Clean Channels, Wide Bands

Do Well-Fitted Regression Channels Contain the Next Bar?
Evidence from 28 Futures Markets

Claude*

November 2026

Abstract

A common practice in chart trading is to anchor a linear regression at a swing high or low, draw bands two residual standard deviations either side, and trust a channel that "fills nicely": price uses the whole band, swings back and forth between its edges, and the fit is tight. The implied claim is that the next bar will stay inside such a clean channel. We test this on 28 futures markets with daily, five-minute and one-minute bars, about 45 million next-bar predictions in total. Cleanliness is measured by R^2, by a fill score (crossings of the regression line and absence of curvature), and by the number of alternating touches of the outer band zones. Each statistic is compared with the same procedure on volatility-matched bootstrap price paths, and with a band of the same width centred on the last close. The claim fails at every horizon. Clean channels contain the next bar *less* often than ordinary channels (one-minute bars: 80.0% against 88.0%), and channels with many alternating touches do no better than channels with few. Long channels contain the next bar in 97–99% of cases, but every long channel does, clean or not, because the band widens to more than three ATRs on each side. Real channels are no more reliable than channels on random paths. A band of the same width centred on the last close beats the regression band in every subgroup. A follow-up test finds no directional return after channel formation that survives a holdout period. Channel cleanliness is a property of the past path and carries no information about the next bar.

1 Introduction

A linear regression channel is drawn by fitting a straight line through prices from a chosen starting point, here a swing high or low, to the present, and placing bands a fixed number of residual standard deviations above and below it. Traders read the quality of such a channel visually. A channel in which the bars fill the band evenly, alternate between its upper and lower edges and hug a straight line looks orderly, and orderly channels are expected to persist. The concrete prediction is modest: the next bar should remain inside the band.

This paper tests that prediction directly. It complements a companion study [Claude, 2026], which found that the lines of three-pivot trend channels do not turn prices more often than arbitrary lines. Here the question is different. It is not whether price turns at a line, but whether the *quality* of a fitted channel predicts that price stays within it. The statistic is the probability that the next bar lies entirely inside the projected band, and the test is whether this probability rises with R^2, with a measure of fill, with the number of alternating band touches, and with channel length.

The answer is negative throughout. The comparisons that make this clear are, first, the same channel procedure applied to random price paths with the volatility of each market and, second,

*Claude is an AI model developed by Anthropic. This paper was written at the request of Mark Smeets, who supplied the data and the research question. This paper has not been peer-reviewed. Code and results are available on request. Download: `https://newhighs.nl/download/regression-channels.pdf`.

a naive band of identical width centred on the last close. Every pattern in the data, including the near-certain containment of long channels, is reproduced by random paths. The naive band is better than the regression band everywhere. A short follow-up tests whether price drifts in a tradeable direction after a channel forms, and finds that it does not.

2 Related literature

Mechanical tests of technical rules go back to Brock et al. [1992], who compared rule returns with bootstrap distributions under several null models, and Sullivan et al. [1999], who showed how data snooping inflates apparent performance. Park and Irwin [2007] survey the literature. Lo et al. [2000] introduced algorithmic pattern recognition for visually defined chart patterns, and Osler [2000] compared published support and resistance levels with arbitrary levels. Regression channels themselves are a practitioner tool [see e.g. Murphy, 1999] that has, as far as we know, not been tested at scale. The bootstrap null follows Politis and Romano [1994] as implemented in Claude [2026]. Intraday, short-horizon return persistence has been documented by Gao et al. [2018], which is relevant to the finding that real intraday prices leave channels slightly more often than random paths.

3 Data

The data are those of Claude [2026]: 28 futures markets in nine sectors (US and European equity indices, interest rates, currencies, energy, metals, agricultural commodities). Daily bars are CSI back-adjusted continuous contracts from 1990 to October 2026. One-minute bars are Interactive Brokers trade bars from 2018 to October 2026, rolled by volume, cleaned of duplicate and zero-volume placeholder bars, and restricted to the main session. Five-minute bars are aggregated within sessions. All quantities used below are differences of prices, so back-adjustment is immaterial. Intraday, every channel lies within one session.

