Is the day's range already made? A random-walk calculation with nothing fitted, tested on index, currency and crypto futures

Fernando Varani Espoladore

Independent researcher

This version: 22 October 2026 (with an online appendix)

Abstract

Part-way through a trading session, how likely is it that the rest of the session leaves the range made so far essentially unchanged? Under a driftless random walk the probability follows from the distances to the high and the low and the volatility still to come; we compute it exactly and evaluate it as a benchmark on about 1,180 sessions each of the Brazilian mini-index (WIN) and mini-dollar (WDO) futures, 2021–2026, against simulated random walks whose volatility is as hard to forecast as in the data. Its average probability is within one point of the observed share on both markets. The registered comparison does not show fitted models ranking readings better, which is not evidence that they are equal; with a noisier volatility estimate they add some skill on WIN. Its probabilities are too extreme (calibration slope 0.58), mostly because an uncertain volatility forecast is plugged into a nonlinear formula: the simulations show the same, and averaging over the forecast's error raises the slope to 0.77–0.78, against 0.94 in simulation. Neither richer descriptions of that error nor volatility that responds to the price closes the remaining gap. A diagnostic with the realised remaining variance, calibrated by simulation, finds WDO consistent with the Brownian law in level, with fewer small extensions than it implies, and, tentatively, WIN's range extended slightly more often than it implies. On Bitcoin and Ether the probability is within two points in level but far more extreme. The calculation is a benchmark that any claim to know whether "the range is in" should be measured against.

Keywords: intraday range; random walk; Brownian motion; double-barrier probability; forecast calibration; index futures

JEL: G13, G14, C53

Introduction

Traders often ask whether "the range is in": whether the market, some hours into the session, will stay inside the high and low it has already made. Market Profile practice calls the first hour's range the "initial balance" and counts later "range extensions" (Drinka & Ptasienski, 1989). Indicators that claim to answer the question are sold, usually without a stated method.

The academic literature has studied neighbouring questions. One is the time of day at which the high and the low occur: under a driftless random walk the extremes fall near the open or the close more often than in the middle of the day. Mok, Lam and Li (2000) cannot reject that law on S&P 500 futures but find extra extremes early in the Hang Seng session [described from the abstract; text not yet read]; Acar, Lequeux and Ritz (1996) find highs and lows on FT-SE 100 futures clustering near the open and close more than a driftless random walk predicts, and Panas (2005) reports the same U-shape on Athens stocks; Copeland, Lam and Jones (2004) reject it in floor and screen trading alike [described from the abstract and a second-hand account; text not yet read]. Another is forecasting the high and the low themselves: Cheung (2007) models the cointegrated daily highs and lows of three US indexes, and Caporin, Ranaldo and Santucci de Magistris (2013) find those of Dow Jones stocks "largely predictable", a predictability that comes, in their words, from "the predictability of second-order moments". A related line forecasts the range as a measure of volatility (Parkinson, 1980; Chou, 2005; Christensen & Podolskij, 2007; Martens & van Dijk, 2007); the joint law of the high, the low and the close of a Brownian path underlies the range-based variance estimators (Garman & Klass, 1980; Rogers & Satchell, 1991) and range-based estimation of stochastic volatility models (Alizadeh, Brandt & Diebold, 2002). Here the same law is used the other way round: given the variance, it states a probability for the extremes. A third assesses trading on the range: opening-range breakouts on crude oil (Holmberg, Lönnbark & Lundström, 2013) and the Nasdaq-100 (Zarattini & Aziz, 2023), and, against a random-walk view, support and resistance levels that predicted intraday trend interruptions in exchange rates (Osler, 2000), with a proposed mechanism in clustered stop-loss and take-profit orders (Osler, 2003).

The probability we study is a classical object. For a Brownian path it is the probability of staying between two levels, the double no-touch probability of option pricing (Kunitomo & Ikeda, 1992; Hui, 1996), which follows from the joint law of the maximum and the minimum (Feller, 1951). Related functionals of the running maximum appear in drawdown and optimal-stopping problems (Taylor, 1975; Lehoczky, 1977; Graversen, Peskir & Shiryaev, 2001; Hadjiliadis & Večeř, 2006; du Toit & Peskir, 2009).

The closest tests in spirit are tests of the continuous-martingale hypothesis by a change of time. If the price is a continuous martingale, its moves measured in units of its own quadratic variation are those of a Brownian motion. Peters and de Vilder (2006) use this time change to test whether S&P 500 futures returns, corrected for drift, are a time-changed Brownian motion, and cannot reject it; Andersen, Bollerslev and Dobrev (2007) derive the distributional implications of returns standardised by realised volatility when there are jumps, leverage effects and noise, and test them on S&P 500 futures. Our diagnostic with

the true remaining variance rests on the same time change (Method), but scores a functional of the path, whether its extremes are extended before the close, rather than the return over an interval. It can therefore react to departures that leave standardised returns close to normal, such as a dependence of the path on its own extremes, and it shares those tests' main difficulty: the variance is measured, with error, on the path being tested. We read it against simulations for that reason, and we do not offer it as a new test of the martingale hypothesis.

What we did not find, in several rounds of search, is this probability used as a forecast on intraday data, with its calibration measured against realised frequencies and compared with fitted models. We do not claim the search was exhaustive.

This paper states the random-walk probability that the rest of a session leaves the range made so far essentially unchanged, computed exactly. It measures the calibration and ranking skill of that probability on two Brazilian futures over five years, against simulated random walks that reproduce the volatility pattern and the volatility-forecast errors of the real data, so that the properties of the method are separated from those of the market. And it compares the probability with fitted models, including a tuned flexible learner. The answer is a benchmark: any indicator that claims to know whether "the range is in" should be compared with this calculation, not with the clock. It also bears on how likely a stop beyond the day's extreme is to be reached before the close, and on how far a Brownian pricing assumption is from realised frequencies at the intraday horizon.

The contribution is not a new barrier probability, nor a new test of the martingale hypothesis. It is an evaluation of a known one as a forecast when the variance still to come is unknown, including how much of its apparent departure from the Brownian law that uncertainty accounts for. Four things are learned that the mathematics alone does not give. First, with no parameter fitted to outcomes, the probability ranks the session's readings about as well as the fitted models tried. Second, used with an estimated variance it is miscalibrated in a predictable direction and by a predictable amount: it is too extreme, by roughly as much as simulated random walks with equally hard-to-forecast volatility show. This applies to any barrier or first-passage probability evaluated at a volatility forecast. The effect itself is familiar (a nonlinear function evaluated at a point estimate of its argument); what is measured here is its size and direction for this probability, which the simulations reproduce. Third, averaging over the forecast's error corrects most of that in simulation but only part of it in the data, and the richer error models we tried do not close the remainder. Fourth, fitted models add information in some settings (a noisier volatility estimate on WIN, the coins) but not reliably in the registered one, so the benchmark is strong but not complete.

Data

WIN and WDO. We use one-minute candles (open, high, low and close) of the continuous mini-index (WIN) and mini-dollar (WDO) futures of B3, from 24 September 2021 to 30 September 2026 (WDO to 24 September 2026), from a broker's MetaTrader 5 history and, for WIN from 30 December 2025, from the Profit charting platform (Nelogica). A session is used when its first candle is at or before 09:05 and its last at or after 17:50, local time; B3

changes its closing time with US daylight saving, so about 28% of sessions end half an hour early. Readings are taken every half hour from 10:00 to 16:30, fourteen a session, when at least 30 candles precede them. The first 40 sessions serve as memory (20 for the fitted model's average daily range, then 20 for the volatility forecast), and the predictive version below needs 20 further sessions of forecast errors, so the comparisons use 16,457 readings on 1,180 sessions for WIN and 16,492 on 1,178 for WDO. Rolls fall between sessions, so no reading spans one. The platform's 2026 series is back-adjusted by a multiplicative factor (about 4%); the calculation uses only ratios of price distances to volatility, and the online appendix (S4) reports 2026 without the sessions whose memory straddles the change of source. Where the two series overlap (184 WIN and 90 WDO sessions in 2026), their candles are identical to the tick once the back-adjustment is divided out, and the outcome of all but one of 3,752 readings is the same on either; bars built from trade-by-trade prints since August 2026 agree with them on WIN, and within one tick on WDO. Large one-minute moves do not cluster on expiry days, and one isolated print (WDO, 26 June 2025) changes no result by more than 0.1 point (online appendix S4).

