

## □ R&D Council Memo — April 29, 2026 (Wednesday)

**Status as of this memo (12:00 UTC): - Quant paper (per Apr 28 trade report):** Model A NAV **\$1,003,527** (11 longs / 5 shorts / \$204K cash) | Model B **\$999,997** (30 longs / 8 shorts / \$271K cash). Combined ~\$2.0M. *Note: prior memos quoted \$5-7M aggregate which was a stale figure — Apr 28 close is the authoritative number.* - **IBKR LIVE:** Funded \$1K seed + ~\$110K Chase wire en route. **STOP file still in place** at scripts/ibkr/STOP (kill-switch on from Apr 27 hardening). No first-rebalance yet. - **Night Shift (06:00 UTC today) delivered:** executor.py --config-v1, dispatch-telegram.py (live test passed, 1 real msg landed), daily-runner.py wired to dispatcher, 9/9 unit tests on the v1 filter, 4 stale alerts archived. □ - **Yesterday's R&D ask (fast-breakdown patch):** Spec'd, backtested, code drafted, **NOT YET MERGED.** ENABLE\_FAST\_BREAKDOWN is not in quant-models/config.py. This is the blocking item for short-side improvements. - **Market regime right now:** SPY 711.69 (+1.08% 5d, +3.29% vs 20MA), VIX 17.97 (<18). **Confirmed bull.** Per yesterday's spec, only fast-breakdown shorts allowed; deep-breakdown + index/sector shorts gated OFF. The 5 active shorts in Model A (SNOW/KTOS/SQQQ/NOW/ACN) opened under the OLD logic — see CIPHER below. - **Live fast-breakdown candidates today (yfinance, 12:00 UTC):** CHTR 5d **-28.45%** / -21.2% vs 20MA, LULU 5d **-14.58%** / -9.7% vs 20MA, LCID 5d **-17.44%** / -26.5% vs 20MA. Same three names as yesterday — still actionable.

---

## □ ADVISOR DEBATE

### ATLAS (Strategist)

Two memos in a row have ended with “approve the patch.” It hasn't been approved or rejected; it's just been *unread*. That's a third-day loop, and it's the same shape as the IBKR funding decision (5 days unanswered) and the Ocala builder RFP (3 memos). Carlos's pattern this week is: night-shift unblocks technical things while strategic decisions sit. Today the council should **stop adding new strategic asks** and instead **collapse the pending decisions into a single 3-tap Telegram**. One message, three buttons, three decisions cleared: 1. **Patch:** ship fast-breakdown to paper-test? (Y/N) 2. **STOP file:** remove and go live with --config-v1 today? (Y/N) 3. **Ocala RFP:** I'll draft and send to 3 builders by Friday — approve? (Y/N)

If we ship today's memo as another wall of analysis, we get another zero-response day. Force the decision surface to be one screen, three taps.

### CIPHER (Quant)

Three things to flag on the books:

1. **The 5 active shorts (SNOW/KTOS/SQQQ/NOW/ACN) are running under the OLD logic.** The Apr 28 trade report shorted them at scores 25-32, but with SPY now in confirmed-bull (above 50MA, VIX <18), the proposed check\_short\_regime\_tiered would have flagged today's environment as “do

**not open new index/sector shorts.”** SQQQ in particular is a levered NDX-inverse — that’s exactly the kind of position the regime gate is supposed to block. Recommendation: **don’t close them today** (would be reactive, slippage + tax inefficient), but the new logic — once shipped — would have prevented them being opened. The patch is doing real work conceptually before any code lands.

2. **The fast-breakdown candidates are still alive.** CHTR -28.45% / 5d, LULU -14.58%, LCID -17.44%. All three would qualify under the proposed thresholds today. CHTR has now been a candidate for 2 sessions running and price has stabilized slightly (-28.45% vs -28.50% yesterday) — meaning the entry window is closing. If the patch had merged Tuesday and paper-tested Wednesday, we’d be sizing into LULU/LCID with 2 days more conviction.
3. **Cost of waiting one more day:** roughly negligible *expected* P&L impact on a single name (avg fwd 5d return on fast-breakdown shorts in backtest: -0.24% per name  $\times$  2.5% sizing  $\times$  \$1M NAV = ~\$50/day per name not yet shorted). So strategically the patch matters; tactically a 1-day slip is not catastrophic. **But** — three names dropping 14-28% in 5 sessions while the index makes new highs is exactly the kind of breadth-divergence signal the patch was designed for. This is the live test the backtest predicted.