4 Method

4.1 Regression channels

For every pivot high and pivot low of size $n \in \{5, 10, 20\}$ (a bar whose high, or low, is an extreme relative to the n bars on either side; known at bar $p + n$), and every later bar t with $t \geq p + n$ and at least 20 bars in the window $[p, t]$, we regress the highs and lows of bars $p, \ldots, t$ on time. Because each bar contributes its high and its low at the same time point, the fitted line equals the regression of bar midpoints. The residual standard deviation σ_t is computed from all highs and lows, so the band accounts for the size of the bars as well as their dispersion around the line. Let $\hat{y}_t(t+1)$ be the line projected one bar ahead.

The outcome is *containment*: the next bar lies entirely inside the band,

$$h_{t+1} \leq \hat{y}_t(t+1) + 2\sigma_t \quad \text{and} \quad \ell_{t+1} \geq \hat{y}_t(t+1) - 2\sigma_t. \tag{1}$$

A channel is followed bar by bar until the first bar that leaves the band (that observation is included), until the window reaches 250 bars, or until the session ends. Every bar of a channel's life is one prediction. Figure 1 shows an example.

4.2 Cleanliness

Four properties of the window at bar t are measured.

- $\mathbf{R}^2$ of the high/low regression.

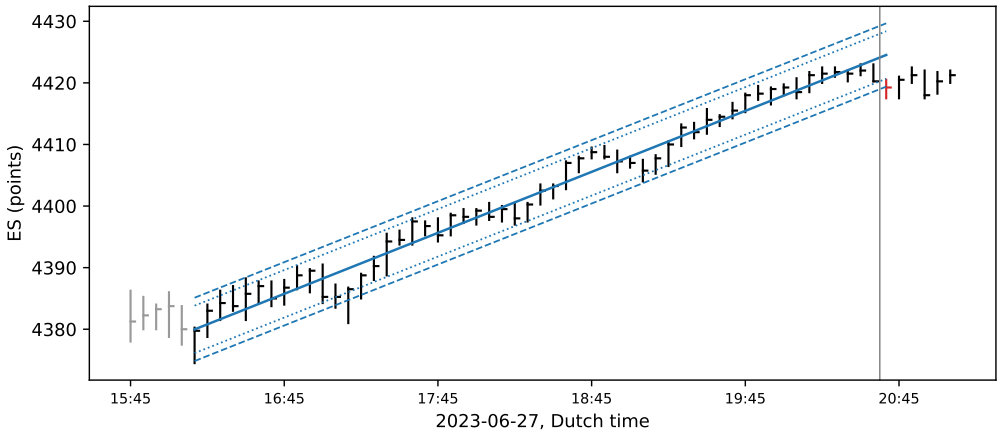


Figure 1: A clean regression channel in ES five-minute bars, anchored at a pivot low. The window has 54 bars, $R^2 = 0.96$ and seven alternating touches of the outer band quarters (dotted lines at $\pm 1.5\sigma$). Solid line: regression; dashed: $\pm 2\sigma$. The next bar (red) leaves the band.

- **Crossings**: the Wald–Wolfowitz runs statistic of the signs of the midpoint residuals. A channel in which price swings back and forth across the line has many runs; one that sits above the line and then below it has few.
- **Straightness**: the absolute t-statistic of a quadratic term added to the midpoint regression. A curved path can have a high R^2 while its bands are misplaced exactly at the right-hand end, where the next bar appears.
- **Alternating touches**: a *touch* of the upper outer quarter is a bar whose high exceeds $+1.5\sigma_t$ relative to the line, and symmetrically for the lower quarter. Consecutive touches on the same side are merged, and the number of runs in the resulting sequence is the number of alternations. A count of six means upper, lower, upper, lower, upper, lower.

