WORKING PAPER NO. 01 · NOVEMBER 2026 · VERSION 1.0

The Opening Gap and Intraday Breakouts in Nasdaq-100 Futures

Day Selection, Replication and Data Quality

Ludwig J. Marx ludwig@ludwigjmarx.com · ludwigjmarx.com

BOTTOM LINE

Trading the opening gap of the Nasdaq-100 does not work as a rule for every day: the gap direction adds little over a random one. The breakout that follows the open earns on wide trend days and loses on quiet ones. A rule that selects days by the direction and the volume of the opening minutes reaches a Sharpe ratio of 1.00 on QQQ minute bars and 16.8% a year with real orders in Micro Nasdaq-100 futures. It fails the registered hurdle (deflated Sharpe ratio 0.83 against 0.95). In the test sample 2015–2019, opened once, it lost money and did no better than a random direction.

KEY JUDGMENTS

Moderate confidence. The overnight gap is a weak guide to direction. It beat a random direction in 4 of 24 rules, all of them weak ones; the first minutes of the session give a better direction.

Moderate confidence. Breakout rules after the open lose 27 basis points a day in the quietest fifth of days and earn 37 in the widest. The result depends on which days are traded.

Moderate confidence. Research model and replication agree trade by trade (0.212 R on 620 common days) after three sources of difference were found.

Moderate confidence. The selected rule does not hold out of sample: it beat 99.8% of random-direction runs in 2020–2026 and 30% in 2015–2019.

ABSTRACT

I test whether the opening gap of the Nasdaq-100 can be traded within the day: enter in the direction of the gap when price leaves the opening range, and hold to the close. The rules were developed in R on vendor data. Every reported result is recomputed on QuantConnect, with real orders in Micro Nasdaq-100 futures. The gap direction adds little, and the breakout earns only on wide days. A rule that selects days by the first minutes earns 16.8% a year with real orders. It does not meet the registered hurdle and fails in the test sample 2015–2019 (Sharpe ratio −0.18).

Keywords: opening gap, opening range breakout, intraday, futures, day selection, replication, data quality, QuantConnect · JEL: G13, G14, C12, C58

Confidence levels. High: Holds out of sample, in both sub-periods and in the independent replication (R against QuantConnect). Moderate: Holds in sample and in the replication; weaker or not yet tested out of sample. Low: First finding, based on few cases or a single period. Not yet a basis for decisions.

Research for discussion. Not investment advice. Backtests are labelled as such and are not live track records.

Full rule and code are published with this paper.

1 Question

Index futures often open away from the prior close. A common trading idea holds that this gap shows where the day is going: if the market opens higher, buyers are in control and the move continues. A simple rule follows from it. Wait until price leaves the range of the first minutes in the direction of the gap, enter there with a stop on the other side of the range, and hold until the close.

Such rules are easy to describe and widely traded. Whether they earn money after costs is less clear, because most days do not trend. This paper asks: **does the direction of the opening gap carry into the session of the Nasdaq-100, and if not, what decides whether a breakout after the open pays?** I build the rules in R on QQQ minute bars, recompute every variant on QuantConnect and trade the final rule there with Micro Nasdaq-100 futures (MNQ). The paper reports the full path, including the rules and controls that failed.

2 Related literature

Opening range breakout. Zarattini and Aziz (2023) test an opening range breakout on QQQ from 2016 to 2023 that takes its direction from the first minutes of the session, cuts losses at the other side of the opening range and exits at the close. They report a Sharpe ratio of 1.12 before slippage. The candidate of this paper belongs to the same family. Their rule takes the direction of the first 5-minute candle, skips flat candles, places the stop at the other side of that candle, risks 1% of the account per trade with leverage capped at four times and sets a target of ten times the risk. Our candidate differs in the entry (a stop order at the edge of the range instead of the open of the next candle) and in the opening-volume filter. It is tested here with registration and deflation of all variants and replicated with real futures orders. Zarattini, Barbon and Aziz (2024a) restrict the same breakout to single stocks with unusually high volume on the day (“stocks in play”); the volume filter here is the index version of that idea.

Intraday momentum. Gao, Han, Li and Zhou (2017) show for the S&P 500 ETF that the return of the first half hour, measured from the prior close, predicts the return of the last half hour, more strongly on volatile days, on days with high volume and on days with macroeconomic news. Baltussen, Da, Lammers and Martens (2020) find the same pattern in more than 60 futures markets and link it to hedging demand of option market makers. Zarattini, Aziz and Barbon (2024b) trade intraday momentum on SPY with trailing stops. The day features in section 5 point the same way: the breakout earns on wide days, on days with high opening volume and on release days.

Overfitting. The deflated Sharpe ratio (Bailey and López de Prado 2014) corrects the best of many tested variants for the number of trials; its expected maximum under the null follows López de Prado and Bailey (2018). Bailey, Borwein, López de Prado and Zhu (2015) show how often the best variant in sample disappoints out of sample. Both motivate the registration and the test sample used here.

3 Data

Minute bars. QQQ, the ETF on the Nasdaq-100, serves as the research instrument. Financial Modeling Prep (FMP) supplied one-minute bars of the regular session from January 2015 to 20 November 2026: 2,990 days and 1.16 million bars, without duplicates. A request that contains no trading day returns the last available days instead of nothing,

so the download keeps only rows inside the requested window. FMP also supplied the US economic calendar (labour market, CPI, PCE, ISM, Federal Reserve).

