sessions earlier and later. But the forecast gain is negligible. The loss ratio of the dose model over the baseline is 0.9981 on dosed sessions and 0.9969 on all sessions, against a required 0.98 or lower. These ratios are point estimates without a test statistic. On WIN the slope was not detected: −0.007, 95% CI −0.086 to +0.073, $t = -0.17$. That interval excludes the Ibovespa point estimate. A calculation from the recorded spreads shows that the 0.98 mark could not have been met on the Ibovespa even with the predicted slope known exactly, so the registered pass marks were inconsistent with each other. The dose regressor also uses filing

dates that were known only after the fact. The result is a small, statistically significant association in one historical sample of the index, with no useful forecast gain.

Keywords: earnings announcements; realised volatility; range-based volatility; HAR model; Ibovespa; B3; index futures; pre-registration; forecast evaluation; negative result

Introduction

The Ibovespa is a concentrated index, and a handful of companies carry a large share of its weight. When one of them reports results, the firm-level news may be large enough to move the index. A day-ahead volatility forecast built only from past volatility, such as the heterogeneous autoregressive (HAR) model of Corsi (2009), has no term for this calendar. The study asks a narrow question: does a weight-scaled calendar of large-company filings add forecasting value to such a model for the Ibovespa and for the WIN mini-index future traded on B3?

Related work. Several strands of research suggest that this effect should exist. Information arrival drives return variance (French and Roll 1986). Firm-level earnings news carries systematic information. Patton and Verardo (2012) find that stock betas rise on announcement days, and Savor and Wilson (2016) relate earnings announcements to systematic risk. Scheduled macroeconomic releases are a related case. Index and bond volatility, and average returns, differ on scheduled announcement days (Jones, Lamont and Lumsdaine 1998; Andersen, Bollerslev, Diebold and Vega 2003; Lucca and Moench 2015). Forecast evaluation work explains why such an association may still not improve forecasts. Daily volatility proxies are noisy, so even a correct model shows a small out-of-sample R^2 (Andersen and Bollerslev 1998; Patton 2011). When the larger model nests the baseline, as here, estimation noise alone pushes the larger model's out-of-sample squared error up (Clark and West 2007). Our result fits both strands: a statistically significant association in the direction that work predicts, too small to show up as a forecast gain.

This is a forecasting question only. Nothing in the study is traded, the paper describes no trading rule, and nothing in it is advice.

Terms used. The *registration* is a structured record written before any outcome was seen. It states the data, the method, the pass marks and the predicted numbers, and later it collects every run and its numbers. *KILLED* is the registration's label for a claim that failed at least one registered pass mark. *Tries* are the runs on real data that the registration records. The *counting run* is a run on real data that computes only sample sizes, dose quantiles and the spread of residuals, with no outcome means. The *look* is the single run that computed the outcome statistics. *Evaluated sessions* are the sessions whose forecasts enter the statistics. The registration calls them "judged sessions".

Timing and disclosure. The registration was time-stamped at 04:10:48 (UTC+1) on 24 November 2026. The counting run followed at 04:11:53 and the look at 04:13:11. The registration, the analysis code and the pass marks were drafted by an AI assistant working without step-by-step supervision, under a standing instruction from the author. Before registration, a separate session of the same AI model reviewed the code and the registration text. Its scope and limits are described under Reproducibility. The pass marks were not amended after the look. The data are historical and existed before registration, so the forecast comparison is out of sample in time but not on data that arrived after the registration. The author had used the same volatility series before (see "Forking paths" below).

Data

All data end on 30 September 2026.

