

Reversal vs. Breakout: Cross-Symbol Pattern Analysis

Research findings, 2026-09-18. Prepared for review. No bot code touched by this analysis itself; a resulting prototype exit (exit_velocity_experimental.py) is documented separately below.

Question

Across all 30 of that day's scanned candidates (not just the 9 the bot actually held a position in), is there anything in the recorded order-flow data that reliably distinguishes a stock whose rally is about to reverse from one that's merely pausing mid-breakout? This follows directly from the PRF and resistance-timer exit experiments, both of which assumed imbalance level/direction at a stalled peak was the relevant signal and both of which underperformed the no-extra-exit baseline on every configuration tested.

Method

Source data: data/simulations/2026-09-18_full_stream_record.jsonl (138,180 rows — every real logged poll, ~5s cadence, for all 30 candidates, with price, bid/ask, trade imbalance, buy/sell volume, sub-bar tick counts, and cumulative volume). “Peak” events were defined, after two looser passes that produced no usable signal, as: a fresh intraday high that then fails to be re-exceeded for 15 seconds — the same definition exit_prf_experimental.py itself uses for its own peak tracking. This produced 436 usable events across the 30 symbols (275 later recovered/re-tested that level, 161 did not — a 63.1% base recovery rate).

What did NOT discriminate

Signal	Result
Imbalance level at the stall	Flat, non-monotonic: 51.7%–66.7% recovery across every bucket, strongly negative to strong
Imbalance flipping negative after the stall	Flat, non-monotonic: 57.1%–72.9% recovery
Spread % (top-of-book)	No differentiation at all: 0.26% vs 0.25% before/after, both classes

This is itself a meaningful result: it directly undermines the core assumption both exit_prf_experimental.py and exit_resistance_timer_experimental.py were built on — that imbalance level or direction at a stalled peak carries information about what happens next. It does not, in this data.

What DID discriminate: participation ratio

Defined as (buy_vol + sell_vol) in the 90 seconds after the stall ÷ the same in the 90 seconds before it. Clean, monotonic gradient:

Ratio bucket	n	Recovered
0.00–0.30	52	44.2%
0.30–0.60	73	54.8%
0.60–1.00	107	58.9%
1.00–1.50	84	70.2%

1.50–3.00	72	72.2%
≥3.00	48	79.2%

Collapsed to a single cut at ratio ≥ 1.0 vs < 1.0 : **73.0% recovery (n=204) vs 54.3% (n=232), z=4.04, p=0.00005** — statistically real on this one day's sample. Rising participation after a stall predicts recovery; fading participation predicts failure. A secondary, weaker signal pointed the same direction: sub-bar tick count (trade pace) drops off faster after failed stalls (105.0→87.7, –16.5%) than after ones that recovered (108.1→79.0, –26.9%) — not independently validated as a leading (pre-peak) signal in this pass.

Critical limitation: this does not explain GLOO

Checked directly against the case this entire line of investigation started from: GLOO's own real top (10:58:46 ET, \$5.58) shows imbalance still **+73%** and a near-neutral participation ratio of **0.92** — deceptively healthy-looking order flow right up to the actual reversal. The participation-ratio finding is a genuine aggregate tendency across 30 symbols on one day, not a rule that would have caught the specific case that motivated PRF, the resistance-timer exit, or this analysis itself. Any prototype built from this finding should be understood as a different tool for different cases, not a GLOO fix.

Web research

Absorption / Wyckoff “effort vs. result” / VSA. High volume with little price progress signals absorption at a defended level; a breakout on narrow spread and shrinking volume fails the effort-vs-result test and should be read as a trap until volume confirms continuation.

Order-flow / cumulative delta divergence. Cumulative delta making a lower high while price makes a new high is described as an early tell of an exhausted breakout — directionally consistent with the participation-ratio finding above.

VPIN / order-flow toxicity (Easley/López de Prado/O'Hara). A volume-clock measure of persistent buy/sell imbalance as an adverse-selection early warning; registered its highest readings an hour before the 2010 Flash Crash. Requires volume-bucketed (not wall-clock) sampling — a real implementation gap versus this bot's fixed 5s poll.

Iceberg order detection. Inferred from traded volume at a level exceeding displayed size and repeated absorption without progress — structurally the same pattern this analysis's “stall” events already target, but distinguishing it from ordinary consolidation needs order-book depth this bot doesn't collect.