Bitcoin and Ether. Five-minute candles of the BTCUSDT and ETHUSDT perpetual futures of Binance (USD-M), from the exchange's public archive (data.binance.vision), 2020 to 2023, with the UTC day as the session and readings every two hours from 04:00 to 20:00: 12,789 readings on 1,421 days for each coin in the registered test's script, and 12,609 on 1,401 in the analyses of this paper, which keep 20 more days as the predictive version's memory.

Three other index markets. Hourly bars of the E-mini S&P 500 and E-mini Nasdaq-100 futures (09:00 to 16:00 New York) and of the DAX (09:00 to 17:30 Frankfurt), for 2026 to 9 October, from Yahoo Finance.

The Binance data is public. The broker, platform and Yahoo data are vendor data: we share code and derived statistics, never the rows.

Method

The target

Take a reading at time t of session d, the session running from 0 to its close T. Let H and L be the highest and lowest prices so far, $R = H - L$ the range so far, and p the last close. Let M and m be the highest and lowest prices of the rest of the session. The extension *ext* is the total distance by which the rest of the session pushes beyond the range, and the range HOLDS ($Y = 1$) if it is no more than a fraction κ of R:

$$\text{ext} = \max(M - H,\ 0) + \max(L - m,\ 0), \qquad Y = \mathbf{1}\{\text{ext} \le \kappa R\}, \qquad \kappa = 0.2 \tag{1}$$

The allowance κ is a convention, not a consequence of the theory. A fifth of the range was fixed in the first registered test of this research (10 October 2026), before any of the measurements reported here, as a practical reading of "the range is essentially unchanged"; it was not tuned, and it was kept so that the registered tests apply. The online appendix repeats the analysis at $\kappa = 0$ (the plain double no-touch event), 0.1 and 0.3, and the WIN departure reported below depends on the choice.

The probability under a random walk

The null hypothesis is narrower than "a random walk". From the reading to the close, the price is taken to be a driftless continuous martingale whose quadratic variation over that time is a number *s*² known at the reading: given everything known then, the rest of the path is such a martingale, and *s*² is not random. By the theorem of Dambis (1965) and Dubins and Schwarz (1965), such a martingale is a Brownian path run at a changed speed, and over the rest of the session it covers exactly that path run for a total variance *s*². The formula is then exact however the variance is spread over the afternoon, even when the spreading depends on the path. It is not exact when the remaining variance is itself random and related to the path, for example when volatility rises after a new extreme: the variance then acts as a random clock that depends on the path. The comparisons can therefore detect drift (including momentum and mean reversion), jumps, any dependence of the path on its own extremes, and such a feedback between the path and its variance. The forecast of s is studied separately. Suppose, then, that the rest of the session is a driftless Brownian path starting at p with total standard deviation s. In units of s, the distances to the two levels and the allowance are

$$a = \frac{H - p}{s}, \qquad b = \frac{p - L}{s}, \qquad c = \frac{\kappa R}{s} \tag{2}$$

For a standard Brownian motion on [0, 1], the probability of staying below u and above $-v$ ($u, v > 0$) is the classical series (Feller, 1951)

$$F(u, v) = \sum_{k=-\infty}^{\infty} \{[\Phi(u + 2kw) - \Phi(-v + 2kw)] - [\Phi(2u + v + 2kw) - \Phi(u + 2kw)]\} \tag{3}$$

where $w = u + v$ and Φ is the standard normal distribution function. The range holds either if the low is not broken and the high is exceeded by at most c, or if the low is broken by some $y \le c$ and the high is exceeded by at most $c - y$. Since $F(u, v)$ is the probability that the path's maximum is at most u and its minimum at least $-v$, $\partial F/\partial v$ is the density of the minimum at $-v$ jointly with a maximum of at most u, and the probability that the range holds is

$$P(a, b, c) = F(a + c,\ b) + \int_0^c \frac{\partial F}{\partial v}(a + c - y,\ b + y)\, dy \tag{4}$$

which we evaluate with the series truncated at $|k| \le 6$ and Gauss–Legendre quadrature (48 nodes). The result agrees with Monte Carlo simulations of 4,000- and 20,000-step paths to within their sampling error once the Broadie, Glasserman and Kou (1997) continuity correction is applied (`revision/validate_exact.py`). A second implementation by a different route, the spectral (Fourier sine) expansion of the joint law with adaptive quadrature, sharing no code with the first, agrees on every WIN reading to within 5×10^{-9} (`revision/check_exact.py`); the truncated series loses accuracy, by up to 0.002, only when the distances and the allowance add up to less than 0.2 volatility units, which no reading approaches (smallest 0.35). An earlier 400-step routine overstated the probability by up to about one point; every number below uses the exact formula.

The volatility still to come

s^2 has to be stated before the fact. Realised variance between two times of session d is the sum of squared changes of the close p, in price points, sampled every Δ minutes:

$$V_d(t_1 \to t_2) = \sum_{j=1}^{n} (p_{t_1 + j\Delta} - p_{t_1 + (j-1)\Delta})^2, \qquad n = \frac{t_2 - t_1}{\Delta}, \qquad \Delta = 5 \text{ minutes} \tag{5}$$

The plug-in forecast scales the average remaining variance of the previous 20 sessions, at this time of day, by how active today has been so far relative to those sessions:

$$\hat{s}^2_{d,t} = \bar{V}_d(t \to T)\, \frac{V_d(0 \to t)}{\bar{V}_d(0 \to t)}, \qquad \bar{V}_d(t_1 \to t_2) = \frac{1}{20} \sum_{i=1}^{20} V_{d-i}(t_1 \to t_2) \tag{6}$$

The average carries the level of volatility and its intraday pattern (Andersen & Bollerslev, 1997). The five-minute interval was chosen after a one-minute version had produced a spurious finding on WDO; results at other intervals are in the online appendix (S3). The variance the rest of the session actually had, and the forecast's log error, are

$$s^{*2}_{d,t} = V_d(t \to T), \qquad e_{d,t} = \ln \frac{s^{*2}_{d,t}}{\hat{s}^2_{d,t}} \tag{7}$$

Three versions of the probability are compared:

- **Plug-in**: $P(\hat{a}, \hat{b}, \hat{c})$, equation (4) with $s = \hat{s}$. This is the calculation of the registered tests.
- **True variance**: $P(a*, b*, c*)$, equation (4) with $s = s*$. It is not a forecast. It removes forecast error, but the variance it uses is measured on the same path whose extremes are scored, so it is a diagnostic read against the simulations, not a direct test of the Brownian law (see Results).
- **Predictive**: the plug-in probability averaged over the forecast's own error. We take e to be normal with mean μ_t and standard deviation σ_t estimated, at each reading time, from the previous 250 sessions only, and average over it (nine-point Gauss–Hermite); a variance error x divides every distance in (2) by $\exp(x/2)$:

$$P^{\text{pred}}_{d,t} = \int_{-\infty}^{\infty} P\left(\hat{a}\, e^{-x/2},\ \hat{b}\, e^{-x/2},\ \hat{c}\, e^{-x/2}\right) \frac{1}{\sigma_t}\, \phi\!\left(\frac{x - \mu_t}{\sigma_t}\right) dx \tag{8}$$

A **conditional** variant lets the error's mean depend on the morning's scaling ratio r, through a least-squares line fitted at each reading time over the previous 250 sessions; μ in (8) is then the line's prediction $\alpha_t + \gamma_t \ln r$, and σ the standard deviation of its residuals:

$$e_{d,t} = \alpha_t + \gamma_t \ln r_{d,t} + \varepsilon_{d,t}, \qquad r_{d,t} = \frac{V_d(0 \to t)}{\bar{V}_d(0 \to t)} \tag{9}$$

Nothing from the future enters either version.