- **Filing dates.** We use the CVM open-data index files of quarterly statements (ITR, 2011-2026) and annual statements (DFP, 2010-2025). A filing is one combination of company, reference date and statement type. Its date is the receipt date of version 1, the original filing, so later restatements do not move it. Filings received fewer than 10 or more than 120 days after the reference date were dropped. Of 42,186 filings, all had a version-1 row and 95.7% passed the lag filter, leaving 40,373.
- **Stock prices.** We use B3's public historical quote files (COTAHIST), 2009-2026: cash market, round lots, common, preferred and unit shares. Prices are divided by the quote factor. The fields used are high, low, close, traded value, and the distribution counter, which voids moves that span a dividend or corporate action.
- **Code-to-company map.** We map B3 trading codes to company tax identifiers (CNPJ). The author built the map earlier, for other work, from the trading codes companies list in their CVM registration-form (FCA) data, then by issuer-code root and a unique name match. The map includes codes of delisted companies. **Coverage:** only 10,973 of the 40,373 kept filings (27%) belong to mapped companies, and 153 companies produced dosed events. Many CVM registrants have no listed shares, so a low share of filings need not mean a low share of traded value. However, the share of cash-market traded value, or of index weight, held by mapped companies was **not measured** for this version. **Point in time:** the map is a single undated table. A company whose ticker changed is covered if both codes are in the map, because its weight sums the traded value of all its mapped codes. A code later reused by another company would be assigned to only one of them. Mergers and delistings are reflected only in that filings stop. These cases were not audited.
- **Ibovespa.** We use daily high, low and close from 2005 to September 2026, from a commonly available public daily history. Stale rows, where the close repeats with zero volume, were dropped, and the high had to exceed the low. Target: $y(t) = \ln(\ln(H(t)/L(t)))$.
- **WIN.** We use one-minute bars of a vendor-built continuous WIN series, September 2021 to September 2026. The series switches to the next contract around each bimonthly expiry under the vendor's own rule, which the study did not document, and it is **not** back-adjusted. Sessions are restricted to bars between 09:00 and 18:25 Brasília time, with at least 200 bars (median 565). The realised measure is $RV5(t) = \text{sqrt}(\Sigma\, r^2)$, where the r are log returns between every fifth one-minute close inside that window. It therefore excludes the overnight gap, the move from the previous close to the first bar, and any price step at a contract switch *between* sessions. Roll sessions were **not** excluded. We did not check whether the vendor's switch ever falls *inside* a session, which would put one artificial jump into that day's RV5. B3 changes its session hours during the year to follow US daylight-saving time. The fixed clock window was not adjusted for this, so the number of bars per session varies with the

schedule. A session enters the evaluation only if it has at least 80% of the median bar count. This removes short sessions but not the smaller seasonal shifts. Target: $y(t) = \ln(RV5(t))$.

- **Implied volatility.** We use the daily implied volatility of the BOVA11 (Ibovespa ETF) option with the strike closest to the ETF price, in the nearest monthly expiry at least seven days away, built from closing option prices in the same B3 quote files and carried forward at most three sessions. BOVA11 options traded well before 2021, but the author's series was built only from 2021 onward. It therefore cannot serve an Ibovespa baseline that starts in 2012, and it is used only in the WIN baseline, whose sample starts in September 2021.

Method

Weight. For a filing, the company's weight w is its codes' summed traded value over the 252 sessions (at least 200) ending at the last month-end before the filing's month, divided by the traded value of all cash stocks over the same sessions. This is a point-in-time, traded-value stand-in for index weight. Only filings with w of at least 0.25% enter, with one event per company and receipt date.

Event window and typical move. The event covers $s0$, the first session on or after the receipt date, and $s1$, the next session, so that both pre-open and after-close releases are caught. The event move is $|\ln(\text{close at } s1 \,/\, \text{close the session before } s0)|$ of the company's most liquid code. It is a raw move, not adjusted for the market. A company's typical move m is the median of its own earlier valid moves, counting only events that ended before the current one starts and requiring at least four. Otherwise m is the median of all earlier moves in the universe, requiring at least 20.