## SPARK (Creator)

Two notes:

1. **Day-1 IBKR Reel is still pre-loaded and unfired.** Wire hasn’t landed → trigger hasn’t tripped → Reel is sitting in Postiz queue. Tracking the wire arrival via IB Gateway is a SPARK responsibility now — when balance jumps from \$1K to \$111K, that’s the trigger. I’ll set a `daily-runner.py --notify-on-deposit` hook for FORGE if needed.
2. **New angle just opened: “My AI saw this 2 days before the news did.”** If we ship the patch and the next memo includes a paper trade on CHTR/LULU/LCID at today’s prices, then 1-2 weeks from now we have a clean before/after content piece: “April 28 — my AI flagged CHTR as a short. April 29 — I approved a code patch. May 6 — here’s what happened.” That’s the most honest, defensible AI-trading content I’ve seen in the queue. It only works if we ship.

## FORGE (Operator)

Three engineering items:

1. **scripts/ibkr/STOP is still in place.** Removing it is a 1-line action (`rm scripts/ibkr/STOP`). The night-shift log explicitly listed this as Carlos’s morning task. Without removal, `--config-v1` refuses to fire any orders even if Carlos runs it. **This is the single biggest blocker to live deployment, more than the wire itself.**
2. **The fast-breakdown patch is ready to ship as a single PR.** Estimate ~45 minutes of work: add 4 config constants, add `detect_fast_breakdown()` and

check\_short\_regime\_tiered() to quant\_engine\_lib.py, wire into the scoring loop with the bottom-quartile-momentum prefilter (FORGE's rate-limit mitigation from yesterday), add 6-8 unit tests. **ENABLE\_FAST\_BREAKDOWN = False** by default — the merge does not change live behavior; flipping the flag is a separate decision.

3. **Borrow-availability check** — yesterday's memo flagged this for the IBKR executor. I confirmed: executor.py does not currently call reqMatchingSymbols or check shortable inventory before placing a SHORT order. If a fast-breakdown candidate is hard-to-borrow (LCID is on multiple HTB lists), the order will reject silently. **Add an ib.reqContractDetails() + shortable check before any SHORT in executor.py.** ~30 min add. Should ship same PR as the fast-breakdown patch.

---

## □ TOP 3 RECOMMENDATIONS

### 1. Collapse all pending decisions into ONE 3-tap Telegram — get a same-day yes/no on the patch, the STOP file, and the Ocala RFP

**Proposed by:** ATLAS (debate consensus) | **Impact:** Unblocks 3 stalled tracks at once; estimated >\$25K of cumulative weekly opportunity cost cleared | **Effort:** Low (5 min to draft) **The idea:** Yesterday's memo, today's memo, and yesterday's night-shift log all asked Carlos for the same three things in different paragraphs. The friction is interface, not content. Send a single Telegram with three numbered yes/no questions and tell Carlos he can reply "1Y 2Y 3Y" in 6 keystrokes. No more analysis paragraphs in the asks themselves; analysis is in this memo for reference. **Debate summary:** All four advisors agree this is the right move. CIPHER notes the patch decision is the most time-sensitive (live signal active 2 days). FORGE adds: STOP-file removal must follow patch decision (rm STOP without the patch means tonight's run uses the old short logic on the SAME candidates that already triggered Apr 24 cushion losses). **Next steps:** This memo's Telegram summary IS that 3-tap. Below in the Telegram block.

### 2. Ship the fast-breakdown patch + borrow-check in ONE PR today — paper-only via ENABLE\_FAST\_BREAKDOWN=False

**Proposed by:** FORGE + CIPHER | **Impact:** Validated backtest = +0.24% avg fwd 5d on shorts, ~3 candidates/week; on \$1M NAV at 2.5% sizing = ~\$15-25K/yr in expected short alpha | **Effort:** ~75 min (45 min patch + 30 min borrow check + tests) **The idea:** Land the code with the flag OFF. Zero behavioral change. Carlos can flip the flag when he's reviewed the live signal vs. paper output for 5 sessions. Bundling the borrow-availability check ensures that when the flag flips, we don't get silent rejections on HTB names (LCID being today's case study). **Debate summary:** SPARK supports because it unblocks the content angle. ATLAS supports because it converts an open decision into a reversible one (flag is the gate, not the merge). Only friction: this is night-shift candidate #1 for tonight per yesterday's plan — **decision needed: ship as part of today's R&D run or wait for Wednesday night shift?**

Recommendation: wait for night shift; today's R&D budget is best spent on the IBKR live deployment unblock. **Next steps:** Tonight (06:00 UTC Thursday), night shift ships the patch + borrow check + tests + paper-test invocation. Memo Thursday will show the paper output side-by-side with live shorts.