Samples. Development uses 2 January 2020 to 20 November 2026. The years 2015 to 2019 form the test sample. They were to be opened once, after a rule had passed the development hurdle. No script read prices before 2020 except as warm-up for indicators. The candidate did not pass the hurdle. I opened the test sample on 22 November 2026 regardless, once and for the candidate only, with the pass criterion written down before the run.

QuantConnect. QuantConnect supplies its own QQQ minute bars and MNQ futures data. All reported results use them, for January 2020 to 20 November 2026; versions 1 to 3 in section 7 end in August 2026.

Code. The QuantConnect algorithms and the R scripts that produce every number and figure in this paper are published with it (see Appendix B). The price data are licensed and are not included.

4 Method

Registration. Each step was written down before its first run: rule family, measure and pass criterion. Changes were added as dated notes. A variant counter records every rule that was computed: 24 base variants, 18 filters, 8 stop rules and one data correction, 51 in total.

Computation. The rules were written and selected in R on the FMP data (the research model). Every variant was then recomputed on the QuantConnect backtest cluster with its QQQ minute bars, in one program that runs all 51 variants side by side at the trigger price. The same program, run on the FMP data, reproduces all 51 Sharpe ratios of the research model to four decimals. The final rule was also traded on QuantConnect with real orders in MNQ. Unless stated otherwise, numbers in this paper are QuantConnect results; section 7 compares them with the research model.

Execution and costs. A variant enters with a stop order at the edge of the opening range and fills at the trigger price. It exits at the stop or shortly before the close, with one trade per day. If the stop lies within reach in the minute of entry, it counts as hit. Costs are 0.01% of the price per round trip, which covers commission and one tick of slippage per side in MNQ. Each trade risks a fixed share of the account, with leverage capped. Real orders on QuantConnect use its own fill model and fees.

Measures. The main measure is the Sharpe ratio of daily returns after costs, with days without a trade counted as zero. At the trigger price it is reported without a risk-free rate. For real orders QuantConnect deducts the risk-free rate, which lowers its values. Results per trade are given in multiples of the initial risk (R), so a full stop is −1 R.

Hurdle. A rule passes development if three conditions hold: a Sharpe ratio of at least 1.0, a result better than 95% of 1,000 runs with random direction on the same days, and a deflated Sharpe ratio (Bailey and López de Prado 2014) of at least 0.95 over all counted variants. In the test sample it needs a Sharpe ratio of at least 0.5, a mean R per trade with a positive bootstrap lower bound, and a result better than 90% of random runs.

Controls. The backtester had to reproduce seven trades computed by hand, including stops through price gaps and short trades. In synthetic prices with a built-in trend in the gap direction it had to find a clear edge (Sharpe ratio 5.98), and in prices without that trend none (0.59, within noise).

5 Trading the gap

Base rules. The first registered step tested 24 breakout rules in the direction of the gap: three lengths of the opening range, two entry types, two stop levels, and with or without a move of the stop to break-even.

Base family. Opening range 5, 15 or 30 minutes from 9:30. Entry (a) stop order at the range edge in gap direction, (b) first 5-minute close beyond the edge, filled in the next minute; no entry after 11:30. Stop (a) opposite edge of the range, (b) middle of the range. Break-even after +1 R: yes or no. Exit at the stop or at 15:55. Risk 1% of the account per trade, leverage at most 4 times.

None of them passed. The best reached a Sharpe ratio of 0.67 with a largest loss of −16%. The gap direction beat a random direction at the registered level in 4 of 24 variants, all with Sharpe ratios of 0.40 or less; in the five best variants it did not (rank 0.66 to 0.80 against 0.95 required). A rule that simply takes the first break of either side did better than the gap rule (0.88 against 0.67 with the longest range). On the Nasdaq-100 in 2020–2026 the gap is a weak guide to direction. **Moderate confidence.**

The mean result per trade was positive in all 24 variants, but only by 0.01 to 0.09 R, which is too little after costs. The median trade was a full loss or close to zero. A breakout edge exists, but it is small and lives on a few large days.

Where the breakout pays. Sorting the days by their range after the close shows the pattern (Figure 1). In the quietest fifth of days the 24 variants lost 27 basis points a day in the median; in the widest fifth they earned 37. Every variant follows the same line. The breakout earns on trend days and pays small stops on all others. A day's range is known only after the close, so this is no rule. It shows where an edge has to come from: from information available at the open that separates wide days from quiet ones. **Moderate confidence.**