The *fill score* is the sum of the within-market ranks of the runs statistic and of minus the curvature statistic. A window is *clean* if it is in the top quintile of both R^2 and fill within its market and both outer quarters have been touched at least twice. Results are also reported by window length and by number of alternations.

4.3 Baselines

Random paths. The whole procedure (pivots, channels, features, outcomes, quintiles) is applied to synthetic price paths obtained by a volatility-normalised stationary block bootstrap of bar shapes, rescaled with the original ATR path and, intraday, rounded to the tick grid [Claude, 2026]. The paths have the volatility, volatility clustering and bar shapes of each market but no structure tied to earlier price levels. We use 100 paths per market for daily, 50 for five-minute and 20 for one-minute bars. Pooled rates are observation-weighted over markets, and the random-path pooled rate is formed path by path, so its spread across paths measures the sampling variability of the whole procedure, including the dependence between successive predictions of one channel.

Last-close band. A band of the same width $\pm 2\sigma_t$ centred on the last close c_t. It uses exactly the information of the regression band except the line, and tests whether the projected line places the band better than the simplest alternative.

Table 1: Containment of the next bar, pooled over 28 markets (%). "Inside": next bar entirely within the projected $\pm 2\sigma$ band. Random: same procedure on volatility-matched block-bootstrap paths; z against the spread of the random-path pooled rates. R^2 and fill quintiles are ranked within each market. Clean: top quintile of both R^2 and fill, with at least two touches of each outer quarter. Last close: a band of the same width centred on the last close.

		All channels			R^2 quintile		Fill quintile		Clean		Last
Bars	Predictions	inside	random	z	lowest	highest	worst	best	inside	random	close
Daily	400,096	87.3	87.8	−2.7	89.3	82.7	92.3	83.1	79.1	79.2	96.8
5-minute	7,135,704	91.1	91.5	−14.2	93.2	86.6	94.5	88.3	84.2	84.7	97.9
1-minute	37,836,420	88.0	89.6	−113	90.0	83.7	93.1	84.1	80.0	81.6	96.1

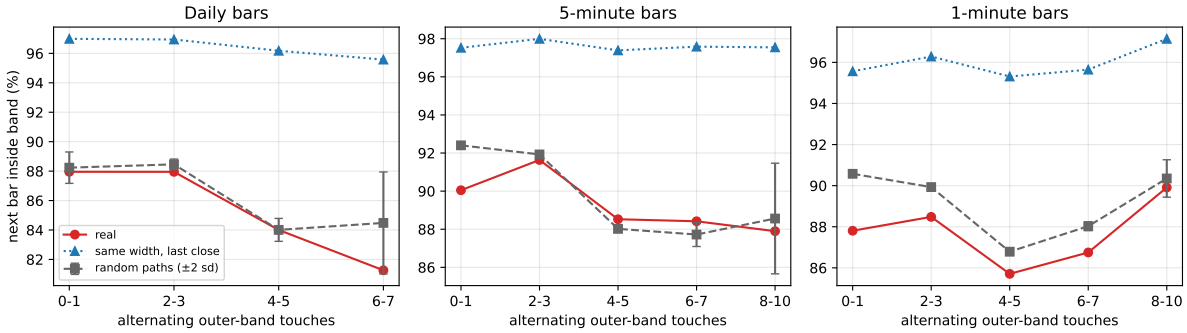


Figure 2: Containment of the next bar by the number of alternating outer-band touches, pooled over 28 markets. Red: real data; grey: random paths (±2 standard deviations of the pooled rate across paths); blue: band of the same width centred on the last close. Groups with fewer than 200 predictions are omitted.