### 3. Pull TurboTenant vacancy list TODAY — peak leasing season is half over

**Proposed by:** ATLAS (4th memo running) | **Impact:** 15 vacant houses × ~\$1,800/mo = \$27K/mo of NOI not landing. One week of vacancy cleared = \$1.7K of net rent recovered | **Effort:** Low (TurboTenant login + 15-min review) **The idea:** This has been on the carry-over list since Apr 25. Spring (April-early June) is the highest-conversion leasing window in Florida SFR. Every week of vacancy in late April converts roughly 1.5× to weeks of vacancy stretching into the slow summer season. Pull the list, identify the 5 longest-vacant, drop their listed rents 10%, re-list. If 3 lease in the next 2 weeks, that's ~\$5K/mo of recurring NOI back online. **Debate summary:** CIPHER frames it as: a 10% rent cut on a 90-day-vacant house is a +infinity% IRR on the immediate cashflow vs. the vacant baseline. SPARK: "Florida Spring rental tour" reel from one of the rentals doubles as inventory marketing. **Next steps:** Carlos logs into TurboTenant (creds in TOOLS.md), pulls vacancy report, replies with the list — or grants ATLAS read access to do the pull and recommend the 5 to discount.

---

#### ☐ QUICK WINS (do today)

- **Reply to the 3-tap Telegram below** — 1Y 2Y 3Y or any combination. Six keystrokes clears the week.
- **rm scripts/ibkr/STOP** the moment you've decided on #2. The dispatcher and --config-v1 are already proven via Tuesday's tests; they're idle waiting on this file.
- **Cron the dispatcher during market hours:** \*/5 13-21 \* \* 1-5 cd /home/ubuntu/.openclaw/workspace && python3 scripts/ibkr/dispatch-telegram.py --once (per night-shift recommendation — 30s setup).
- **Pull TurboTenant vacancy list** (or grant the council read access).
- **Pre-fire the Day-1 IBKR Reel** — wire is expected today/tomorrow; do not let the moment pass.

#### ☐ IDEAS IN DEVELOPMENT

- **Slippage tracking schema** — needed before live promotion of fast-breakdown shorts. Schema sketch: data/slippage/YYYY-MM-DD.csv with ticker, side, signal\_price, fill\_price, ms\_latency, signal\_type. FORGE can ship alongside the patch.
- **Sector-rotation short v2** — if fast-breakdown ships clean and sec\_below\_20ma >= 8 fires once, then we've got a real second signal. Today: 2/11 sectors below 20MA — far from threshold. Park this until breadth deteriorates.
- **Ocala builder RFP** — 4-memo carry-over. ATLAS will draft Friday if Carlos approves Q3 of the Telegram. Draft target: apps/ocala-rfp/RFP-2026-05-01.md.

## □ MOONSHOTS

- **Public live track record dashboard** — once IBKR has 30 real fills, publish a redacted statement page on charlywhite.com. This is the one piece of infrastructure that turns “Carlos has an AI” into “Carlos has a fundable AI strategy.” Required for any future LP raise (Gobundance is the obvious first audience).
- **Bunker BTR niche product line** — 52 lots × \$50K equity creation per lot if the prepper-niche premium even half materializes = \$1.3M in latent equity. Needs ONE comp. RFP unblocks this.

## □ This Week’s Focus

**Stop accumulating analysis. Start shipping decisions.** This memo is shorter on net-new ideas than the last 3 because the bottleneck is execution, not ideation. Three patches are spec’d and unshipped (fast-breakdown, borrow-check, slippage-tracking). One file deletion (STOP) blocks the entire IBKR live deployment. One TurboTenant login resolves \$27K/mo of bleeding NOI. One email resolves a \$1.84M idle land position. None of these need more thinking — they need a “yes” and 30 minutes. **Today’s success metric: the 3-tap Telegram gets a reply, and at least 2 of 3 are approved.**