A simulated random walk with realistic volatility

To separate what the method does from what the market does, the whole pipeline (readings, forecast, probabilities, measures) is run on 20,000 simulated sessions in which the truth is a driftless Brownian path whose volatility is as uncertain as in the data. The volatility of each minute combines a persistent daily level, WIN's measured intraday pattern and a within-day factor that the morning cannot forecast, with dispersions set to reproduce the standard deviation of log daily realised variance (0.55) and of the log ratio of the true remaining variance to its forecast (0.66). Appendix B gives the details. We call the measures on these sessions the method's own values: a gap or a slope in the data is read against them, not against zero or one.

Measures

For readings $i = 1, ..., N$ with outcome Y_i, stated probability $\hat{P}_i$ and reading time $t(i)$:

- **Level** (calibration in the large): the gap, the share of readings where the range held minus the average probability.
- **Calibration slope** (Cox, 1958): the coefficient β of a logistic recalibration with a separate intercept for each reading time. $\beta = 1$ for calibrated probabilities; $\beta < 1$ means they are too extreme. Being on the log-odds scale, the slope is driven by the tails.

$$\text{Gap} = \frac{1}{N}\sum_{i=1}^{N}\left(Y_i - \hat{P}_i\right), \qquad \text{logit}\,\Pr(Y_i = 1) = \alpha_{t(i)} + \beta\,\text{logit}\,\hat{P}_i \tag{10}$$

- **Brier score** and its reliability term over the ten tenths k of the readings, with N_k readings, average probability $\bar{P}_k$ and share held $\bar{Y}_k$ (Murphy, 1973; Gneiting & Raftery, 2007):

$$\text{BS} = \frac{1}{N}\sum_{i=1}^{N}\left(\hat{P}_i - Y_i\right)^2, \qquad \text{REL} = \frac{1}{N}\sum_{k=1}^{10} N_k\left(\bar{P}_k - \bar{Y}_k\right)^2 \tag{11}$$

- **Ranking skill at the same time of day**: the area under the ROC curve within each reading time t, with N_t readings, averaged with weights equal to the number of readings (ties count one half). A forecast that knows only the clock scores 0.50. Ranking is not evidence for the random-walk law; the law is tested by calibration.

$$\text{AUC}_{\text{st}} = \sum_{t}\frac{N_t}{N}\,\Pr\left(\hat{P}_i > \hat{P}_j \mid Y_i = 1,\ Y_j = 0,\ t(i) = t(j) = t\right) \tag{12}$$

The fourteen readings of a session overlap, and the volatility memory makes forecast errors persist across sessions. Every confidence interval is therefore a percentile interval from a moving-block bootstrap over sessions in date order, with blocks of 20 sessions (Künsch, 1989; 1,000 resamples). Comparisons of squared error use a Diebold and Mariano (1995) statistic on the per-session mean loss differences ℓ (n sessions, mean $\bar{\ell}$, autocovariances $\hat{\gamma}_k$) with a Newey and West (1987) long-run variance over 20 lags, and two-sided p-values:

$$\text{DM} = \frac{\bar{\ell}}{\sqrt{\hat{\Omega}/n}}, \qquad \hat{\Omega} = \hat{\gamma}_0 + 2\sum_{k=1}^{20}\left(1 - \frac{k}{21}\right)\hat{\gamma}_k \tag{13}$$

What was registered in advance

Before measuring, the author wrote down four tests with their pass marks and a prediction each, in a private register (dated text files in a version-controlled repository; not a public registry, so the dates cannot be verified independently): the fitted model against the random walk on WIN; a lead on WDO; the hourly version on three other markets; and the calculation on Bitcoin and Ether (Appendix A). The same markets and years had been used in earlier work by the author, so the tests are confirmatory only in the narrow sense that their pass marks were fixed before the measurement. Everything else in this paper is exploratory.

The register counts every run of a test's script as a try. The registered WIN comparison, the paper's main test, took 2 counted tries: one count of the readings and one measurement, each run once. The WDO lead and the hourly test took 2 tries each, and the coins 5 (one count, and one run of the measurement whose three extra readings, shown but not judged, were counted as tries). The forward test has taken none. No registered test stated a p-value, so the register gives none of them a false-discovery q across the author's other tests; the after-registration comparisons of Table 2 are reported with their own multiplicity (Results).

Results

Level and calibration

Table 1 gives the three versions of the probability beside the same measures on simulated random walks. In level, the plug-in probability is close to the outcome: it says 56.1% on WIN and the range held in 55.5% of readings, a gap of −0.6 points (confidence interval −2.4 to +1.5); on WDO it says 58.2% against 59.0%, +0.7 (−0.8 to +2.6). The method's own gap on the simulated random walk is −1.4 (−1.8 to −1.1), and −0.8 and −0.7 in simulations calibrated to each market with its short sessions. Read against these, both markets hold their range about 0.2 to 2.1 points more often than the procedure shows under a random walk. That apparent excess is explained by the forecast's bias: in the data the remaining variance averaged 8% less than its forecast (mean log ratio −0.08), against 3–4% more in the simulations, and an overstated variance understates the probability that the range holds. When the simulated forecast is given the data's mean bias, the simulated plug-in gap becomes +1.5, and against that reference WIN's plug-in level is 2.1 points below the random walk's and WDO's 0.8 below, the same direction as the true-variance results below.

Table 1. Level, calibration and ranking, all sessions (five-minute realised variance; allowance a fifth of the range)

	Readings (sessions)	Range held	Version	Says	Gap (95% CI)	Calibration slope (95% CI)	Brier score (reliability)	Same-time AUC (s.e.)
WIN	16,457 (1,180)	55.5%	plug-in	56.1%	−0.6 (−2.4 to +1.5)	0.58 (0.45 to 0.71)	0.1661 (0.0012)	0.697 (0.013)
			predictive	57.1%	−1.6 (−3.4 to +0.4)	0.73 (0.58 to 0.89)	0.1651 (0.0007)	0.695 (0.013)
			predictive, conditional	57.0%	−1.4 (−3.3 to +0.5)	0.77 (0.62 to 0.90)		
			true variance	57.4%	−1.9 (−3.5 to −0.3)	1.11 (1.00 to 1.24)	0.1317 (0.0007)	0.838
WDO	16,492 (1,178)	59.0%	plug-in	58.2%	+0.7 (−0.8 to +2.6)	0.58 (0.48 to 0.65)	0.1636 (0.0009)	0.679 (0.011)
			predictive	59.1%	−0.1 (−1.6 to +1.7)	0.74 (0.62 to 0.83)	0.1622 (0.0002)	0.677 (0.011)
			predictive, conditional	58.8%	+0.4 (−1.3 to +2.1)	0.78 (0.68 to 0.87)		
			true variance	59.5%	−0.6 (−1.8 to +0.8)	1.12 (1.00 to 1.23)	0.1297 (0.0002)	0.828
Simulated random walk (reference)	279,160 (19,940)	51.3%	plug-in		−1.4 (−1.8 to −1.1)	0.67 (0.65 to 0.69)	0.1699 (0.0008)	0.690
			predictive		−0.5 (−0.8 to −0.1)	0.85 (0.83 to 0.88)	0.1689 (0.0000)	0.688
			predictive, conditional		−0.3 (−0.8 to +0.1)	0.94 (0.91 to 0.96)		

			true variance	−0.3 (−0.7 to −0.0)	0.99 (0.97 to 1.01)	0.1433 (0.0000)	0.811
Simulated, WIN-calibrated, short sessions	167,160 (11,940)	52.6%	plug-in / true variance	−0.8 / +0.3	0.66 / 1.01		0.689
Simulated, WDO-calibrated, short sessions	167,160 (11,940)	51.9%	plug-in / true variance	−0.7 / +0.2	0.69 / 1.00		0.696
Simulated, bias-matched forecast	279,160 (19,940)	51.3%	plug-in	+1.5	0.66		0.690

Gaps in percentage points (held minus says). CIs: moving-block bootstrap over sessions, blocks of 20, 1,000 resamples (400 for the conditional version; simulations 100–200, which measure simulation noise). The conditional version needs its own 20 sessions of history, so its gap is not exactly the held share minus its average. The true-variance version is not a forecast; its AUC is shown for completeness. The market-calibrated and bias-matched simulations are described in Appendix B.