Pause vs. failure — the exact question asked. The clearest practitioner rule found: watch the VELOCITY of the return move plus its volume. A fast snapback on rising volume is a failed breakout; a move that holds and builds on volume is a genuine pause. Shape/speed-based, not a fixed time window — structurally different from both PRF and the resistance-timer, which each use a flat clock or read-count regardless of how fast price is actually moving.

Faster-than-confirmation / leading signals. “Crumbling quotes” (displayed depth withdrawn faster than replenished, needs book-cancellation data this bot doesn't record) and “order-flow exhaustion” defined as measurable deceleration in trade pace/arrival rate even while price keeps printing — faster than waiting for price or imbalance itself to turn, and the likely explanation for the weaker tick-count finding above.

Prototype built from this: `exit_velocity_experimental.py`

A fifth experimental exit layer, wired into `simulate.py` via `--velocity-exit` (plus `--velocity-min-pct-per-min` and `--velocity-max-participation` to sweep both knobs). Deliberately a different SHAPE of rule than PRF or the resistance-timer: no fixed time window and no imbalance gate at all. Instead, once price is off its own peak (armed the same way, +3% MFE), it fires a single-read exit the moment BOTH: (1) price's decline over a short trailing 30s window exceeds a velocity threshold (%/min), reacting to how FAST price is actually giving back rather than waiting out a clock, and (2) the same 90s-before/after participation ratio from the analysis above is below 1.0 (volume fading, not picking back up) — gating the velocity trigger so an ordinary sharp-but-recovering wiggle with volume already returning doesn't fire.

Backtest results (2026-09-18, max_positions=3)

Baseline (no extra exit layer): 10 trades, 5W/5L, \$559.56. Full sweep of both knobs:

Velocity ≥ (%/min)	Participation < 0.5	< 0.8	< 1.0	< 1.5
1.5	\$385.60 (3 fires)	\$361.52 (3)	\$312.85 (4)	\$268.25 (3)
2.0	\$344.70 (2)	\$344.70 (2)	\$393.55 (4)	\$365.06 (5)
3.0	\$559.56 (0)	\$559.56 (0)	\$523.47 (1)	\$523.47 (1)
4.0	\$559.56 (0)	\$559.56 (0)	\$559.56 (0)	\$559.56 (0)
6.0	\$559.56 (0)	\$559.56 (0)	\$559.56 (0)	\$559.56 (0)

“0 fires” cells trivially tie the baseline because the layer never triggers at that threshold, not because it correctly stayed out of the way of a real reversal — GLOO's own decline never exceeded ~2.9%/min over any 30s window that day (see limitation above), so any threshold high enough to avoid false triggers is also high enough to never fire on the one case this line of work started from. Every configuration that actually fires (velocity ≤ 3.0) finishes the day below baseline. The least-bad active config, velocity ≥2.0%/min + participation <1.0 (\$393.55, 4 fires), still clips a genuine GLOO continuation: it exits a re-entered GLOO position at \$5.14 (14:33 ET) on a real but temporary 2.9%/min dip — price went on to \$5.49 before the actual later reversal. Same failure mode as PRF and the resistance-timer: fast enough to react is fast enough to also cut a stock that was simply having a rough 30 seconds inside an intact move.

Recommendation

- Do not wire `exit_velocity_experimental.py` into the live bot — it joins PRF, the giveback exit, and the resistance-timer as a sixth-day-running result below the no-extra-exit baseline on 2026-09-18. This is the fourth distinct rule shape tested against this exact problem (structure/imbalance state machine, fixed giveback floor, fixed recovery-window clock, and now short-window velocity) with the same outcome.
- The participation-ratio finding underneath it is still the first cross-symbol signal in this investigation with a clean statistical gradient rather than noise around a base rate — but as implemented here (a memoryless single-read velocity trigger) it doesn't translate into a working exit, and it explicitly fails on the motivating GLOO case regardless.
- Before any further tuning: validate against a second day (`data/decisions/2026-09-17.jsonl` and `data/candidates/2026-09-17_scanner.json` exist; tick cache is not yet fetched for that day).
- Most promising untested idea from the research: a genuine pre-peak leading version of trade-pace deceleration (fires before the stall completes), rather than the post-hoc average this pass computed.

- VPIN and crumbling-quotes are blocked on data this bot doesn't currently record (volume-bucketed sampling, order-book depth/cancellations) — not directly implementable today.