---

## □ QUANT MODEL REVIEW

**Model A** (paper NAV \$1,003,527 | Cash \$204,438 / 20% | 11 longs + 5 shorts as of Apr 28 close) - **Top-10 scorers (Apr 28):** CGNX 90, EQIX 86.7, TER 85, BLSH/BWET/CIEN/NBIS 82, LUNR 81.7, ADI 81.6, IRM 81.1 - **Bottom-5 (current shorts):** SNOW 31.8, KTOS 29.4, SQQQ 28.8, NOW 26.9, ACN 25.0 - **CIPHER:** SQQQ shouldn’t be in this book in confirmed-bull regime; if patch ships, the next rebalance flags it for cover. SNOW and ACN are score-driven shorts with deep momentum decay (Mom60 -22%/-32%) — those stay valid even under stricter regime gating because they’re score-bottom, not regime-bet, shorts. KTOS at 28 is borderline; price stabilized in the last 3 sessions, monitor for cover candidacy. - **ATLAS:** Bottom names (NOW, ACN) are Carlos’s direct competitive thesis names — IT services and enterprise software facing AI-driven margin compression. Aligned with the broader thesis. SQQQ is a hedge instrument that doesn’t belong in a long-bias rebalance. - **FORGE:** Engineering improvements, ranked: (1) fast-breakdown patch + borrow check (tonight, Thursday memo), (2) slippage tracking schema (must precede live promotion), (3) IB Gateway auto-restart (was crashing twice in test runs; needs supervisord wrapper before unattended live runs), (4) --config-v1 synthesizer for missing-from-report tickers needs a unit test on the case where fewer than 10 tickers exist in a stale report (edge case from yesterday’s test pass). - **SPARK:** CHTR/LULU/LCID is the content trifecta if the patch fires this week. NOT yet mentionable publicly because nothing has shipped. Hold the post.

## □ Quant Model Actions

- **Model A: HOLD** the 5 existing shorts; do not close reactively. Patch evaluation will determine whether SQQQ specifically should be covered next rebalance.

- **Model A: Pull universe-coverage check** on top-10 scores vs. deployment-config-v1.json top-10. If divergence >2 names, alert Carlos.
  - **Model B:** No changes. Continues as the diversifying paper-only reference (30 longs / 8 shorts / commodity & international tilt).
  - **Universe (both models):** No new short candidates approved for live execution today. CHTR/LULU/LCID earmarked for first paper-fire AFTER Thursday morning patch merge.
- 

## □ BUSINESS ACTIONS

- **IBKR (CRITICAL):** (a) rm scripts/ibkr/STOP, (b) check IB Gateway for wire arrival (~\$110K Chase→IBKR expected today/tomorrow), (c) when wire posts, run python3 scripts/ibkr/daily-runner.py --live --config-v1 (or --dry-run first), (d) confirm dispatcher cron is set, (e) trigger Day-1 Reel.
  - **Quant patch (PRIORITY 1):** Decision needed: ship tonight's night shift (recommended) or accelerate into today's day-shift run. ATLAS recommends night shift to keep R&D council and engineering separated.
  - **Ocala SFR (15 vacant houses):** Pull TurboTenant vacancy list TODAY. Discount the 5 longest-vacant by 10%. 4-memo carry-over.
  - **Ocala BTR / 52 lots:** Builder RFP — Q3 in Telegram. If approved, ATLAS drafts Friday, sends 3 emails by EOD Friday. **Question for Carlos (re-asking):** were the 5 closings (3/27) completed? Portfolio = 72 or 77 houses? (3rd memo asking)
  - **Schwab/HSBC margin (\$52.3M @ 4.25%):** No action this week. Will re-review on the 1st-of-month memo.
  - **Content/Social:** (a) Day-1 IBKR Reel pre-loaded — fire on wire, (b) Tuesday Angel Publishing thread (carry-over from 4/27 — confirm with Carlos: did this post or not?), (c) "AI flagged CHTR before the news" content holds until patch ships.
  - **Legal (divorce):** No surfaced update. Carry-over question: any deadline this week?
  - **Health:** Apple Health Auto Export still STALE (last sync flagged Mar 23 in 4/28 memo). Re-sync the iOS app — 30 sec. Skipping further protocol commentary until data refreshes.
- 

*ORACLE (Grok) contrarian module — XAI\_API\_KEY not loaded in this subagent environment. Carlos can re-run manually with: source ~/.bashrc && python3 scripts/grok-council-advisor.py memory/rd-council/2026-04-29.md apps/rd-council/context-brief.md*