Dose. The dose of an event is $100 \times w \times m$, in percent of the index. The dose of a session, $x(t)$, is the sum over all events whose window contains t, and zero otherwise. It is defined from 2012; 2011 events only seed the typical moves. This produced 4,341 dosed events from 153 companies, 531 of them from companies with w of at least 2% (the "heavyweight" events).

Look-ahead in the dose. The weight and the typical move use only earlier data. **The event dates do not.** They are CVM receipt dates, which are known only once the filing arrives. A filing received during or after session $t+1$ enters $x(t+1)$, but it was not known when the forecast for $t+1$ was made at the end of session t. Most large companies announce their release dates in advance, but **no announced release calendar was used**, and the gap between announced and actual dates was not measured. The dose model is therefore an **oracle-calendar forecast**: it is given the true filing dates. This favours the dose model. Since even this favoured version did not improve the forecast, the negative forecast result is, if anything, stronger. But the study does not show that the dose can be built in real time.

Forecasts. For each market the baseline is a HAR model: $\hat{y}(t+1) = \beta 0 + \beta 1 \cdot y(t) + \beta 5 \cdot \text{mean}(y(t-4..t)) + \beta 22 \cdot \text{mean}(y(t-21..t))$. For WIN it also includes $\gamma \cdot \ln IV(t)$. The test model (HARX) adds $\delta \cdot x(t+1)$ and so nests the baseline. Both are refitted every day by least

squares on an expanding window of earlier rows only, with at least 250 rows, on exactly the same rows. The residual is $r(t+1) = y(t+1) - \hat{y}(t+1)$, the log ratio of the realised measure to the forecast; rx is the same for HARX. Evaluated sessions run from January 2013 to September 2026 for the Ibovespa, and from the first WIN forecast (about October 2022) to September 2026 for WIN.

Pass marks (all six had to pass).

1. Ibovespa: the Newey-West (1987) t, with 5 lags, of the slope b in the full-sample regression $r = a + b \cdot x$ is at least 2 in absolute value. Its sign s fixes the direction for the other marks; the test was two-sided by registration.
2. Placebo control: $s \times t$ of b in $r = a + b \cdot x + c \cdot x(-5) + d \cdot x(+5)$ is at least 2. Here $x(-5)$ is the same calendar moved so that every event window sits **five sessions earlier**, so that $x(-5)(t) = x(t+5)$: at session t it holds the dose of the week before the releases. Likewise $x(+5)$ moves every window **five sessions later**, so that $x(+5)(t) = x(t-5)$. The control separates the release sessions themselves from the sessions five days around them. It does not control for season-wide or longer-horizon patterns.
3. Forecast gain: the loss ratio $\Sigma\, rx^2 / \Sigma\, r^2$ on dosed sessions is at most 0.98, for the worse of the two markets.
4. Stability: $s \times t$ in each Ibovespa half (2013-2019 and 2020-2026) is at least 1.
5. WIN agreement: $s \times t$ of the dose slope against WIN residuals is at least 1.
6. Calendar sanity: for heavyweight events, the company's own $\ln(\ln(H/L))$ averaged over s0 and s1, minus its mean over sessions s0−22 to s0−3 (at least 10 valid), has a median of at least 0.20. This is a raw, not market-adjusted, range of a volatility measure. It has no sign and is **not** multiplied by s.

Before registration, the measuring code was run on synthetic data: pure noise failed at least one mark, and noise with a planted dose effect passed all six. The counting run then checked the power of the real sample. The minimum detectable effect, defined as $2.8 \times \mathrm{sd}(r) / (\mathrm{sd}(x) \cdot \sqrt{n})$ times the 90th-percentile dose on dosed sessions, was 0.056 log on the Ibovespa and 0.057 log on WIN, against a required maximum of 0.10.

Results

Table 1. The six registered pass marks at the single look.