The plug-in probabilities are too extreme: the calibration slope is 0.58 on both markets. This is mostly a property of the method. On the simulated random walk, where the formula is exactly right and only the volatility is uncertain, the plug-in slope is 0.66 to 0.69: plugging a point forecast of an uncertain variance into a function as nonlinear as P stretches the probabilities towards zero and one. Averaging over the forecast's error raises the slope to 0.73 and 0.74 in the data and to 0.85 in the simulation. The error is not independent of the reading's state: a morning that happened to be active has both a larger range and an overstated forecast (the log error falls by about 0.5 per unit of the log scaling ratio, in the data and in the simulation). Conditioning on it raises the simulated slope to 0.94 (0.91 to 0.96) and the data's to 0.77 (WIN, 0.62 to 0.90) and 0.78 (WDO, 0.68 to 0.87). The data remain about 0.16 of slope below the simulation. Measured directly, with the uncertainty of both the data and the simulation bootstrapped, the difference is −0.16 on WIN (−0.30 to −0.01) and −0.15 on WDO (−0.28 to −0.05); for the plug-in version it is −0.09 on both (WIN −0.21 to +0.04, WDO −0.20 to −0.00) (`revision/round5.py`). That residual is either a departure of prices from the random walk or a feature of the data's forecast error that the simulation lacks (its mean, its shape, its dependence on other features of the morning), and the paper cannot say which. Two richer descriptions of the error do not close it: replacing the normal by the error's own past residuals, or letting its mean and spread also depend on whether recent volatility is high or low, leaves the slopes at 0.77 to 0.80 on WIN and 0.78 on WDO, against 0.93 to 0.94 in the simulation (online appendix S3). These two alternatives do not explain the residual; other shapes, regimes or features of the morning that the simulation

lacks might. A feedback from the price to its own volatility, which the reference simulation lacks, was tested separately and does not (below). That weakens the second reading without excluding it.

Figure 1 shows why a slope well below one coexists with probabilities that look close to right. In the fourth to sixth tenths the observed share lies within two points of the stated probability. The lowest tenth on WIN says 1.4% and the range held 4.4% of the time; the second-highest says 92.4% and it held 86.7%. The predictive version removes most of the lowest-tenth gap (4.0% said, 3.9% held) and part of the upper one (90.7% against 86.0%). The Brier score's reliability term is 0.0012 or less of a score of about 0.165 in every version. By period, the plug-in gap moves within its intervals except for WDO in 2025 (+4.0, 0.5 to 6.7; +0.2 with the true variance; online appendix S1).

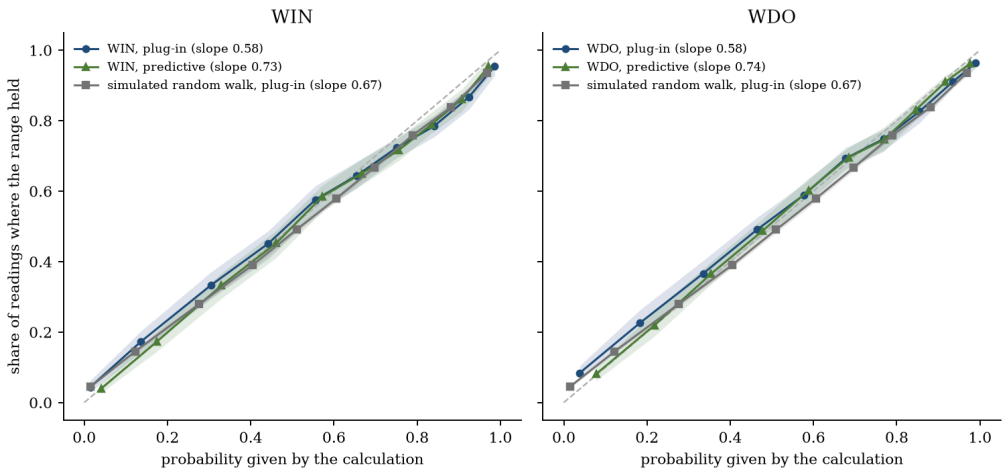

Figure 1. Probability given by the calculation against the share of readings where the range held, by tenths of the readings, with 95% moving-block bootstrap bands. The plug-in and predictive versions on the data, and the plug-in version on the simulated random walk (calibrated to WIN). Each tenth of the data holds 1,645–1,650 readings.

The Brownian law with the true remaining variance

With the variance the rest of the session actually had, forecast error is removed and the comparison bears more directly on the Brownian law. It is not an unconditional test of that law. The realised variance is measured on the same path whose extremes are scored, so the two are dependent, and five-minute sampling misses part of the variance within each interval. What the Brownian law implies for this comparison is therefore read from the simulations, which apply the same procedure to Brownian paths: the comparison is a diagnostic calibrated by simulation, and it is only as good as the simulation's account of the joint behaviour of realised variance and extremes. On the simulations this version is nearly unbiased (gaps −0.3 to +0.3, slopes 0.99 to 1.01). On WDO it is close to that (gap −0.6, −1.8 to +0.8; 0.8 points beyond the WDO-calibrated simulation). On WIN the range is extended more often than the law implies: the gap is −1.9 (−3.5 to −0.3), 1.6 points beyond the reference simulation and 2.3 beyond the WIN-calibrated one. Its interval excludes zero with bootstrap blocks of 10, 20 and 40 sessions, it is negative in every period and with the sessions of a single closing time (online appendix S1, S4), and the bias-matched plug-in

comparison points the same way (on the same readings, so not independent evidence). The true-variance calibration slope is 1.11 on WIN (1.00 to 1.24) and 1.12 on WDO (0.998 to 1.23), against 0.99 to 1.01 in the simulations; the direct differences from the reference simulation are +0.12 on both markets (+0.01 to +0.24 and +0.00 to +0.24).

We treat the WIN departure as tentative. It is significant at about two standard errors in one specification; it is absent at allowances of zero and a tenth and with one- and two-minute variance (online appendix S3, S4), not significant in 2025, absent on WDO, and one of many comparisons made after the registered tests.

Figure 2 shows it by time of day. The WIN gap is −2.8 points (−4.3 to −1.3) at the four reading times from 15:00 and −1.6 (−3.3 to +0.1) before; the late-minus-early contrast, chosen after seeing the figure, is −1.2 (−2.7 to +0.3). The departure is spread through the day, somewhat larger late. In the reference simulation the true-variance gap lies between −0.5 and −0.2 points at every reading time.

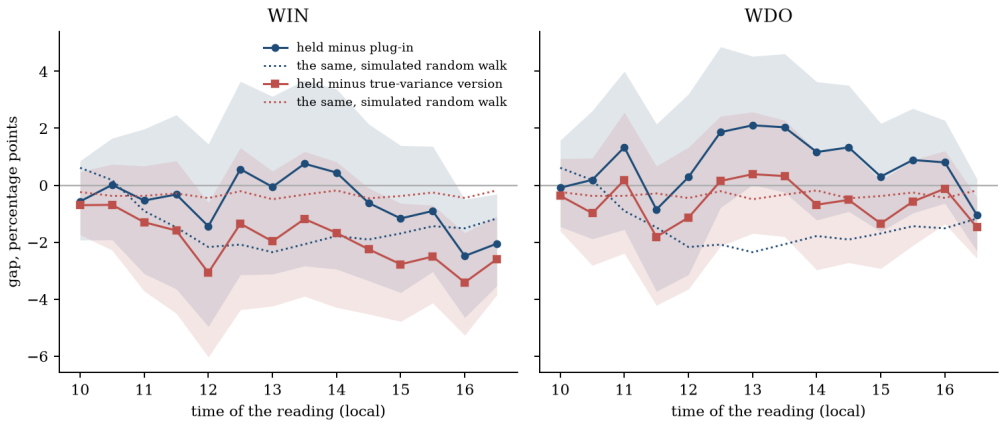

Figure 2. Gap between the share of readings where the range held and the probability, by time of the reading: plug-in (blue) and true-variance (red) versions on the data, solid with 95% moving-block bootstrap bands, and the same on the simulated random walk calibrated to WIN, dotted.

