02 Stopping Apex

Five strategies, zero dollars lost: ending a research program on evidence

PERIOD
February – June 2026

STATUS
Decommissioned 2026-06-16

MY ROLE
Set the business bar, ruled on pre-registered verdicts, drove the decision to stop

IN ONE LINE
5 strategy generations tested and killed · paper only · $0 real capital at risk

Steven Fellows · Co-Founder, AI Strategy & Operations, Jiyu Technology Inc · stevenpfellows@gmail.com · 206-483-4426 · stevenpfellows.com · One-page extract from the AI Transformation Portfolio; figures are from company records.

BUSINESS PROBLEM

Apex was a research program: a paper-trading platform for Kalshi prediction-market contracts, built to find out whether a durable edge existed before any real capital was committed. By spring 2026 it had a real engine, a dashboard, a database recording millions of evaluation decisions, and a team that had spent months on it. Live trading was never enabled. The question was no longer "can the system trade on paper?" It was "is there a business here, and who will say no if there is not?"

SYSTEM AND APPROACH

Every strategy generation was pre-registered: hypothesis, kill and advance criteria, and the sample size required, written before the P&L was seen. The first generations (15-minute directionals; eight epochs, about 100 days, 22 hypotheses) reached fee parity and no edge. Strategy 3 (an hourly bracket curve model) reached calibration and stayed unprofitable. For Strategy 4 I reset the bar from "prove capability" to a business outcome: roughly $100 a day of meaningful income. That bar funded a venue-wide census against a full year of real prints (all five taker families failed the confirm bar) and a maker fill simulation (the passive quoter lost money in the venue's two best cells, minus $2.64 and minus $6.31 a day, at every size). A pre-committed exit written on 2026-06-10 said that if both halves failed we stop. Both failed. We stopped.

MY ROLE

I held the business decisions. I set the income bar that replaced the capability bar, ruled on the pre-registered verdicts (recording the taker path as closed venue-wide and approving the maker study as the last live path), and drove the decision to stop when the pre-committed exit fired. My co-founder built and operated the platform and executed the decommission on 2026-06-16, including a later, separate kill of the successor research program in August, which was his call and not mine. Every figure here is from the company's closing memo.

OUTCOME

Five strategy generations tested and killed with evidence. Zero real capital at risk at any point. The engine, dashboard, and API were stopped and disabled, the database volume deleted, 31 open issues closed, and the repository archived. The methodology (pre-registration, pre-committed exits, adversarial recompute, business-outcome bars) was documented and carried into the next research program. The closing memo's own summary: this was not a failed build; every gate the system was engineered to enforce worked.

ENTERPRISE RELEVANCE

Enterprises rarely kill AI pilots; they let them fade, which costs more. This case is the template for a pilot exit: pre-registered success criteria tied to business value, a pre-committed stop clause, an independent recompute of claimed results, a named executive who rules, and a closing memo that turns a dead project into a reusable prior. It is also what AI value realization looks like when the honest answer is "no value here at our cost structure."