Mark	Required	Result	Passed
1. Ibovespa \	t\	of dose slope (sign s = +1)	≥ 2
2. Placebo-controlled t	≥ 2	3.66	yes
3. Worse dosed-session loss ratio (HARX / baseline)	≤ 0.98	0.9981 (Ibovespa); WIN 0.9976	**no**
4. Worse Ibovespa half t	≥ 1	2.84 (halves 3.11 / 2.84)	yes
5. WIN signed t	≥ 1	−0.17	**no**
6. Stock sanity lift (median, 523 events)	≥ 0.20	0.284	yes

Two marks failed, so under the registered rule **the registered test KILLED the claim**. The forward stage on new sessions that the registration foresaw was not run.

Table 2. Detail by market, as recorded at the look. Newey-West t with 5 lags; two-sided p from the normal approximation. The 95% intervals are b ± 1.96·b/t, computed afterwards by arithmetic from the recorded b and t.

Quantity	Ibovespa	WIN
Evaluated sessions / with dose > 0	3,405 / 1,358	971 / 355
Dose slope b (log per 1 point of dose)	+0.160	−0.007
Standard error; 95% CI	0.039; [0.084, 0.235]	0.041; [−0.086, +0.073]
t (p)	4.13 (p ≈ 0.00004)	−0.17 (p ≈ 0.87)
t of b with x(−5) and x(+5) in the regression	3.66 (p ≈ 0.0003)	−0.69
Placebo alone: x(−5), windows 5 sessions earlier	t = 1.47	t = 0.41
Placebo alone: x(+5), windows 5 sessions later	t = 0.64	t = 1.21
Halves 2013-2019 / 2020-2026: slope (t)	+0.211 (3.11) / +0.130 (2.84)	—
Loss ratio, dosed sessions	0.9981	0.9976
Loss ratio, all evaluated sessions	0.9969	1.0048
Top-decile dose sessions: mean r minus undosed mean r	+0.087	−0.005
90th percentile of dose on dosed sessions	0.48	0.60
Minimum detectable effect at that dose (counting run)	0.056 log (slope ≈ 0.115)	0.057 log (slope ≈ 0.096)

The loss ratio is a point estimate, not a test. The registered comparison against the baseline is the loss ratio, which comes without a standard error. A Diebold-Mariano test (Diebold and Mariano 1995) of the loss difference was registered only for the forward stage, which was never reached. Because HARX nests the baseline, the Clark-West (2007) adjusted statistic would be the right descriptive test here. **Neither statistic is reported in this version.** Each needs the per-session losses, which the look did not save, so computing either one means a new, unregistered run on the data. That run has not been

made. The ratio is therefore not evidence of a significant gain; the calculation below shows why only a small gain was possible anyway.

Table 3. Registered predictions against outcomes.

Quantity	Predicted	Observed	Inside range?
Ibovespa slope	+0.08 to +0.20	+0.160	yes
Ibovespa t	2.5 to 5	4.13	yes
Placebo-controlled t	2 to 4	3.66	yes
Half t	1.5 to 3.5 each	3.11, 2.84	yes
WIN signed t	1 to 3	−0.17	**no**
Dosed-session loss ratio	0.96 to 0.99	0.9981 (Ibovespa), 0.9976 (WIN)	**no** (weaker than the weak end)
Top-decile lift	+0.05 to +0.15	+0.087	yes
Stock sanity lift	+0.30 to +0.60	+0.284	**no** (just below)
Dosed events / heavyweight events	2,500-4,500 / 400-700	4,341 / 531	yes
Dosed sessions, Ibovespa / WIN	1,200-1,800 / 300-500	1,358 / 355	yes
Prior chance the claim survives	about 0.3	killed	—

Why the 0.98 mark was out of reach on the Ibovespa

The registered predictions for the slope (0.08 to 0.20) and for the dosed loss ratio (0.96 to 0.99) were not consistent with each other. The calculation below uses only recorded numbers, and was done after the look.