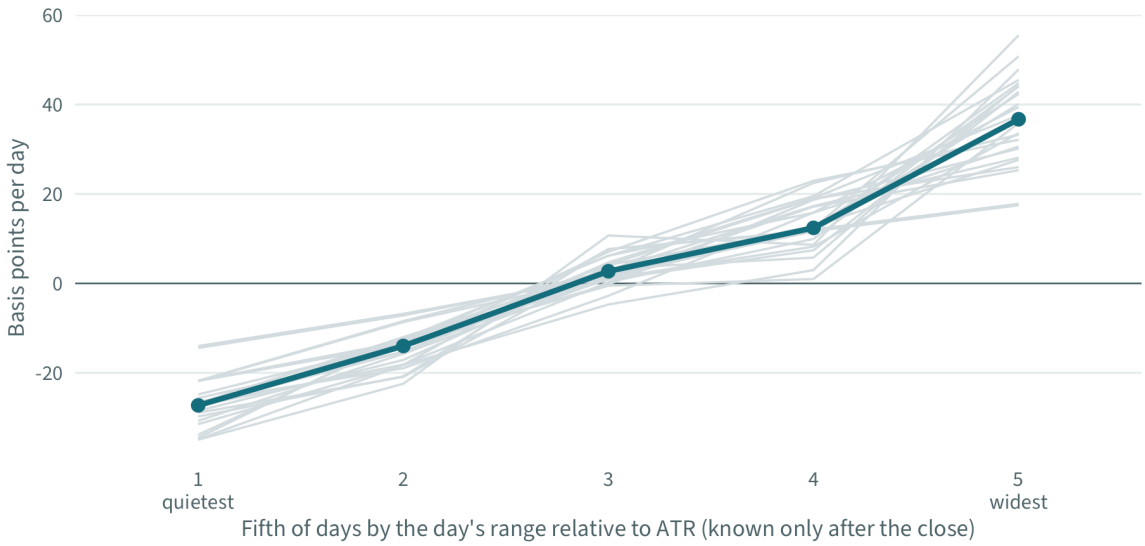


Figure 1 · Mean daily result of the 24 gap rules by fifth of days, sorted by the day's high-low range relative to its 14-day average true range. Grey: single variants; blue: median. · Source: QuantConnect QQQ minute bars, 2020–2026; author's calculations · As of 21 Nov 2026 · Backtest, net of costs

6 Choosing the day

Day features. For each of the 1,732 development days I recorded only information known at the end of the opening range: gap, volatility, width and direction of the opening range, opening volume and economic releases. The breakout did better when the first candle ran in the direction of the gap (gap rule +14 against −2 basis points a day), when opening volume was in the top fifth (first-break rule +18 against −7 in the bottom fifth) and on days with an 8:30 release (+15 against +4). It did worse on days of a Federal Reserve decision (first-break rule −51 basis points a day). The width of the first 15 minutes separated best (+47 against −19), but it is known only at 9:45 and was not used.

Filters. The second registered step combined three direction rules with six day filters, 18 variants in all. None passed. In the research model the best rule took its direction from the first candle and traded only on days with high opening volume; this is the rule that was carried forward. On QuantConnect's data it reaches a Sharpe ratio of 0.84 and beats 99.5% of the random-direction runs, but its deflated Sharpe ratio is 0.74. On these data a different filter ranks first, the first candle with a minimum gap (1.03). In every combination the first candle was at least as good a guide to direction as the gap, and opening volume helped only with the first-candle and first-break directions. **Moderate confidence.**

Selected candidate. Direction: sign of the first 5-minute candle of QQQ (close 9:34 against open 9:30); no trade if it is flat. Day filter: QQQ volume of the opening minutes at least equal to its mean over the 14 previous sessions. Since the data correction in section 7 the volume is taken from the bars 9:31 to 9:34. Entry: stop order at the high (long) or low (short) of the 5-minute range, valid until 11:30. Stop: opposite edge of the range, not trailed. Exit: stop or 15:55. Size: 1% of the account at risk, at most 4 times leverage. In MNQ the signal comes from QQQ, and the orders go to the front contract.

7 Stops

Stop rules. I registered eight stop rules on the selected filter: two initial stops (the other side of the range, or 15% of the average true range) and four ways to trail them (none, break-even after 1 R, a trail 1 R behind the 5-minute close, and a trail along a smoothed price line on 5-minute bars). None raised the Sharpe ratio by a relevant amount. Moving the stop to break-even after 1 R gave 0.85 against 0.84 and lowered the largest loss from −19% to −13%. Trailing by 1 R cut the big winners and fell to 0.15–0.48, trailing along the smoothed line to 0.60–0.73. The rule lives from the few days that run to the close. **Moderate confidence.**

8 Replication: where research model and market differ

The candidate was replicated on QuantConnect with MNQ futures, although it had not passed the development hurdle. The replication is the most instructive part of this study. The rule stayed the same throughout; what changed were data, instrument and execution. Four versions were needed before research model and replication agreed (Table 1), and each difference had a cause that no single system would have shown (Table 2).

VERSION	RETURN P.A.	SHARPE	MAX. LOSS	TRADES	SLIPPAGE
1: signal and volume from the future, minute data	−4.6%	−0.31	−41.9%	867	0.12 R
2: signal from QQQ, minute data	6.4%	0.17	−24.7%	607	0.16 R
3: signal from QQQ, second data	13.1%	0.47	−16.9%	618	0.04 R
4: as 3, volume measure corrected	**16.8%**	**0.63**	**−16.6%**	**649**	**0.04 R**
Research model in R, QQQ (to 20 Nov 2026)	16.2%	0.90	−23.9%	645	

Table 1 · Four versions of the replication on QuantConnect from January 2020; versions 1 to 3 to August 2026, version 4 to 20 November 2026. Sharpe ratio on QuantConnect after the risk-free rate, in R without it. Slippage: mean cost per trade against the trigger price.