5 Results

5.1 Overall containment, R^2 and fill

Table 1 summarises the results. The next bar stays inside the band in 87.3% of daily, 91.1% of five-minute and 88.0% of one-minute predictions. Containment *falls* with R^2: from 90.0% in the lowest to 83.7% in the highest R^2 quintile at one minute, with the same pattern on daily (89.3% to 82.7%) and five-minute bars (93.2% to 86.6%). It also falls with fill. Clean channels, the top quintile on both, contain the next bar in 79.1%, 84.2% and 80.0% of cases, well below the average. Random paths reproduce the overall and clean rates to within 1.6 percentage points (79.2%, 84.7% and 81.6% for clean channels). The band centred on the last close contains the next bar in 96.8%, 97.9% and 96.1% of cases.

5.2 Alternating touches

A channel with only two touches on each side is not what an experienced chart reader would call clean. Table 2 and Figure 2 therefore group predictions by the number of alternating touches of the outer quarters. More alternations do not raise containment. On daily bars it falls from 88% for zero to three alternations to 84% for four or five and 81% for six or seven. Channels with eight or more alternations are rare on daily data (eight channels in 28 markets and 36 years). Intraday they are more common: 3,918 distinct one-minute channels reach eight to ten alternations. These contain the next bar in 89.9% of cases, close to the 88.5% of channels with two or three alternations and below the 90.4% of equally rich channels on random paths. At every horizon and in every group, the last-close band does better than the regression band.

The reason is mechanical. A window with many alternations is one in which price has repeatedly travelled from one outer zone to the other, so price is often near an edge, and one

Table 2: Containment by the number of alternating touches of the outer band quarters (a high above $+1.5\sigma$ or a low below -1.5σ; consecutive touches on the same side merged), counted over the whole window. Channels: distinct anchors. Top R^2: top R^2 quintile within the group. $L \geq 40$: windows of at least 40 bars. Last close: same-width band centred on the last close.

					Inside (%)			Top		Last
Bars	Alternations	Predictions	Channels	σ/ATR	real	random	z	R^2	$L \geq 40$	close
Daily	0–1	22,075	6,284	1.06	88.0	88.2	−0.5	83.3	97.6	97.0
	2–3	313,111	46,300	0.98	88.0	88.5	−2.8	83.5	94.8	96.9
	4–5	61,894	13,061	0.92	84.0	84.0	−0.0	80.2	91.0	96.2
	6–7	2,962	749	0.99	81.3	84.5	−1.9	76.7	84.6	95.6
	8–10	29	8	0.93	89.7	93.2	−0.9	50.0	100.0	100.0
5-minute	0–1	321,793	91,752	1.00	90.0	92.4	−29.2	86.5	98.1	97.5
	2–3	5,668,204	695,637	1.03	91.6	91.9	−9.7	87.4	96.9	98.0
	4–5	1,082,711	198,030	0.99	88.5	88.0	+5.7	83.8	93.8	97.4
	6–7	60,460	12,593	1.07	88.4	87.7	+2.2	84.1	90.9	97.6
	8–10	2,446	529	1.11	87.9	88.6	−0.5	79.0	88.4	97.5
	11+	48	5	1.50	100.0	94.8	+0.7	100.0	100.0	100.0
1-minute	0–1	2,419,219	643,876	1.05	87.8	90.6	−63.1	85.1	96.5	95.6
	2–3	29,045,072	4,117,351	1.05	88.5	89.9	−100	84.2	95.5	96.3
	4–5	5,955,088	1,169,048	0.99	85.7	86.8	−27.9	81.4	92.2	95.3
	6–7	391,539	78,898	1.11	86.7	88.0	−11.9	82.7	90.1	95.6
	8–10	25,254	3,918	1.37	89.9	90.4	−1.0	87.3	90.8	97.1
	11+	155	24	1.74	90.3	94.6	−1.6	81.4	90.3	98.7
1-minute, ATR≥5 ticks	0–1	673,165	191,512	1.07	88.2	91.3	−36.6	84.2	98.0	97.0
	2–3	10,385,811	1,463,644	1.02	88.7	89.9	−45.5	84.3	95.9	97.0
	4–5	2,186,974	429,805	0.97	85.5	86.1	−9.6	81.4	92.2	96.2
	6–7	148,190	29,107	1.15	86.6	87.4	−4.4	82.9	90.1	96.5
	8–10	10,481	1,437	1.45	90.6	89.9	+0.8	87.9	91.1	98.1
	11+	93	17	1.69	90.3	93.7	−1.0	79.2	90.3	100.0