The counting run's minimum detectable effect fixes the ratio of residual to dose spread: $\mathrm{sd}(r)/\mathrm{sd}(x) = \mathrm{MDE}\cdot\sqrt{n} / (2.8 \times p90)$. For the Ibovespa this is $0.0555 \times \sqrt{3405} / (2.8 \times 0.483) \approx$ **2.39**. For WIN it is $0.0570 \times \sqrt{971} / (2.8 \times 0.596) \approx$ **1.07**. The two variances themselves were not saved. Suppose the dose's true slope b were known exactly and the dose term were the only source of gain. Then the dose would remove a share $b^2\cdot\mathrm{var}(x)/\mathrm{var}(r)$ of the baseline's squared error on all sessions. For the Ibovespa at $b = 0.16$ that share is $(0.16/2.39)^2 \approx$ **0.0044**, i.e. 0.44%. That gain cannot exceed the total gain, and dosed sessions carry roughly 40% of all sessions (1,358 of 3,405), so the dosed-session gain is at most about $0.0044/0.40 \approx 1.1\%$. **The best achievable dosed loss ratio is therefore about 0.989.** Across the predicted slope range the bound is 0.997 at $b = 0.08$ and about 0.983 at $b = 0.20$. The 0.98 mark would need $b \geq$ about 0.21 even in this best case, and the predicted 0.96 would need $b \approx 0.30$. With the slope estimated on an expanding window, the achievable gain is smaller still. **The forecast-gain mark was very likely to fail even if the association was exactly as predicted.** This matters more than the arbitrary choice of 0.98: the registered pass marks were internally inconsistent. The counting run came after the registration, but the same arithmetic could have been done before it.

For WIN the same arithmetic gives $(0.16/1.07)^2 \approx 2.3\%$ on all sessions. The 0.98 mark was within reach there if the predicted slope had appeared, but it did not. These calcula-

tions use the referee's rough reasoning with the study's own spread ratio, and they rest on two approximations. First, the dosed sessions' share of squared error is taken to be close to their share of sessions; a share between 30% and 50% changes the Ibovespa bound only to between about 0.985 and 0.991. Second, the gain is taken to come only through the dose term.

Where the small Ibovespa gain came from

The all-session ratio is a weighted average of the dosed and undosed ratios, with weights equal to each group's share q of the baseline's squared error. Because the all-session ratio (0.9969) is lower than the dosed ratio (0.9981), **the undosed-session ratio must be lower than 0.9969**. For q between 0.30 and 0.50 it is 0.9957 to 0.9964, about 0.996 at the session share $q \approx 0.40$. Over the same range of q, dosed sessions contribute only 18% to 30% of the total reduction in squared error, so **the small gain that exists does not come mainly from release nights.** A plausible mechanism is this. When x is left out, the baseline's intercept absorbs the average dose effect, so its forecasts on undosed days are slightly too high. Adding x lets the intercept and HAR coefficients shift and lowers those forecasts. Meanwhile the noisily estimated slope helps little on the dosed days themselves. The data at hand do not separate these channels. For WIN the same identity gives an undosed ratio of about 1.008 to 1.011. There, the dose model is worse on undosed sessions and slightly better on dosed ones, with an estimated slope near zero. This is what refitting noise in a nested model looks like (Clark and West 2007).

Ibovespa versus WIN

The WIN slope's 95% interval, [−0.086, +0.073], excludes the Ibovespa point estimate of 0.160. It also excludes the lower end of the predicted range, 0.08, though barely. If the two estimates were independent, their difference would have a z of about 3. They are not independent, because the WIN sessions overlap the Ibovespa's last four years, so this figure is only indicative. The two slopes are also not directly comparable. The targets differ (log daily range versus log five-minute realised volatility), the baselines differ (the WIN baseline includes implied volatility), and the samples differ (2013-2026 versus 2022-2026). The study cannot tell whether the difference is due to the market, the target, the implied-volatility term or the period. Two exploratory checks would separate the period from the market: the Ibovespa slope estimated only on the WIN evaluation sessions, and the Ibovespa slope with the BOVA11 implied-volatility term added to its baseline over the same years. **Neither check is in this version.** Both would be post-hoc runs on the data and have not been made. The Ibovespa's second-half slope, +0.130 ($t = 2.84$, 2020-2026), suggests the association did not vanish in recent years, but that period is not the WIN sample.