Fill model. With minute data QuantConnect fills a stop order at the close of the minute in which it triggers. On a breakout this close often lies well beyond the trigger. Entry slippage averaged 7.6 points and stop slippage 4.4 points, together 0.16 R per trade. With second data the same orders filled 1.3 and 1.6 points away, 0.04 R per trade. The research model fills at the trigger price and is somewhat optimistic. Second data are closer to a real fill than either.

Signal source. The first version took direction and volume from the futures contract. Futures volume selects different days than QQQ volume does: 192 days appeared only in R (+0.44 R on average) and 408 only on QuantConnect (−0.08 R). The day filter is tied to QQQ volume and does not transfer to futures volume. This weakens the case for the volume filter as a general effect.

Vendor data. After versions 2 and 3, R and QuantConnect still disagreed on many days from 2025 on. In 2025 there were 23 trades only in R and 28 only on QuantConnect, against one to three a year before. The cause was a change in the FMP data. The volume of the opening one-minute bar jumps by a factor of four in March 2025, falls to a tenth in May 2026 and jumps back in September 2026 (Figure 2). The likely reason is that the opening auction is sometimes included and sometimes not. The following bars show no such break.

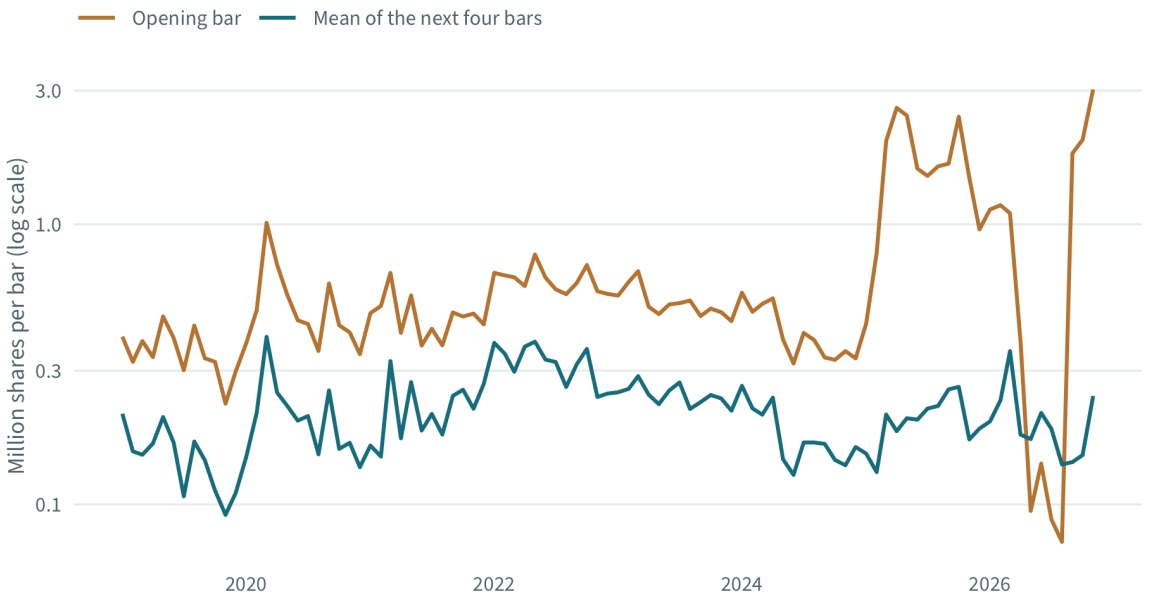


Figure 2 · QQQ volume of the opening one-minute bar and mean volume of the next four bars, monthly median, million shares, log scale. · Source: FMP QQQ minute bars, Jan 2019–Nov 2026; author's calculations · As of 21 Nov 2026

A filter that compares opening volume with its recent average reacts to such a jump as if the market had changed. After an upward jump it passes almost every day for some weeks, after a downward jump almost none. In 2025 and 2026 the research model therefore selected days by the state of the data. I replaced the volume measure by one that leaves out the opening bar, in R and on QuantConnect alike, and counted the change as a variant. The Sharpe ratio in R fell from 0.95 to 0.90, and most of the loss fell on 2026 (1.54 to 0.66). The error had raised the overall result only a little, but it had made the latest year look much stronger than it was.