The size of extensions points the same way. Under the Brownian law with the true variance, the distribution function of a reading's extension is G_i, which has an atom at zero (no extension). Its probability integral transform, the standard check of a forecast distribution (Diebold, Gunther & Tay, 1998), is randomised on the atom (Czado, Gneiting & Held, 2009, discuss this device and a non-randomised alternative for discrete outcomes): U_i, with V_i uniform on [0, 1], is uniform under the law, and so, given that the range was extended at all, is U^+_i:

$$G_i(x) = P\left(a_i^*,\ b_i^*,\ x/s_i^*\right), \qquad U_i = G_i(\mathrm{ext}_i)\ \text{ if }\ \mathrm{ext}_i > 0, \qquad U_i = V_i\, G_i(0)\ \text{ if }\ \mathrm{ext}_i = 0 \tag{14}$$

$$U_i^+ = \frac{U_i - G_i(0)}{1 - G_i(0)} \quad \text{given}\ \ \mathrm{ext}_i > 0 \tag{15}$$

On the simulations the shares of U^+ in its top and bottom tenths lie between 9.2 and 10.1%. On WIN, given that the range was extended, 13.0% of the transforms fall in the top tenth of

the conditional law (11.3 to 15.0; simulated 10.0): when the range is extended, it tends to be extended somewhat further than a Brownian path with the same variance would extend it. On WDO the top share is 9.9%, as under the law, but small extensions are rarer than the law implies (6.7% in the bottom tenth, 5.9 to 7.7, against 9.8 to 10.1% simulated).

We did not find the cause (online appendix S2). It is not momentum. The variance ratio (Lo & MacKinlay, 1988) compares the squared 30-minute moves of the close with the squared 5-minute moves, pooled over sessions; it is one under a random walk and above one under momentum:

$$\mathrm{VR} = \frac{\sum_{d}\sum_{k}\left(p_{d,\,30k} - p_{d,\,30(k-1)}\right)^2}{\sum_{d}\sum_{j}\left(p_{d,\,5j} - p_{d,\,5(j-1)}\right)^2} \tag{16}$$

It is 0.99 on WIN (0.97 to 1.02) over the whole session, and from 15:00 to the close it is 0.93 (0.88 to 0.99), so late-session moves partly reverse, which should make the range hold more. Intraday momentum of the kind documented by Gao, Han, Li and Zhou (2018) and Baltussen, Da, Lammers and Martens (2021) does not explain the gap. Nor is it a feedback from the price to its own volatility, which the reference simulation, drawing volatility independently of the price, lacks by construction. Volatility that rises after down moves, or after a new high or low, was simulated with its strength set from statistics of WIN's volatility alone (online appendix S2). In the data both statistics are close to zero, and matched to them the simulations behave as the reference does. Even far stronger feedbacks raise the conditional slope (to 0.97) rather than lower it, and strong breakout feedback makes the range hold more often with the true variance (+1.0 point), the opposite of WIN's −1.9. Fat tails within five-minute intervals and behaviour around B3's closing call remain untested candidates.

Ranking, and fitted models

Within each time of day the plug-in probability ranks readings with an area under the ROC curve of 0.697 on WIN (standard error 0.013) and 0.679 on WDO (0.011); the simulated random walk gives 0.690.

The registered test compared the probability, with variance from one-minute closes, against a logistic model of six facts of the day, fitted once on 2021–2024 and frozen. Its pass mark was a gain of at least 0.03 in both 2025 and 2026; the model gained +0.002 and −0.030, and the test failed. That is a failure to show the model better, not a demonstration that it adds nothing. Table 2 gives the gains with intervals for three competitors: the logistic model; an encompassing regression of the outcome on the probability's log-odds and the six facts; and a gradient-boosted learner on the log-odds, the six facts, the clock and three momentum variables. Each competitor's settings were chosen by blocked validation inside 2021–2024 only (validating on 2023 and 2024 by log-loss), refitted on 2021-11 to 2024 and read once on 2025 and 2026, with one-minute and with five-minute variance.

Table 2. Same-time AUC gain over the random-walk probability, and squared error, on the two unseen years (competitors tuned inside 2021–2024)

Market, variance	Year	Random walk AUC	Logistic: gain (95% CI)	Encompassing: gain (95% CI)	Encompassing: Brier ratio (p)	Gradient boosting: gain (95% CI)	Gradient boosting: Brier ratio (p)
WIN, 1 min (registered)	2025	0.685	+0.002 (−0.036 to +0.044)	+0.009 (−0.009 to +0.026)	0.976 (0.06)	+0.002 (−0.019 to +0.017)	1.003 (0.86)
	2026	0.735	−0.030 (−0.067 to +0.019)	+0.002 (−0.009 to +0.018)	1.004 (0.85)	+0.010 (−0.013 to +0.027)	1.002 (0.91)
WIN, 5 min	2025	0.650	+0.037 (−0.009 to +0.080)	+0.027 (+0.007 to +0.047)	0.958 (0.04)	+0.027 (+0.008 to +0.050)	0.984 (0.42)
	2026	0.680	+0.025 (−0.023 to +0.098)	+0.023 (+0.006 to +0.050)	0.969 (0.25)	+0.029 (+0.001 to +0.056)	0.978 (0.39)
WDO, 1 min	2025	0.710	−0.019 (−0.045 to +0.015)	+0.007 (−0.010 to +0.022)	0.921 (0.01)	+0.004 (−0.014 to +0.028)	0.932 (0.007)
	2026	0.683	−0.042 (−0.077 to +0.016)	−0.002 (−0.028 to +0.027)	0.949 (0.02)	+0.016 (−0.011 to +0.052)	0.940 (0.004)
WDO, 5 min	2025	0.681	+0.009 (−0.023 to +0.043)	+0.021 (+0.000 to +0.044)	0.964 (0.01)	+0.025 (−0.001 to +0.045)	0.966 (0.05)
	2026	0.681	−0.040 (−0.104 to +0.025)	−0.005 (−0.045 to +0.048)	0.989 (0.71)	+0.020 (−0.019 to +0.066)	0.971 (0.26)

Brier ratio: competitor over random walk (below one = competitor better); p: two-sided Diebold–Mariano test with a Newey–West variance (20 lags). Every row uses the exact probability and refitted competitors, so the registered row's logistic gains are close to, not computed in the same way as, the registered test's. Validation chose the edge of the grid for the logistic penalty on WIN ($C = 10$) and for the boosted learner in all four rows (learning rate 0.03, 7 leaves, 100 trees). The table reports 48 intervals and tests.

Three things follow. With the registered one-minute variance, no competitor shows a reliable gain in ranking on either market; the intervals do not exclude gains of up to about 0.04, so the evidence is "no detectable gain", not equivalence. On WDO the encompassing regression and the boosted learner lower squared error by 5 to 8% in both years (p from

0.004 to 0.024): the one-minute plug-in level there is biased by about five points, and a fitted intercept removes the bias without improving the ranking. Better calibration and better ranking are different gains, and this one is calibration only. Across the sixteen squared-error tests of the table, none survives a Holm correction at 5% (smallest adjusted p 0.065), so these gains are descriptive; their consistency across both years and both learners is what gives them weight. With five-minute variance the random-walk probability ranks worse, and both the encompassing regression and the boosted learner improve the ranking on WIN in both unseen years (+0.023 to +0.029, intervals excluding zero, not adjusted for the table's other comparisons): the probability is then not a sufficient summary of what the day has shown. This is an exploratory result. One-minute variance ranks better; five-minute variance is less biased in level.

Robustness

The online appendix gives the tables; we summarise them here.