The calendar check

Heavyweight companies' own daily range on s0 and s1 is about 0.28 log (roughly a third) above its level over the preceding month (median over 523 events). This confirms that the

receipt dates line up with real price activity, but it fell just short of the registered range of 0.30 to 0.60. The 523 events are not a subset of the 531 heavyweight events. The 531 are heavyweight events with a known typical move. The 523 are heavyweight events in the same years, whether or not a typical move was known, whose most liquid code had a valid range on s0 or s1 and at least 10 valid sessions in the baseline month. The run did not record the breakdown of the difference. The company range is raw: on release days it also contains any market-wide move, though the index effect (about 0.08 log at the top-decile dose) is small next to 0.28.

In plain terms, the dose effect on the Ibovespa at the 90th-percentile dose is about 0.16 $\times$ 0.48 $\approx$ 0.08 log, roughly an 8% wider daily range than the forecast. That is large enough to be detected over 14 years. It is small next to the day-to-day noise of a daily range, and the forecast error barely changes.

What this does and does not show

It shows that, in 2013-2026 Ibovespa data, a weight-scaled calendar of large-company filings is statistically significantly associated with a higher-than-forecast daily range. The association holds in both halves of the sample and beyond the same calendar shifted ± 5 sessions. Its size is close to the slope predicted before the look. Other predictions missed their ranges: the WIN t (-0.17 against 1 to 3), the stock sanity lift (0.284 against 0.30 to 0.60) and the loss ratio (weaker than the weak end of its range). The study also shows that the association does not translate into a meaningful day-ahead forecast improvement on either market, even with the true filing dates given in advance. An effect on WIN five-minute realised volatility over 2022-2026 was not detected: the 95% interval excludes an effect as large as the Ibovespa estimate, but not a small one. Under the rule fixed in advance, the claim that this calendar term improves the forecasts is rejected.

It does not show that the association is causal or that it holds outside this sample. The Ibovespa baseline has no implied-volatility term, which might already capture part of a known calendar, and it is built from a series the author had used before. Nor does the study show why WIN differs (see above), that a dose built from announced calendars would perform the same, or that a differently designed forecast could not use this information.

Limitations. The filing receipt date stands in for the public release date and may differ from it; the two-session window is meant to absorb this. The weight is a traded-value proxy, not the official index weight, and the code-to-company map covers an unmeasured share of traded value. The WIN series is a vendor's unadjusted continuous series, with roll sessions included and the session window not adjusted for schedule changes. The forecast gain was judged only by a loss ratio without a standard error, and that pass mark, as shown above, was effectively out of reach on the Ibovespa.

Forking paths beyond the registration. The registration counts two tries for this study. Before them, the code was developed and checked on synthetic data; those runs used generated series only, and their number was not recorded. After the verdict, the same script was run once more to reproduce the numbers, and nothing changed. More broadly,

the author had registered at least nine earlier tests on the same or closely related Ibovespa and WIN day-volatility series, with 65 recorded tries among them. Seven were on WIN day-ahead volatility, including tests of a dozen published HAR additions, among them a generic calendar term, and of more implied-volatility terms. One was the Ibovespa daily-range HAR used here as the baseline. One tested a calendar of US-holiday absences on both markets. There was also earlier unregistered exploratory work on these series whose extent is not fully recorded. None of the earlier tests used earnings dates on these index series. An earlier test on individual stocks' daily ranges had named earnings-day jumps as its main problem, and that motivated this study. Several further calendar terms were tested on the same series after this study, the same day; they are not part of this paper. A reader should treat the Ibovespa t of 4.13 as one result among many tests on these series.

