The Panic Below the Line

TSLA capitulation sessions and the 65-day moving average filter

AQI-ES-2026-TSLA-01 | July 24, 2026 | Instrument: TSLA (daily) | Study window: Jan 2021 - Jul 2026

Situation. TSLA fell 14.5% on July 23, 2026 following its earnings report, closing at 15% of the session range with the stock already trading below its 65-day moving average. Since 2021, sessions this violent and this weak into the close have occurred ten times.

Tension. A -14% earnings gap reads as the start of something. The data says the opposite: when the panic occurs *below* trend, it has marked late-stage liquidation, not an opening act. But the sample is five completed events, three of them siblings from one bottoming cluster, and the one analog that failed (April 2022) failed badly.

Options. Stand aside and demand a 65-day MA reclaim; buy the dip outright and wear the tail risk; or split the trade by horizon: sell the front-loaded, statistically reliable 10-session bounce, and hold small funded convexity against the back-loaded, cluster-dependent 50-session tail.

Recommendation. The split. Core short put spread (30-45 DTE, short strike at or below the event low of $291.72) at discounted quarter-Kelly of 8-10%, plus a satellite long call spread (50-70 DTE) at roughly one quarter of core risk, financed from the credit. Stage the second tranche on a 65-day MA reclaim within 10 sessions. Invalidation: a closing violation of the event low.

I. The screen

Two conditions define a capitulation session: a close in the bottom quartile of the day's range, and a single-session loss of 11% or worse. From January 2021 through July 24, 2026 this fires ten times. The second-stage filter keeps only events where the prior day's close was already below the 65-day simple moving average — panic arriving into an established downtrend rather than interrupting an uptrend. Six of the ten qualify, including the live July 23 event.

Date	Return	Range pos	Below 65d	+5d	+10d	+21d	+34d	+50d
2021-11-09	-11.99%	0.07	No	+3.05	+8.36	-1.92	+6.13	-7.78
2022-01-27	-11.55%	0.00	Yes	+7.48	+9.10	+4.98	+5.13	+23.69
2022-04-26	-12.18%	0.01	No	+3.75	-8.71	-24.83	-24.39	-14.16
2022-12-27	-11.41%	0.03	Yes	+4.16	+12.94	+63.06	+96.37	+58.97
2023-01-03	-12.24%	0.24	Yes	+9.94	+19.13	+74.16	+85.81	+70.33
2024-01-25	-12.13%	0.20	Yes	+3.41	+3.79	+9.18	-11.02	-5.28
2024-07-24	-12.33%	0.11	No	+7.44	-11.22	-2.47	+5.62	+11.42
2025-03-10	-15.43%	0.06	Yes	+7.14	+25.32	-0.13	+28.69	+54.77
2025-06-05	-14.26%	0.22	No	+12.09	+13.16	+4.60	+11.02	+17.72
2026-07-23	-14.52%	0.15	Yes	—	—	—	—	—

Table 1. All qualifying events, 2021-present. Range pos = (close - low) / (high - low). Forward returns in percent from event close. The July 23, 2026 event qualifies for the filtered cohort.

II. Forward returns

The unfiltered cohort (9 completed events) bounces almost mechanically: +6.5% mean at 5 sessions with a 100% win rate, +8.0% at 10 sessions, then dispersion takes over — +23.3% mean but only 67% positive at 50 sessions against a baseline of +5.1% and 52%. The filtered cohort (5 completed events) is stronger at every horizon and, critically, more reliable exactly where a trade can monetize it: +14.1% mean at 10 sessions, five for five positive, versus +1.1% baseline. At 50 sessions the filtered mean reaches +40.5% (median +54.8%, 80% positive). The filter earns its keep by removing April 2022 — the one panic that occurred above trend and kept falling — while retaining every true liquidation low.