Dependence (S1). With bootstrap blocks of 10 to 120 sessions, or whole calendar months resampled, the conclusions stand. The plug-in gap's interval includes zero on both markets, and the plug-in slope's lies below one. The WIN true-variance gap's interval lies below zero and WDO's includes zero. The conditional slope's interval lies below the simulated value except with 120-session blocks on WIN, where it reaches it.

The volatility estimator (S3). The level depends on the sampling interval by several points, in the data and in the simulation alike, so each interval is read against the simulation at the same interval. On that basis WIN's plug-in level is 0.4 to 2.1 points above the simulation at every interval from one to fifteen minutes, and WDO's 1.2 to 2.4 points from five minutes upwards. At one and two minutes WDO is 6.2 and 3.8 points above: the known inflation of realised variance by microstructure noise at fine sampling (Andersen, Bollerslev, Diebold & Labys, 2000; Aït-Sahalia, Mykland & Zhang, 2005; Hansen & Lunde, 2006; Bandi & Russell, 2008), which produced the spurious WDO finding of the first version. A subsampled five-minute estimator (Zhang, Mykland & Aït-Sahalia, 2005) changes the levels by less than a point.

The allowance, trading hours and data source (S4). At allowances of zero, a tenth and three tenths of the range, the true-variance gaps are within about one point of the simulation on WDO, and on WIN 1.2 to 1.6 points beyond it at a fifth and three tenths, but not at zero or a tenth. Restricting to sessions whose memory has a single closing time, or dropping the 2026 sessions whose memory straddles the change of data source, moves no gap by more than 0.3 points.

Power (S5). Simulated alternatives run through the same pipeline show that the level tests detect momentum or mean reversion strong enough to move the variance ratio of 30-minute to 5-minute moves to 1.17 or 0.88 with probability 0.83 to 0.99, and "sticky" extremes of the support-and-resistance kind (Osler, 2003) with certainty; the calibration slope has little power against any of them. At 80% power the smallest detectable shift of the true-variance gap is 2.2 points on WIN and 1.9 on WDO, so the WIN gap of −1.9 points is of a size these tests detect only part of the time.

Other markets

Bitcoin and Ether. The calculation ran unchanged on the two coins. In level it is within two points: on the 12,609 readings of this paper's analyses, the plug-in probability says 32.5% for BTC and 34.0% for ETH, and the range held in 34.3% and 34.4% of readings, gaps of +1.7 (+0.3 to +3.1) and +0.3 (−1.1 to +1.9) (`revision/other_markets.py`). On the 12,789 readings of the registered script, with the exact probability, the gaps are +1.9 (32.5% said, 34.4% held) and +0.3 (34.0%, 34.3%); with the registered 400-step routine they were +0.4 and −1.1, the "largest gap 1.1 points" of Appendix A being ETH's. But the probabilities are far more extreme than on WIN, with plug-in calibration slopes of 0.38 and 0.40 in this paper's analyses (predictive 0.68 and 0.67); in the surest third (registered script, exact probability) it says 73.4% and 74.6% and the range held 69.0% and 68.2%, so the registered mark of 75% failed. With the true variance the slopes are 1.55 and 1.48. We have no simulation calibrated to the coins, so we cannot say how much of this is the method. The ranking skill is 0.681 and 0.684, and a fitted model adds to it in 2023 (0.735 and 0.791 against 0.679 and 0.685): evidence against a general claim that the random-walk probability exhausts the information in the day's path.

Three index markets, hourly. On one year of hourly bars, with volatility from each bar's high and low (Parkinson, 1980), and net of the method's own gap, the range held 5.3 points more often than stated on the E-mini S&P 500 (0.8 to 9.9), 4.5 on the E-mini Nasdaq-100 (−0.1 to 8.9) and 2.3 on the DAX (−1.8 to 6.4). The registered mark was 5 points; the S&P's registered figure, under the earlier routine, was 5.6 (failed), and the 5.3 here is exploratory. One year of free hourly bars does not settle whether US index futures hold their range more often than a random walk allows. Details are in the online appendix (S6).

Discussion

What the calculation is good for. It is a benchmark. A model, indicator or rule of thumb that claims to know whether "the range is in" should be compared with this calculation, not with the clock: a logistic model of six facts of the day, which ranks readings well against the clock, shows no reliable gain over it in the registered version. The claim this supports is narrower than "the random walk contains everything". The calculation captures most of what the day's path says about whether the range holds, but not necessarily all of it: with five-minute variance, fitted models add about 0.025 of ranking skill on WIN in both held-out years (an exploratory result), and more on the coins. A competing model's worth is its gain over this benchmark, out of sample, not its accuracy alone. Given a volatility forecast, it also estimates how likely a level beyond the day's extreme is to be reached before the close, which bears on where stops and targets sit, and, through the double no-touch connection, it is a reference for how far a Brownian pricing assumption is from realised frequencies at the intraday horizon. A comparison with option-implied double no-touch probabilities, which we could not make, would be the natural next step.

What it is not. It is not a trading signal: a driftless random walk implies nothing about direction. And it is not exact: its probabilities are too extreme as a forecast, mostly for a reason internal to the method (with a residual whose source is open), and on WIN, read

against simulations whose volatility does not respond to the price, the range is extended more often and somewhat further than the Brownian law implies.

Limitations

- The literature search was several rounds of web search. Mok et al. (2000) and Copeland et al. (2004) have no open version we could find and are described from their abstracts. A closer equivalent of this calculation may exist.
- The simulated random walks reproduce the dispersion of daily volatility and of the volatility-forecast error, and each market's long and short sessions, not all features: the forecast's mean bias is matched only in a separate variant, and their volatility does not respond to the price path (two such responses were simulated separately and do not change the conclusions). There is no simulation for the coins.
- The level result depends on the volatility estimator by several points; results are read against the simulation at the same interval.
- The departure on WIN with the true remaining variance is tentative, and measured, not explained.
- The fitted competitors are three; others might find information the probability misses, as the five-minute results and the coins suggest.
- The minute data come from a broker's continuous series and, for WIN from December 2025, a charting platform. They agree with each other, and with bars built from trade-by-trade prints, where they overlap (2026 only), but were not compared with the exchange's own trade files, and only derived statistics can be shared.
- The three index markets rest on one year of hourly bars from a free source; the coins on 2020–2023.
- The readings overlap and sessions are dependent; the bootstrap block of 20 sessions is a choice. WIN and WDO trade the same hours and the same news and are not independent samples.
- Four tests were registered in advance, in a private register whose dates cannot be verified independently, on data the author had used before. Every other analysis is exploratory. No data after 9 October 2026 were used. A forward test of the main version, with its marks fixed, was registered on 17 October 2026 and will be judged once, on the first 250 later sessions of each market (Appendix A); it has no result yet.

Conclusion

Whether the day's range is already made can be calculated, with no fitted parameter, from where the price sits in the range and the volatility still expected, under a driftless random walk.

- On two Brazilian futures over five years, the calculation's average probability is within one point of the observed share, and within one to two points of what the same procedure shows on a simulated random walk.

- Its probabilities are too extreme as a forecast, and simulated random walks with realistic volatility show that most of this comes from plugging an uncertain volatility forecast into the formula; averaging over the forecast's error removes most of it in simulation and part of it in the data.
- With the true remaining volatility, read against simulations whose volatility does not respond to the price, WDO is consistent with the Brownian law in level; on WIN the range is extended more often and somewhat further than the law implies, a tentative departure.
- Tuned fitted models show no detectable gain in ranking over the registered version; with a noisier five-minute volatility estimate they add information on WIN.
- On Bitcoin and Ether the calculation is within two points in level and ranks as well as on WIN, but its probabilities are far more extreme and its registered calibration test failed.