Nothing in this paper is investment or trading advice.

Reproducibility

Registration. The registration, with its timestamps, pass marks, predictions and the log of every run, is stored in the author's private, version-controlled research repository. The analysis script's content hash is recorded with each run. At the time of the runs the script had not yet been committed. No public commit, posted hash or third-party timestamp exists, so **the ordering of registration, counting run and look is attested only by the author's own records and cannot be independently verified.**

Review. Before registration, a separate session of the same AI model (Claude) reviewed the analysis script and the registration text and approved them. It had no access to outcome numbers. "Independent" means only that: a separate session, given the documents and not the outcomes. It is not a human review, and it shares the drafting model's training and likely blind spots.

Replication. After the verdict, the full analysis was run again with the unchanged script on the same data. All 36 numbers shared with the registered record were reproduced exactly. This mainly shows that the computation is deterministic and that the recorded numbers are what the script produces. It does not check the code's correctness, and no independent re-implementation was made.

Data and code. The filing index (CVM open data) and the stock and option quotes (B3 COTAHIST) are public. The Ibovespa daily series is widely available. The WIN one-minute bars come from a non-public vendor feed, and the code-to-company map was compiled by the author; both affect the exact counts. The analysis script is not published with this version.

Not in this version. The referee asked for the following, but they need new runs on the data and none was made for this revision: Clark-West and Diebold-Mariano statistics on dosed and on all sessions; the Ibovespa slope on the 2022-2026 WIN sessions, with and without the implied-volatility term; the traded-value coverage of the code-to-company map; a breakdown of the 531 and 523 heavyweight-event samples; and a check for contract switches inside WIN sessions. All Table 2 intervals and the calculations in the Results are arithmetic on recorded numbers, done after the look.

Parts of the analysis and the writing were done with an AI assistant.

References

Bibliographic details marked † were checked against reference lists in the author's earlier papers. Those marked ‡ come from the author's and the assistant's knowledge of the literature and were not re-verified against the journals for this version.

Andersen, T. G., & Bollerslev, T. (1998). Answering the skeptics: Yes, standard volatility models do provide accurate forecasts. *International Economic Review*, 39(4), 885–905. †

Andersen, T. G., Bollerslev, T., Diebold, F. X., & Vega, C. (2003). Micro effects of macro announcements: Real-time price discovery in foreign exchange. *American Economic Review*, 93(1), 38–62. †

Clark, T. E., & West, K. D. (2007). Approximately normal tests for equal predictive accuracy in nested models. *Journal of Econometrics*, 138(1), 291–311. †

Corsi, F. (2009). A simple approximate long-memory model of realized volatility. *Journal of Financial Econometrics*, 7(2), 174–196. †

Diebold, F. X., & Mariano, R. S. (1995). Comparing predictive accuracy. *Journal of Business & Economic Statistics*, 13(3), 253–263. †

French, K. R., & Roll, R. (1986). Stock return variances: The arrival of information and the reaction of traders. *Journal of Financial Economics*, 17(1), 5–26. ‡

Jones, C. M., Lamont, O., & Lumsdaine, R. L. (1998). Macroeconomic news and bond market volatility. *Journal of Financial Economics*, 47(3), 315–337. ‡

Lucca, D. O., & Moench, E. (2015). The pre-FOMC announcement drift. *Journal of Finance*, 70(1), 329–371. †

Newey, W. K., & West, K. D. (1987). A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. *Econometrica*, 55(3), 703–708. †

Patton, A. J. (2011). Volatility forecast comparison using imperfect volatility proxies. *Journal of Econometrics*, 160(1), 246–256. †

Patton, A. J., & Verardo, M. (2012). Does beta move with news? Firm-specific information flows and learning about profitability. *Review of Financial Studies*, 25(9), 2789–2839. ‡

Savor, P., & Wilson, M. (2016). Earnings announcements and systematic risk. *Journal of Finance*, 71(1), 83–138. ‡