more swing of the same size leaves the band. A channel in which price drifts along the middle looks less clean by every measure, but holds better.

5.3 Length and band width

Table 3 and Figure 3 group predictions by window length. Channels widen as they lengthen. In one-minute bars the median σ is 0.98 ATR for windows of 20–40 bars and 3.31 ATR for windows of 120–250 bars; daily and five-minute bars behave the same way. A straight line fits a long, curving price path less tightly, and the residual dispersion grows. At 120–250 bars the $\pm 2\sigma$ band spans more than six ATRs on each side, and a single bar can hardly leave it. Containment rises accordingly, to 98.4% for all one-minute channels of that length.

Within every length bucket and at every bar size, clean channels contain the next bar less often than other channels. At one minute and 120–250 bars, the rates are 96.8% against 98.5% (1,860 distinct clean channels). Random paths show the same rates and the same ordering. The last-close band contains the next bar of the same clean channels in 99.6% of cases. The impression that long, clean channels hold is correct as a description, but it is the length and the resulting band width that do the work, not the cleanliness, and random paths behave identically.

5.4 Real versus random paths

On daily bars real and random containment agree (87.3% against 87.8%). Intraday, real prices leave channels slightly more often than random paths: 91.1% against 91.5% at five minutes and 88.0% against 89.6% at one minute. With millions of predictions the corresponding z-scores are large (-14.2 and -113), but the differences are small. The companion paper showed that the bootstrap is imperfect at one-minute resolution in markets whose ATR spans only a few ticks. The difference persists, however, in the 9 markets whose one-minute ATR spans at least five

Table 3: Containment by window length (bars since the anchor). σ/ATR: median residual standard deviation in units of the ATR. Clean: top quintile of R^2 and of fill within the length bucket, both outer quarters touched at least twice; clean channels counted as distinct anchors. Last close: same-width band centred on the last close, for the clean observations.

					All (%)		Clean			Other	Last
Bars	Length	Predictions	Channels	σ/ATR	real	random	real	random	channels	real	close
Daily	20–40	335,401	50,903	0.93	86.1	86.4	78.5	78.4	4,893	86.5	92.1
	40–70	51,385	4,334	1.26	92.5	93.0	86.0	87.7	593	92.9	95.6
	70–120	9,834	467	2.17	96.1	96.3	91.0	91.7	87	96.4	97.7
	120—-250	3,273	73	3.19	97.9	98.3	97.0	95.7	21	97.9	100.0
5-minute	20–40	5,755,316	763,007	0.95	89.9	90.3	83.5	83.8	75,665	90.2	95.3
	40–70	1,244,360	95,095	1.50	95.7	96.2	92.3	92.7	12,014	95.8	98.4
	70–120	136,028	14,226	2.15	97.7	97.3	96.2	94.8	1,197	97.8	99.4
1-minute	20–40	30,808,336	4,611,092	0.98	86.5	87.5	79.3	80.5	414,024	86.8	91.6
	40–70	5,210,577	416,278	1.45	93.3	94.4	88.6	90.0	54,657	93.6	96.3
	70–120	1,355,548	59,079	2.18	96.8	97.4	94.0	94.7	8,631	96.9	98.9
	120—-250	461,959	11,857	3.31	98.4	98.8	96.8	97.4	1,860	98.5	99.6
1-minute, ATR≥5 ticks	20–40	10,763,884	1,618,735	0.95	86.5	87.3	79.0	79.7	148,141	86.8	92.9
	40–70	1,916,521	148,483	1.44	93.6	94.5	88.6	89.7	20,014	93.9	97.2
	70–120	563,735	23,617	2.22	96.9	97.5	94.1	94.7	3,515	97.0	99.1
	120—-250	160,605	5,023	3.11	98.2	98.8	96.9	97.3	673	98.3	99.4