Code and data availability

The paper's folder contains `replication/` (the re-run of the registered test's script, unchanged, which matched all nine of its recorded numbers; since 22 October 2026 it runs on a dated snapshot of the minute bars the script read, cut at 30 September 2026, the last day the registered test saw, and checked by its SHA-256, so that later additions to the data cannot move it) and `revision/` (the exact probability, the simulated sessions, every analysis of this version and the figures, with their outputs, and `datacheck.py`, the comparison of the data feeds), with the scripts of the registered tests. The search record is `literature.md`; the referee reports and our responses sit next to the paper. The Binance klines are public, so that part can be reproduced by anyone. The WIN and WDO minute data and the hourly index bars are vendor data; for these we share code and derived statistics only. The code runs on any one-minute or hourly series with high, low and close, and the simulated sessions (`sim.py`) are synthetic data on which the whole pipeline runs without vendor data.

Disclosure

The author trades these markets, including WIN and WDO, and the research was carried out for a private trading fund owned by the author. This is research, not investment advice. Four tests were registered in advance in a private register (Appendix A); the registered comparison failed, and so did the test on the coins (the register records the comparison as "KILLED", meaning failed, at the evidence level it calls "in-sample", meaning that the years had been used before); every other analysis is exploratory. Parts of the analysis and of the writing were done with an AI assistant (Claude, by Anthropic); every number comes from the scripts in the paper's folder. The four referee reports beside the paper were written by a separate session of the same assistant, given the paper and its folder (the scripts and their outputs); they are not peer review. This version answers them; it has been revised since the last report and has not been refereed again.

References

References follow APA style (7th edition). Volume, issue, page and DOI details were checked against the Crossref index on 13, 16, 17 and 22 October 2026; the four works not indexed there were checked on the article itself or on publisher and author pages, as marked.

Acar, E., Lequeux, P., & Ritz, S. (1996). Timing the highs and lows of the day. *LIFFE Equity Products Review*, 2nd quarter, 1996. [Not indexed in Crossref; checked on the article.]

Aït-Sahalia, Y., Mykland, P. A., & Zhang, L. (2005). How often to sample a continuous-time process in the presence of market microstructure noise. *The Review of Financial Studies, 18*(2), 351–416. https://doi.org/10.1093/rfs/hhi016

Alizadeh, S., Brandt, M. W., & Diebold, F. X. (2002). Range-based estimation of stochastic volatility models. *The Journal of Finance, 57*(3), 1047–1091. https://doi.org/10.1111/1540-6261.00454

Andersen, T. G., & Bollerslev, T. (1997). Intraday periodicity and volatility persistence in financial markets. *Journal of Empirical Finance, 4*(2–3), 115–158. https://doi.org/10.1016/S0927-5398(97)00004-2

Andersen, T. G., Bollerslev, T., Diebold, F. X., & Labys, P. (2000). Great realisations. *Risk, 13*(3), 105–108. [Not indexed in Crossref; checked on an author's faculty listing; text not read.]

Andersen, T. G., Bollerslev, T., & Dobrev, D. (2007). No-arbitrage semi-martingale restrictions for continuous-time volatility models subject to leverage effects, jumps and i.i.d. noise: Theory and testable distributional implications. *Journal of Econometrics, 138*(1), 125–180. https://doi.org/10.1016/j.jeconom.2006.05.018

Baltussen, G., Da, Z., Lammers, S., & Martens, M. (2021). Hedging demand and market intraday momentum. *Journal of Financial Economics, 142*(1), 377–403. https://doi.org/10.1016/j.jfineco.2021.04.029

Bandi, F. M., & Russell, J. R. (2008). Microstructure noise, realized variance, and optimal sampling. *The Review of Economic Studies, 75*(2), 339–369. https://doi.org/10.1111/j.1467-937X.2008.00474.x

Broadie, M., Glasserman, P., & Kou, S. (1997). A continuity correction for discrete barrier options. *Mathematical Finance, 7*(4), 325–349. https://doi.org/10.1111/1467-9965.00035

Caporin, M., Ranaldo, A., & Santucci de Magistris, P. (2013). On the predictability of stock prices: A case for high and low prices. *Journal of Banking & Finance, 37*(12), 5132–5146. https://doi.org/10.1016/j.jbankfin.2013.05.024

Cheung, Y.-W. (2007). An empirical model of daily highs and lows. *International Journal of Finance & Economics, 12*(1), 1–20. https://doi.org/10.1002/ijfe.303

Chou, R. Y. (2005). Forecasting financial volatilities with extreme values: The conditional autoregressive range (CARR) model. *Journal of Money, Credit, and Banking, 37*(3), 561–582. https://doi.org/10.1353/mcb.2005.0027

Christensen, K., & Podolskij, M. (2007). Realized range-based estimation of integrated variance. *Journal of Econometrics, 141*(2), 323–349. https://doi.org/10.1016/j.jeconom.2006.06.012

Copeland, L., Lam, K., & Jones, S.-A. (2004). The index futures markets: Is screen trading more efficient? *Journal of Futures Markets, 24*(4), 337–357. https://doi.org/10.1002/fut.10119

Cox, D. R. (1958). Two further applications of a model for binary regression. *Biometrika, 45*(3–4), 562–565. https://doi.org/10.1093/biomet/45.3-4.562

Czado, C., Gneiting, T., & Held, L. (2009). Predictive model assessment for count data. *Biometrics, 65*(4), 1254–1261. https://doi.org/10.1111/j.1541-0420.2009.01191.x

Dambis, K. E. (1965). On the decomposition of continuous submartingales. *Theory of Probability and Its Applications, 10*(3), 401–410. https://doi.org/10.1137/1110048

Diebold, F. X., Gunther, T. A., & Tay, A. S. (1998). Evaluating density forecasts with applications to financial risk management. *International Economic Review, 39*(4), 863–883. https://doi.org/10.2307/2527342 [Last page not seen in Crossref.]

Diebold, F. X., & Mariano, R. S. (1995). Comparing predictive accuracy. *Journal of Business & Economic Statistics, 13*(3), 253–263. https://doi.org/10.1080/07350015.1995.10524599

Drinka, T. P., & Ptasienski, S. M. (1989). Understanding Market Profile, part 6. *Technical Analysis of Stocks & Commodities, 7*(4), 121–123. [Not indexed in Crossref; checked on the publisher's page; text not read.]

du Toit, J., & Peskir, G. (2009). Selling a stock at the ultimate maximum. *The Annals of Applied Probability, 19*(3), 983–1014. https://doi.org/10.1214/08-AAP566

Dubins, L. E., & Schwarz, G. (1965). On continuous martingales. *Proceedings of the National Academy of Sciences, 53*(5), 913–916. https://doi.org/10.1073/pnas.53.5.913

Feller, W. (1951). The asymptotic distribution of the range of sums of independent random variables. *The Annals of Mathematical Statistics, 22*(3), 427–432. https://doi.org/10.1214/aoms/1177729589

Gao, L., Han, Y., Li, S. Z., & Zhou, G. (2018). Market intraday momentum. *Journal of Financial Economics, 129*(2), 394–414. https://doi.org/10.1016/j.jfineco.2018.05.009

Garman, M. B., & Klass, M. J. (1980). On the estimation of security price volatilities from historical data. *The Journal of Business, 53*(1), 67–78. https://doi.org/10.1086/296072 [Last page not seen in Crossref.]

Gneiting, T., & Raftery, A. E. (2007). Strictly proper scoring rules, prediction, and estimation. *Journal of the American Statistical Association, 102*(477), 359–378. https://doi.org/10.1198/016214506000001437

Graversen, S. E., Peskir, G., & Shiryaev, A. N. (2001). Stopping Brownian motion without anticipation as close as possible to its ultimate maximum. *Theory of Probability & Its Applications, 45*(1), 41–50. https://doi.org/10.1137/S0040585X97978075 (Russian original 2000.)