Agreement. After these corrections the two systems agree. Over the full period they share 620 trading days, with 25 days only in R and 29 only on QuantConnect. On the common days the mean result is 0.212 R in both systems, with a correlation of 0.96 and the same direction on 98% of the days (Figure 3). The remaining differences come from volume near the threshold and lie mostly in 2025. Whether FMP or QuantConnect deviates there is open. **Moderate confidence.**

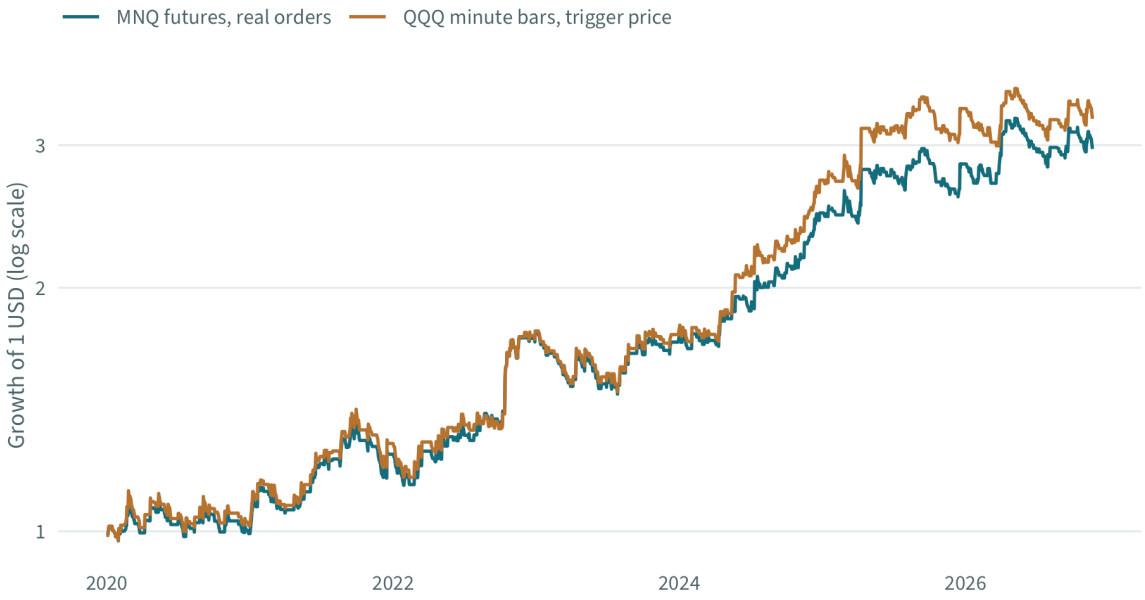


Figure 3 · Growth of 1 USD, January 2020 to November 2026: the candidate at the trigger price on QQQ minute bars and with real orders in MNQ. The MNQ curve is rebuilt from the order fills and matches the reported final equity to the dollar. · Source: QuantConnect QQQ minute bars and MNQ second data; author's calculations · As of 21 Nov 2026 · Backtest, net of costs

One flaw of the replication remains. On shortened sessions the market closes at 13:00, before the exit time of the rule. On 24 December 2024 QuantConnect therefore held a position overnight and closed it on the next trading day. One trade is affected. A live version must tie its exit to the exchange calendar.

DIFFERENCE	CAUSE	EFFECT	HANDLING
Fill of stop orders	Minute data fill at the close of the trigger minute	−0.12 to −0.16 R per trade	Second data: −0.04 R
Signal source	Futures volume differs from ETF volume	Different days, Sharpe −0.31	Signal from QQQ, orders in MNQ
Vendor data	Opening bar of FMP with and without auction	Days selected by data state in 2025–2026	Volume without the opening bar, counted as a variant
Shortened sessions	Fixed exit time after the early close	One position held overnight	Exit from the exchange calendar (open)
Data end	Futures data of the first account level end in August 2026	Three months missing in versions 1 to 3	Full data for version 4
Sharpe ratio	QuantConnect deducts the risk-free rate	Lower values than R	Both reported, labelled
Download	Empty request returns the last days	Duplicated days	Rows outside the window dropped
Minute bars of two vendors	Different highs, lows and volumes	51 variants 0.14 higher in Sharpe on QuantConnect data	Both reported; the code gives identical results on identical data

Table 2 · What separated research model and replication, and how each difference was handled.

Three lessons follow. First, a result is only as good as the fill it assumes; a breakout rule is most sensitive to that, because it enters where price moves fastest. Second, a filter on volume measures the data feed as much as the market, and a change in the feed looks like a change in the market. Third, the differences became visible only because both systems were compared day by day and trade by trade, not by their summary statistics. **Moderate confidence.**

9 The candidate

	QQQ, TRIGGER PRICE	MNQ, REAL ORDERS
Period	Jan 2020–Nov 2026	Jan 2020–Nov 2026
Trades	648	649
Return p.a.	18.6%	16.8%
Sharpe ratio	1.00	0.63 (after risk-free rate)
Largest loss	−17.6%	−16.6%
Hit rate	34.0%	33.6%
Mean result per trade	0.241 R	0.210 R
Median result per trade	−1.0 R	−1.0 R
Better than random direction	99.8% of 1,000 runs	99.7% of 1,000 runs
Deflated Sharpe ratio (51 variants)	0.83	

Table 3 · The candidate on QuantConnect: at the trigger price on QQQ minute bars, and with real orders in MNQ. Random direction: same days, same entry and stop rules, direction drawn at random, computed on the backtest cluster.

The candidate earns 0.21 R per trade with real orders and 0.24 R at the trigger price (Table 3). The median trade is a full loss. Long and short trades contribute about equally: 331 long trades with a hit rate of 34% and 318 short trades with 33%. Short winners are slightly larger (2.70 R against 2.56 R).