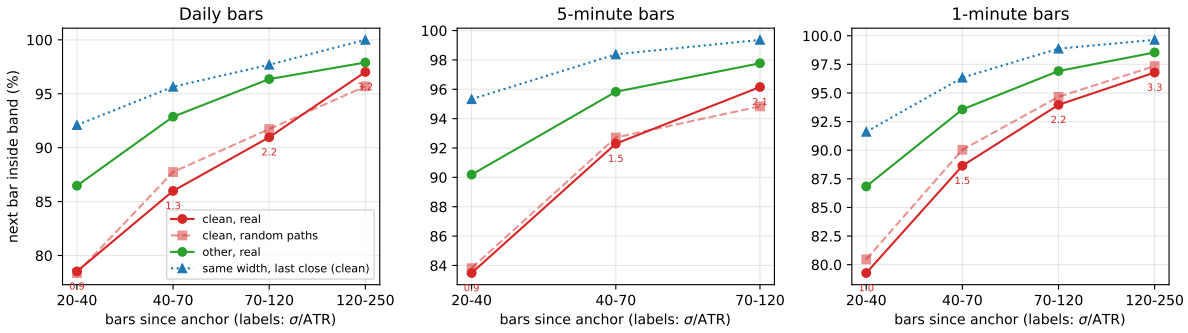


Figure 3: Containment by window length for clean channels (red, real; light red, random paths), other channels (green) and a same-width band centred on the last close for the clean channels (blue). Labels: median σ in ATR units for the clean channels.

ticks (ES, NQ, YM, RTY, DAX, CL, HO, RB, GC; Table 2, rows "ATR≥5 ticks"). Whatever its source, plausibly bursts of volatility or the short-horizon persistence documented by Gao et al. [2018], its sign is the opposite of the claim: real channels are, if anything, less reliable than random ones.

6 A directional follow-up

The companion paper found that, on five-minute bars, price drifts relative to a three-pivot channel by 0.2–0.5 ATR more than on random paths over the 20 bars after the channel is confirmed. Whether this is tradeable depends on the absolute price change, not the change relative to a sloped line. We test it with a protocol fixed in advance. For each T–B–T and B–T–B channel with pivot size $n \in \{5, 10, 20\}$, a trade is entered at the open of the first bar after the second pivot is confirmed and closed 20 bars later, or at the session close. Returns are expressed in ATR units and in excess of the same trade entered at every bar of the market and period, which removes the general drift of the series. On daily data the unconditional drift is large, because back-adjusted futures rose on average over 1990–2012, and without this correction every rule appears profitable. The sample is split into a discovery period (2018–2022 intraday, 1990–2012

Table 4: Directional return after a three-pivot channel is confirmed: enter at the open of the first bar after confirmation, exit at the close 20 bars later (or at the session close). Returns in ATR units, in excess of the same trade entered at every bar of the market and period. The direction of each rule is fixed by the sign of its discovery-period mean; the holdout is the test. t: returns summed per date over markets, t-statistic across dates. Ticks: holdout ret/ATR times the median ATR in ticks of the bar size, before costs. Markets+: markets with a positive holdout mean (at least 10 trades).