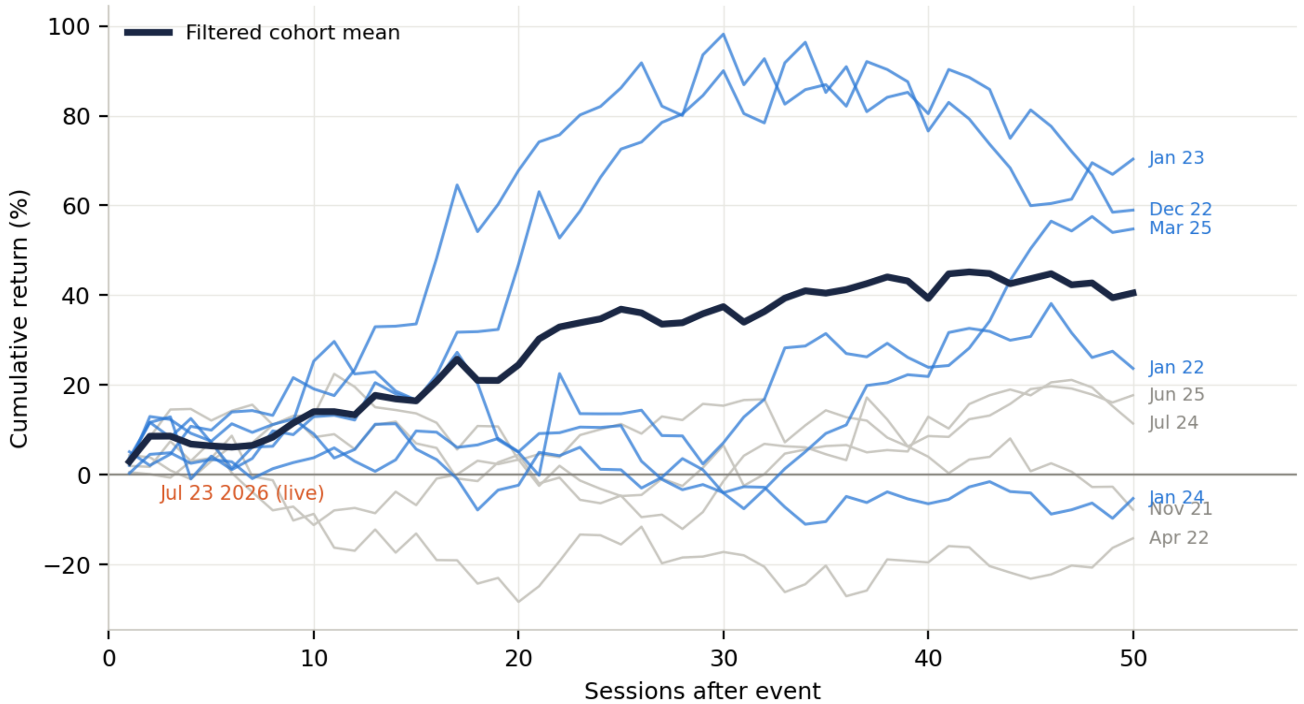

Figure 1. Individual 50-session forward paths. Blue: filtered cohort (prior close below 65d MA). Gray: excluded events. Navy: filtered-cohort mean. Orange dotted: the live July 23, 2026 event (one session elapsed).

III. Statistical validation

Four tests, 20,000 replications each, against the population of all 1,345 overlapping 50-session windows from January 2021 forward.

Test	Design	Observed	Percentile	p-value
1. Random-date placebo	Mean 50d return of 5 random dates	+40.5%	99.6th	0.004
	Median of same	+54.8%	99.9th	0.001
2. Stationary block bootstrap	Synthetic 50d paths, 10-session mean blocks	+40.5%	99.2nd	0.008
3. Cluster-collapsed	Dec 27 / Jan 3 merged; n = 4 episodes	+34.5%	97.9th	0.021
4. Tradeable horizon	Placebo at +10 sessions	+14.1%	98.7th	0.013

Table 2. Bootstrap and placebo results. All tests one-sided against the null of no edge.

The cluster-collapsed test is the load-bearing one. Merging the December 2022 and January 2023 events — whose forward windows share 47 of 50 sessions — drops the effective sample to four independent episodes and the mean to +34.5%, and the result still clears the 97.9th percentile. The cluster inflates the point estimate; it is not what carries the significance. Separately, the five-for-five positive record at 10 sessions has a 3.4% probability of arising by chance given TSLA's 50.9% baseline 10-session win rate. For calibration: the JNK 65-day MA break tested this week (AQI-RN-2026-JNK-01) landed at the 80th percentile — indistinguishable from carry — and the note said stand pat. This cohort clears the 98th on every test, which is what justifies partial size ahead of confirmation.