Hadjiliadis, O., & Večeř, J. (2006). Drawdowns preceding rallies in the Brownian motion model. *Quantitative Finance, 6*(5), 403–409. https://doi.org/10.1080/14697680600764227

Hansen, P. R., & Lunde, A. (2006). Realized variance and market microstructure noise. *Journal of Business & Economic Statistics, 24*(2), 127–161. https://doi.org/10.1198/073500106000000071

Holmberg, U., Lönnbark, C., & Lundström, C. (2013). Assessing the profitability of intraday opening range breakout strategies. *Finance Research Letters, 10*(1), 27–33. https://doi.org/10.1016/j.frl.2012.09.001

Hui, C. H. (1996). One-touch double barrier binary option values. *Applied Financial Economics, 6*(4), 343–346. https://doi.org/10.1080/096031096334141

Kunitomo, N., & Ikeda, M. (1992). Pricing options with curved boundaries. *Mathematical Finance, 2*(4), 275–298. https://doi.org/10.1111/j.1467-9965.1992.tb00033.x

Künsch, H. R. (1989). The jackknife and the bootstrap for general stationary observations. *The Annals of Statistics, 17*(3), 1217–1241. https://doi.org/10.1214/aos/1176347265

Lehoczky, J. P. (1977). Formulas for stopped diffusion processes with stopping times based on the maximum. *The Annals of Probability, 5*(4). https://doi.org/10.1214/aop/1176995770 [Pages not

seen.]

Lo, A. W., & MacKinlay, A. C. (1988). Stock market prices do not follow random walks: Evidence from a simple specification test. *The Review of Financial Studies, 1*(1), 41–66. https://doi.org/10.1093/rfs/1.1.41

Martens, M., & van Dijk, D. (2007). Measuring volatility with the realized range. *Journal of Econometrics, 138*(1), 181–207. https://doi.org/10.1016/j.jeconom.2006.05.019

Mok, D. M. Y., Lam, K., & Li, W. (2000). Using daily high/low time to test for intraday random walk in two index futures markets. *Review of Quantitative Finance and Accounting, 14*(4), 381–397. https://doi.org/10.1023/A:1008323926951

Murphy, A. H. (1973). A new vector partition of the probability score. *Journal of Applied Meteorology, 12*(4), 595–600. https://doi.org/10.1175/1520-0450(1973)012<0595:ANVPOT>2.0.CO;2

Newey, W. K., & West, K. D. (1987). A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. *Econometrica, 55*(3), 703–708. https://doi.org/10.2307/1913610 [Last page not seen in Crossref.]

Osler, C. L. (2000). Support for resistance: Technical analysis and intraday exchange rates. *Federal Reserve Bank of New York Economic Policy Review, 6*(2), 53–65. [Not indexed in Crossref; pages from the author's own citation (Staff Report 125).]

Osler, C. L. (2003). Currency orders and exchange rate dynamics: An explanation for the predictive success of technical analysis. *The Journal of Finance, 58*(5), 1791–1819. https://doi.org/10.1111/1540-6261.00588

Panas, E. (2005). Generalized beta distributions for describing and analysing intraday stock market data: Testing the U-shape pattern. *Applied Economics, 37*(2), 191–199. https://doi.org/10.1080/00036840412331315006

Parkinson, M. (1980). The extreme value method for estimating the variance of the rate of return. *The Journal of Business, 53*(1), 61–65. https://doi.org/10.1086/296071 [Last page not seen in Crossref.]

Peters, R. T., & de Vilder, R. G. (2006). Testing the continuous semimartingale hypothesis for the S&P 500. *Journal of Business & Economic Statistics, 24*(4), 444–454. https://doi.org/10.1198/073500106000000341

Rogers, L. C. G., & Satchell, S. E. (1991). Estimating variance from high, low and closing prices. *The Annals of Applied Probability, 1*(4), 504–512. https://doi.org/10.1214/aoap/1177005835

Taylor, H. M. (1975). A stopped Brownian motion formula. *The Annals of Probability, 3*(2). https://doi.org/10.1214/aop/1176996395 [Pages not seen.]

Zarattini, C., & Aziz, A. (2023). *Can day trading really be profitable?* SSRN Working Paper 4416622.

Zhang, L., Mykland, P. A., & Aït-Sahalia, Y. (2005). A tale of two time scales: Determining integrated volatility with noisy high-frequency data. *Journal of the American Statistical Association, 100*(472), 1394–1411. https://doi.org/10.1198/016214505000000169

Appendix A. The tests registered in advance

Each test below was written down, with its pass marks and a prediction, before its measurement, and measured once (the forward test will be). The register is private (see Method).

Test	Registered	Pass marks	Outcome
The fitted model against the random walk, WIN	2026-10-12 19:02	the fitted model's same-time ranking skill exceeds the random walk's by at least 0.03 in 2025 and in 2026 (one-minute variance, 400-step routine)	Failed: +0.002 and −0.030. The prediction written with the test: "the random walk gets most of it"
The WDO lead	2026-10-12 19:54	with five-minute variance, a gap of at least 4 points with its 95% interval above zero, and the same with the true remaining variance, in 2025 and in 2026 (six marks)	Failed all six
The hourly version on three index markets	2026-10-12 20:06 (pass marks amended at 20:07, before any market data, to read every gap net of the method's own)	level gap within 5 points and ranking skill of at least 0.60 on each market; the hourly version within 5 points of the five-minute one on WIN (seven marks)	Six passed; the E-mini S&P 500 missed the level mark (5.6)
Bitcoin and Ether	2026-10-13 11:21	same-time ranking skill of at least 0.65 on both coins; level gap of at most 5 points; in the third of readings where the calculation is surest, the range held at least 75% of the time (all on 2020–2023)	Ranking (0.681, 0.684) and level (largest gap 1.1 points, ETH, under the registered 400-step routine) passed; the surest-third mark failed (69% for BTC, 68.3% for ETH)
Forward test of the main version, WIN and WDO	2026-10-17 11:08	on the first 250 sessions after 16 October 2026 of each market: plug-in level gap within 5 points, calibration slope below 0.9 and same-time ranking skill of at least 0.62, on both markets (six marks); the true-variance gap reported, not judged	Pending (about October 2027)

Before these tests, the author had registered and failed five trading rules that acted on the fitted model's confidence. They are listed for completeness; this paper makes no claim from them. A separately registered test of a one-hour question on WIN, and a version reported with it, are in the online appendix (S7).

Appendix B. The simulated sessions

Each simulated session has 565 one-minute candles (09:00 to 18:24, the long WIN session), each built from 24 Brownian sub-steps so that its high and low are close to the path's extremes. The standard deviation of the move in minute m of session d, which falls in half-hour block $k(m)$, is

$$\sigma_{d,m} = \exp(h_d)\,\sqrt{\pi_m}\,\exp\left(z_{d,k(m)}\right) \tag{17}$$

where π_m is WIN's measured average squared one-minute change in minute m of the session, normalised to mean one, and the daily level h and the within-day factor z are AR(1) processes:

$$h_d = 0.72\,h_{d-1} + \eta_d, \qquad \eta_d \sim N\left(0,\ 0.23^2\,(1-0.72^2)\right) \tag{18}$$

$$z_{d,k} = \tilde{z}_{d,k} - \tau^2, \qquad \tilde{z}_{d,k} = 0.5\,\tilde{z}_{d,k-1} + \zeta_{d,k}, \qquad \zeta_{d,k} \sim N\left(0,\ 0.75\,\tau^2\right), \qquad \tau = 0.32 \tag{19}$$

The coefficient 0.72 is measured on WIN; h has stationary standard deviation 0.23 and the stationary $\tilde{z}$ has standard deviation τ, and the shift by τ^2 makes $\mathrm{E}[\exp(2z)] = 1$. z is the part of the remaining variance that the morning cannot forecast. The two dispersions reproduce the standard deviation of log daily realised variance (0.55) and of the log forecast error (0.66). The same pipeline as for the data then runs unchanged.

The **market-calibrated** runs (12,000 sessions each) take each market's own dispersions from a small grid (WIN 0.23 and 0.32; WDO 0.26 and 0.32), intraday pattern, and long and short sessions in blocks that mimic B3's daylight-saving schedule. The **bias-matched** variant multiplies the reference simulation's forecast variance by $\exp(0.115)$, the difference between the data's mean log forecast error (-0.08) and the simulation's. The online appendix (S8) gives the full specification, the hourly version's simulation and the alternatives of the power analysis.