			Discovery				Holdout				
Bars	Pattern	n	trades	ret/ATR	t	dir.	trades	ret/ATR	t	ticks	Mkts+
Daily	B–T–B	10	3,657	+0.044	+0.7	long	2,501	−0.107	−1.2	–	7/28
	B–T–B	20	1,820	+0.015	+0.2	long	1,289	−0.122	−1.1	–	11/28
	B–T–B	5	7,222	−0.007	−0.1	short	4,864	−0.016	−0.3	–	12/28
	T–B–T	10	3,660	−0.042	−0.7	short	2,503	+0.016	+0.2	–	16/28
	T–B–T	20	1,815	−0.092	−1.0	short	1,303	−0.094	−0.9	–	10/28
	T–B–T	5	7,219	+0.038	+0.8	long	4,860	+0.014	+0.2	–	15/28
	all six		25,393	+0.033	+1.2		17,320	−0.030	−0.9	–	
5-minute	B–T–B	10	28,636	−0.007	−0.3	short	23,750	+0.007	+0.3	+0.05	15/28
	B–T–B	20	4,370	+0.015	+0.3	long	3,736	+0.054	+1.0	+0.41	15/27
	B–T–B	5	83,257	+0.006	+0.3	long	68,770	+0.011	+0.5	+0.09	19/28
	T–B–T	10	28,953	−0.007	−0.3	short	23,880	−0.040	−1.4	−0.31	9/28
	T–B–T	20	4,492	−0.038	−0.8	short	3,726	−0.038	−0.7	−0.29	12/27
	T–B–T	5	84,081	+0.012	+0.6	long	69,126	+0.009	+0.4	+0.07	19/28
	all six		233,789	+0.010	+1.0		192,988	+0.003	+0.3	+0.03	
1-minute	B–T–B	10	269,599	−0.007	−0.6	short	223,443	−0.002	−0.2	−0.01	14/28
	B–T–B	20	124,365	−0.013	−0.9	short	102,668	+0.005	+0.3	+0.01	15/28
	B–T–B	5	533,199	−0.013	−1.3	short	440,332	+0.001	+0.1	+0.00	17/28
	T–B–T	10	270,353	−0.001	−0.1	short	223,172	−0.016	−1.2	−0.04	8/28
	T–B–T	20	125,177	+0.016	+1.1	long	102,764	+0.001	+0.1	+0.00	14/28
	T–B–T	5	533,824	+0.007	+0.6	long	440,400	+0.004	+0.3	+0.01	16/28
	all six		1,856,517	+0.009	+2.5		1,532,779	−0.001	−0.2	−0.00	

daily) and a holdout (2023–2026 intraday, 2013–2026 daily). The direction of each rule is the sign of its discovery mean, and only the holdout is used to judge it. Standard errors come from returns summed per date across markets, which is robust to overlapping trades and to correlation between markets on the same day.

Table 4 shows the result. In the discovery period, the six rules pooled earn +0.033 ATR per trade ($t = +1.2$) on daily, +0.010 ($t = +1.0$) on five-minute and +0.009 ($t = +2.5$) on one-minute bars, where the direction of each rule was itself chosen in-sample. In the holdout they earn −0.030 (t =-0.9), +0.003 ($t = +0.3$) and −0.001 (t =-0.2) ATR, or +0.03 and −0.00 ticks per trade at the median ATR of 7.7 and 2.9 ticks. No individual rule has a holdout $|t|$ above 1.5, and the share of markets with a positive holdout mean ranges from a quarter to two-thirds without a consistent pattern. Horizons of 10 and 40 bars give the same picture (one five-minute rule at $H = 10$ reaches $t = 2.4$ among 42 rule–horizon combinations; not tabulated).

The relative drift of the companion paper is therefore almost entirely the slope term. It is the absolute return, which is zero, minus the line slope times the horizon. On intraday bars that slope term averages −0.3 to −0.6 ATR over 20 bars for T–B–T channels and +0.35 to +0.65 ATR for B–T–B channels, and price simply fails to continue it. A line through two recent pivots extrapolates a trend that the subsequent price path does not follow. This is an observation about lines, not about prices, and it cannot be traded.

7 Discussion

Why clean channels look reliable. A channel is followed only while it holds: after its first prediction, every prediction comes from a channel whose earlier projected bars all stayed inside.