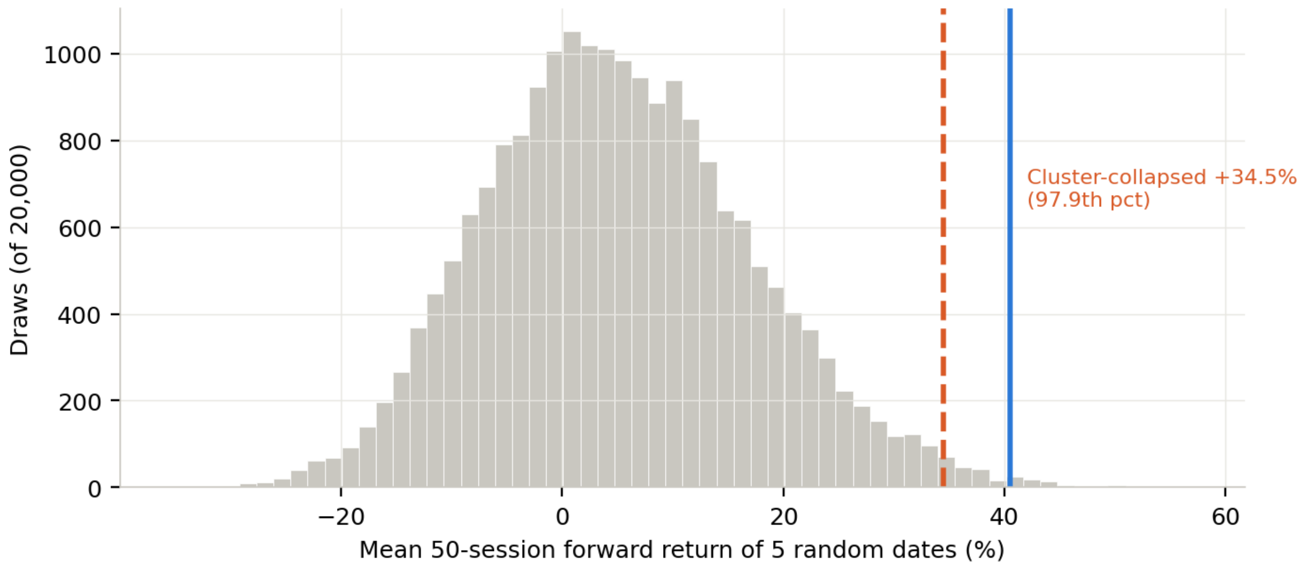

Figure 2. Null distribution of placebo means (Test 1) with observed and cluster-collapsed values.

IV. Honest accounting

Three caveats belong in print. First, n = 5 completed events, effectively four episodes; one bad future instance moves every statistic materially. Second, three of the five best outcomes trace to the late-2022 bottoming cluster, so the median path is the more honest central tendency than the mean. Third, the live event has already broken pattern once: July 24 closed -2.7%, only the second next-day decline in the ten-event sample — either a better entry than the pattern usually offers, or an early tell toward the January 2024 dud path (the one filtered event that finished -5.3% at 50 sessions). The staging logic below exists so the position never has to decide which.

V. Mechanics

Figure 3. TSLA daily close vs the 65-day MA since mid-2023. Blue markers: filtered events. Gray: excluded events. Orange: July 23, 2026.

Core (sells the reliable part). Short put spread, 30-45 DTE, short strike at or below the event low of $291.72. The statistics being monetized: 100% positive at 5 and 10 sessions in the filtered cohort, median +7 to +13%, sold into post-earnings elevated implied volatility. The event low is the invalidation shelf — every completed filtered event held it, and the analog that did not hold its shelf (April 2022) is precisely the above-trend case the filter excludes. A closing violation ends the trade; no averaging.

Satellite (owns the lottery ticket). Long call spread, 50-70 DTE, risk approximately one quarter of core, financed in part from the put-spread credit. This is the honest expression of a +40.5% mean that decomposes into one dominant cluster: real enough to want exposure, too concentrated to size as thesis.

Staging (borrows the JNK reclaim logic). The event fired below the 65-day MA — that is what placed it in the strong cohort — so the first genuine trend confirmation is a reclaim of the MA within 10 sessions (on or before approximately August 6). Deploy partial size now on the statistical edge; add the second tranche on the reclaim. The cost is entry price; the purchase is meaningful protection against the January 2024 path.

VI. Sizing

Step	Basis	Fraction
Raw Kelly	p = 0.80, avg win +52%, avg loss -5% (filtered cohort, 50d)	~78%
Quarter-Kelly	Standard fractional discount	~19%
Sample-size discount	n = 5 events, ~4 independent episodes; halve again	8-10%
Tranche 1 (now)	Half of discounted allocation on statistical edge	4-5%
Tranche 2 (on reclaim)	Balance on 65d MA reclaim within 10 sessions	4-5%

Table 3. Kelly-Thorp sizing ladder with explicit sample-size discount. Raw Kelly on five events is quoted only to be distrusted; if the effective sample is treated as three independent episodes, even the 8-10% terminal figure is generous.

Decision summary. Enter tranche 1 core + satellite at current levels. Add tranche 2 on a 65-day MA reclaim on or before ~August 6. Exit all on a closing violation of $291.72. Reassess at 10 sessions if neither trigger has fired.

Methodology: BATS TSLA daily OHLC, Mar 2020 - Jul 24, 2026 (1,603 sessions). Event screen and forward returns computed on closes. 65-day simple moving average verified against source MA column (exact match). Bootstrap: 20,000 replications per test; stationary block bootstrap with geometric block lengths, mean 10 sessions. This document is quantitative research, not a recommendation to any person to buy or sell any security. Options involve risk and are not suitable for all investors.