Exposure. Position size follows the distance to the stop, so each trade risks the same share of the account. A narrow opening range gives a near stop and a large position; a wide range gives a small one. With real orders the position was 2.4 times the account in the median, 2.5 for long and 2.4 for short trades, and 25% of the long and 21% of the short trades were above 3.5 times (Figure 4). The account is flat on the 63% of days without a

trade and every night. Averaged over all days, exposure is about once the account, but concentrated in a few hours of a third of the days.

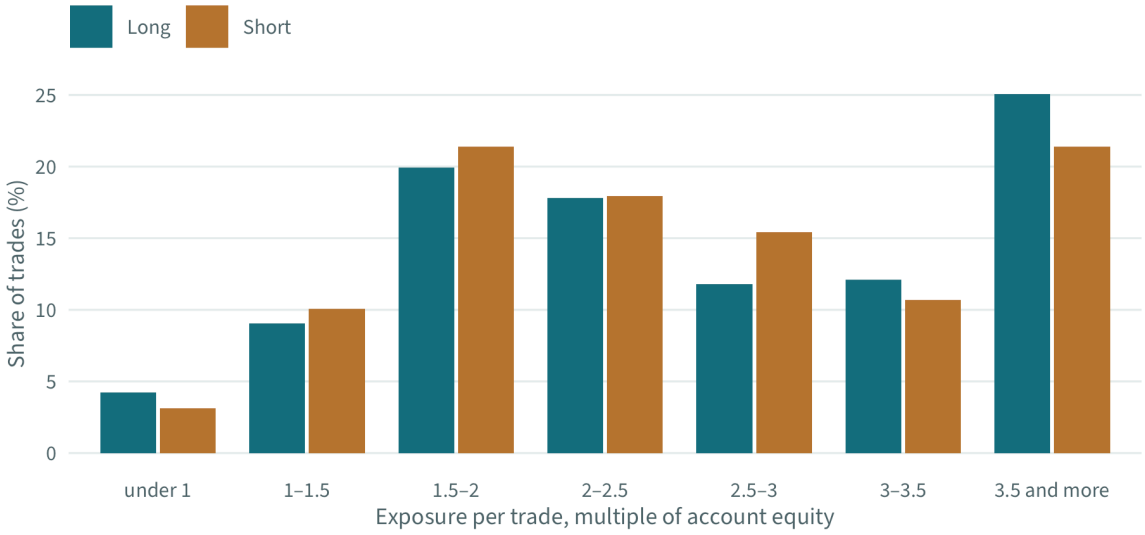


Figure 4 · Exposure per trade as a multiple of account equity at the prior close, long and short trades, MNQ real orders 2020–2026. · Source: QuantConnect MNQ second data; author's calculations · As of 21 Nov 2026 · Backtest, net of costs

Against the hurdle. The candidate beats random direction clearly. Its Sharpe ratio at the trigger price (0.997) falls just short of the required 1.0, and its deflated Sharpe ratio of 0.83 is below 0.95; it rises to 0.86 when skewness and kurtosis of the daily returns are taken into account. With 51 variants tried, a result of this size can still come from selection.

Test sample 2015–2019. The registered criterion for the test sample required three things: a Sharpe ratio of at least 0.5, a mean result per trade above zero with a bootstrap lower bound (5% quantile) above zero, and a result better than 90% of random-direction runs. The candidate was computed on QuantConnect's QQQ minute bars, as in development. Over 1,258 days it made 426 trades. The Sharpe ratio was −0.18, the return −3.6% a year and the largest loss −33%. The mean result per trade was −0.005 R with a lower bound of −0.14 R, and the rule beat only 30% of the random-direction runs. It failed all three criteria. By year the Sharpe ratio was −1.17 in 2015, −1.13 in 2016, 0.22 in 2017, 1.05 in 2018 and −0.57 in 2019. With real orders in E-mini Nasdaq-100 futures (NQ; MNQ started trading only in May 2019) and an account of 1 million USD, the same rule entered the same 426 trades and lost 3.9% a year, with a largest loss of −31.7% and a Sharpe ratio of −0.33 after the risk-free rate. Run as one account from January 2015 to 20 November 2026 in NQ, the rule made 1,073 trades and earned 8.5% a year, with a largest loss of −31.7% and a Sharpe ratio of 0.31 after the risk-free rate. On the common days it agrees with the two separate runs: all 426 trades of 2015–2019 match, and 647 of the 649 MNQ trades of 2020–2026 fall on the same days with the same direction. **Moderate confidence.**

Years. The result is uneven (Figure 5). Sharpe ratios at the trigger price by year: 2020 0.18, 2021 1.30, 2022 1.60, 2023 0.01, 2024 2.37, 2025 1.09, 2026 −0.10. With real orders 2024 is the strongest year and 2023 the only year with a loss.