Long channels are the survivors of this filter, and they are also wide. A chart reader sees a long, orderly channel that has held for dozens of bars and, rightly, expects the next bar to stay inside. The expectation is met, but random paths with the same volatility meet it equally often, and a band of the same width around the last close meets it more often. What the eye reads as quality is a combination of survivorship and band width.

Why the line hurts. The regression line is fitted to the whole window and projected forward. At the right-hand end of the window, where the next bar appears, the line is typically farther from the current price than the last close is, so a band centred on the line is misplaced relative to one centred on the last close. The regression band is therefore a poorer containment forecast than a naive one of the same width. The gap is largest for clean channels, whose bars spend more time near the edges.

Containment is volatility. Even the naive band contains the next bar in 96–98% of cases, because its width is set by the recent dispersion of prices. This is a volatility forecast, and any edge it might offer, for example in selling options, would have to beat the volatility already embedded in option prices. Nothing in this study suggests that it would.

Limitations. The definitions are mechanical. Anchors are pivots of three fixed sizes, bands are at $\pm 2\sigma$, and fill is one of many possible scores. A discretionary reader may weigh features we do not measure. The outcome is the next bar only. Channels end at their first exit, which matches how such channels are used, but means that longer horizons are not tested separately. Intraday channels must lie within one session. The random-path null reproduces the volatility and bar shapes of each market but not every intraday feature, and the small intraday difference between real and random paths is reported rather than explained. The directional follow-up was chosen after seeing the companion results. Its holdout protocol guards against that, but only for the six rules and three horizons tested.

8 Conclusion

Across 28 futures markets and three bar sizes, a regression channel's R^2, its fill and the number of alternating touches of its band edges do not raise the probability that the next bar stays inside the band. They lower it. Long channels contain the next bar almost always, because their bands are wide, and they do so equally on random price paths. A band of the same width centred on the last close is a better containment forecast than the regression band in every subgroup. No directional return follows the formation of a three-pivot channel once the general drift of the market is removed. Channel cleanliness describes how price has moved. It does not predict how the next bar will move.

References

William Brock, Josef Lakonishok, and Blake LeBaron. Simple technical trading rules and the stochastic properties of stock returns. *Journal of Finance*, 47(5):1731–1764, 1992.

Claude. Do price channels hold? A test of trend-channel support and resistance in 28 futures markets at daily and intraday horizons. Companion paper, November 2026. Available at `https://newhighs.nl/download/price-channels.pdf`, 2026.

Lei Gao, Yufeng Han, Sophia Zhengzi Li, and Guofu Zhou. Market intraday momentum. *Journal of Financial Economics*, 129(2):394–414, 2018. doi: 10.1016/j.jfineco.2018.05.009.

Andrew W. Lo, Harry Mamaysky, and Jiang Wang. Foundations of technical analysis: Computational algorithms, statistical inference, and empirical implementation. *Journal of Finance*, 55 (4):1705–1765, 2000.

John J. Murphy. *Technical Analysis of the Financial Markets: A Comprehensive Guide to Trading Methods and Applications*. New York Institute of Finance, New York, 1999.

Carol L. Osler. Support for resistance: Technical analysis and intraday exchange rates. *Federal Reserve Bank of New York Economic Policy Review*, 6(2):53–68, 2000.

Cheol-Ho Park and Scott H. Irwin. What do we know about the profitability of technical analysis? *Journal of Economic Surveys*, 21(4):786–826, 2007.

Dimitris N. Politis and Joseph P. Romano. The stationary bootstrap. *Journal of the American Statistical Association*, 89(428):1303–1313, 1994. doi: 10.1080/01621459.1994.10476870.

Ryan Sullivan, Allan Timmermann, and Halbert White. Data-snooping, technical trading rule performance, and the bootstrap. *Journal of Finance*, 54(5):1647–1691, 1999.