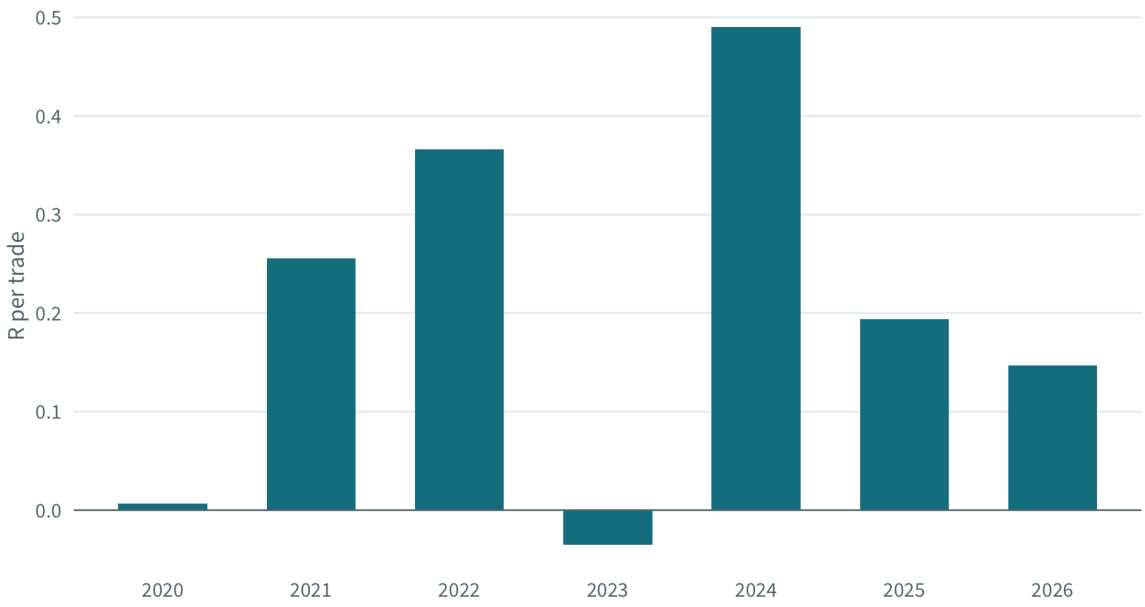


Figure 5 · Mean result per trade by year, MNQ real orders, in multiples of the initial risk; 2026 to 20 November. · Source: QuantConnect MNQ second data; author's calculations · As of 21 Nov 2026 · Backtest, net of costs

Concentration. The ten best trades carry 66% of the total profit, and the best 2% of trades carry 78% (Figure 6). Without them a fifth of the profit would remain. This is the price of a rule that cuts every loser at −1 R and lets every winner run: it needs the rare trend day and has to sit through long stretches without one.

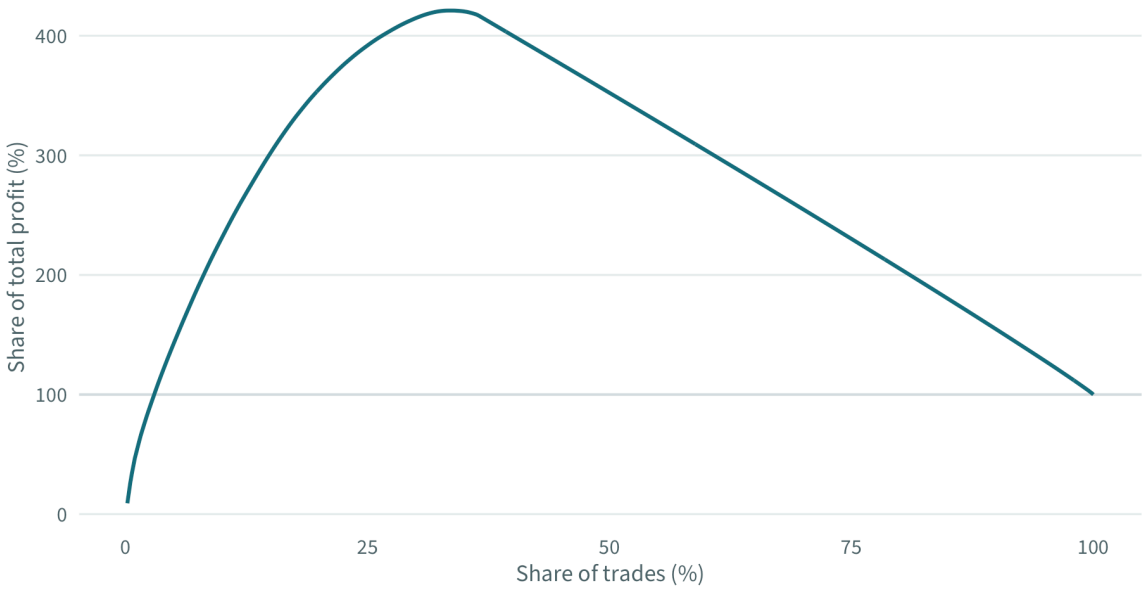


Figure 6 · Cumulative share of total profit, MNQ real orders, trades sorted from best to worst. The curve rises above 100% because the losing trades follow; it ends at 100%. · Source: QuantConnect MNQ second data; author's calculations · As of 21 Nov 2026 · Backtest, net of costs

10 Limits

- **Number of trials.** 51 variants were computed. The registration limits selection bias, but the deflated Sharpe ratio shows that the candidate's result is not safe from it.
- **Selection and data.** The candidate was selected in the research model. On QuantConnect's data the same 51 variants score 0.14 higher in Sharpe on average, and a different filter ranks first. Which variant looks best depends on the data vendor.
- **Few decisive days.** A rule of this type lives from rare trend days. Seven years contain only a handful of them, and two or three years decide the result.
- **Day filter.** The volume filter selects different days with futures volume and reacted strongly to a change in the vendor's data. A filter this sensitive to its input is fragile.
- **Instrument.** Signals come from QQQ, orders go to MNQ. The overnight gap of the future differs from that of the ETF, because the future trades through the night.
- **Replication.** Shortened sessions are not handled correctly. The random-direction test for real orders was computed at the trigger price of the same days (mean result 0.25 R, bootstrap lower bound 0.12 R).

11 Conclusion

The opening gap of the Nasdaq-100 does not tell where the day is going. A breakout after the open in gap direction is hardly better than one in a random direction, and it loses money on most days. What decides the result is the day: the breakout earns on wide trend days and pays small stops on quiet ones. A rule that chooses days by the first candle and the opening volume reaches a Sharpe ratio of 1.00 on QQQ minute bars and 16.8% a year with real orders. That is close to what Zarattini and Aziz (2023) report for a similar rule on QQQ before slippage. It fails the registered hurdle, and in the untouched years 2015–2019 it lost money and did no better than chance. Their sample of 2016–2023 includes four of these years. The two rules differ in the entry and in the volume filter, so this is no direct test of their result, but it shows how much a rule of this family depends on the period. The result of 2020–2026 is best read as a property of that period and of the selection among 51 variants.

One result reaches beyond this rule. A change in vendor data can create or inflate an effect without any visible error in the code. Here only the comparison with a second data source and a second execution model exposed it.

The test sample is now used. A new rule on the same question needs data that no step of this study has seen, for example the years after November 2026.

References

Bailey, D. H., Borwein, J. M., López de Prado, M. and Zhu, Q. J. (2015): The Probability of Backtest Overfitting. SSRN Working Paper 2326253.

Bailey, D. H. and López de Prado, M. (2014): The Deflated Sharpe Ratio: Correcting for Selection Bias, Backtest Overfitting and Non-Normality. *Journal of Portfolio Management* 40 (5), 94–107. SSRN Working Paper 2460551.

Baltussen, G., Da, Z., Lammers, S. and Martens, M. (2020): Hedging Demand and Market Intraday Momentum. SSRN Working Paper 3760365.

Gao, L., Han, Y., Li, S. Z. and Zhou, G. (2017): Market Intraday Momentum. SSRN Working Paper 2440866.

López de Prado, M. and Bailey, D. H. (2018): The False Strategy Theorem: A Financial Application of Experimental Mathematics. *American Mathematical Monthly* 125 (7). SSRN Working Paper 3221798.

Zarattini, C. and Aziz, A. (2023): Can Day Trading Really Be Profitable? Evidence of Sustainable Long-term Profits from Opening Range Breakout (ORB) Day Trading Strategy vs. Benchmark in the US Stock Market. SSRN Working Paper 4416622.

Zarattini, C., Barbon, A. and Aziz, A. (2024a): A Profitable Day Trading Strategy for the U.S. Equity Market. SSRN Working Paper 4729284.

Zarattini, C., Aziz, A. and Barbon, A. (2024b): Beat the Market: An Effective Intraday Momentum Strategy for S&P500 ETF (SPY). SSRN Working Paper 4824172.

Appendix A: Registered steps (results on QuantConnect)

STEP	VARIANTS	RESULT
Gap rules: range length, entry, stop, break-even	24	Failed (best Sharpe 0.67)
Direction and day filters	18	Failed (selected 0.84, deflated 0.74; best 1.03)
Stop rules on the selected filter	8	Failed (best 0.85 against 0.84 for the plain stop)
Volume measure without the opening bar (data correction)	1	Failed (1.00, deflated 0.83)
Test sample 2015–2019, candidate only	0	Failed (Sharpe −0.18, better than 30% of random runs)
Real orders on QuantConnect, four versions	0	Agrees per trade with the research model (0.212 R in both)

Appendix B: Rule of the candidate and code

Instrument for signals: QQQ, regular session, New York time. Instrument for orders: MNQ front contract (open-interest mapping). Opening range: 9:30 to 9:34. Direction: close of 9:34 against open of 9:30; no trade if equal. Volume filter: sum of QQQ volume 9:31 to 9:34 divided by its mean over the 14 previous sessions; trade only if at least 1.0. Entry: buy stop at the MNQ range high (long) or sell stop at the range low (short), placed at 9:35 and cancelled at 11:30 if not filled. Stop: opposite range edge, placed on fill and never moved. Exit: market order at 15:55. Size: whole contracts, risk 1% of equity to the stop, notional at most 4 times equity, MNQ point value 2 USD. Costs in R: 0.01% of price per round trip. QuantConnect: default futures fees and fill model, MNQ in second resolution, QQQ in minute resolution.

Code: all QuantConnect algorithms and R scripts are published at github.com/LudwigJ-Marx/quant-research (folder wp01-opening-gap-breakouts) with a list of which file produces which number. The price data of FMP and QuantConnect are licensed and not included. The full run with real orders in NQ from 2015 to November 2026 is public on QuantConnect, with code, orders and statistics: https://www.quantconnect.cloud/backtest/737aabdee9588fffd28b21c70ea61438/.