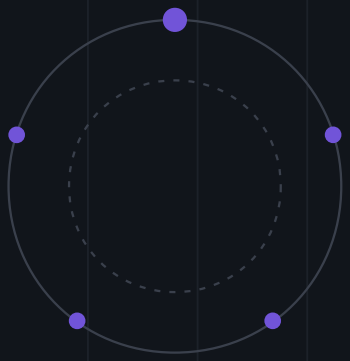


PUBLIC EDITION

QSI Labs

An AI-native quant neolab focused on solving RSI (recursive self-improvement)

Every trading decision gets graded by the market. The loop learns from the grades. Every trade passes a deterministic, AI-free decision system built to catch AI's own mistakes.

SHUBHAM SINGH, FOUNDER
007shubhamkumarsingh@gmail.com

WEBSITE
quantitativesuperintelligencelabs.com

DATE
October 2026

IMPORTANT NOTICE. This paper is for information only. It is not an offer to sell or a solicitation of an offer to buy any security or interest in any fund, and it is not investment, legal or tax advice. All performance figures are internal and unaudited. Past performance is not indicative of future results. Targets and hypotheses are labelled as such and are not forecasts.

§ 01

The claim

> **QSI Labs is an AI-native quant neolab focused on solving RSI (recursive self-improvement). The company is a training loop.** Our agents make forecasts, propose strategies and propose trades. The market grades every decision. Those grades feed back into the loop, so each turn of it starts from more real outcomes than the last. QSI Labs makes money by trading, nothing else.

Every trade passes through a deterministic decision system that contains no AI. Models propose; they never decide. Between every model output and the market sits a rule-based system built for one job: catching AI's own bad decisions before they cost money. The same inputs always give the same decision, and every decision is logged. It runs live today. It is internal and never sold.

> **An early live record from the QSI Labs trading system: about two and a half months, small capital.** The record is the output of QSI Labs' own quant trading system: our agents found the strategies, and the deterministic decision system approved and ran them. Over about 2.5 months, from August to October 2026, three such strategies returned +16.98% net of fees and costs, about 112% annualised (compounded from 2.5 months; not a forecast), with a Sharpe ratio of 2.58 and a maximum drawdown of about 5%, over about 800 trades, on about $10k of my own capital, trading on Kalshi, Interactive Brokers and Public. These are internal, unaudited figures, and past performance is not indicative of future results. The agents currently run on ChatGPT and Claude models; those are only their underlying models today. Section 7 sets out what the record does and does not show.

> **The headline milestone is risk-adjusted returns by Month 6 of the plan:** a live and paper record that clears a Sharpe target of 2, with returns that beat the S&P 500, after deflating for every strategy variant we tried. Month 1 is the first month of the plan.

The QSI Labs trading system is live: agents find strategies, and the deterministic decision system approves and runs them. The next step is proving the full loop, that the market's grades make the agents better. This paper sets out the thesis and what would prove it wrong, the opportunity, the approach and how it improves itself, the live record, why now, the founder, and, at length, what can fail.

§ 02

The thesis, from first principles

When every firm runs the same models, the edge moves to the decision record and the rate of learning from it.

Frontier AI commoditises research capability. Reading filings, writing code, running backtests and drafting a trade thesis are becoming things any firm can rent. The largest funds are already rolling out the same frontier models to their teams [1][2][3][4]. A capability everyone can buy is not an edge. Worse, the same model given the same data finds the same ideas, so commodity research tends towards crowded trades. Shared ideas decay fast: published trading signals have lost more than half their returns once widely known [22]. So the edge can't be the model itself.

Markets grade everything. A forecast resolves. A strategy makes or loses money next period. A prediction market settles yes or no. Every decision gets a label, every day, at the cost of the data needed to score it.

So the scarce asset is the proprietary, time-stamped, market-graded decision record. A decision made in real time, before its outcome existed, with a real information set, real costs and real fills, cannot be bought or backfilled later (Section 8). Each firm's record is its own. It is the one input to an AI-run firm that rivals cannot rent.

Markets never finish, so the loop compounds. Coding and maths have fixed answers that can be checked against a fixed body of examples. Markets have no final dataset. Companies list, regimes shift, participants adapt, strategies decay. New information arrives daily. A firm that learns from its graded decisions faster than the market changes keeps its edge; one that stops learning loses it. The edge is the rate of learning, not the starting model.

Learning has to go deeper than prompts. Prompts and memory help a model look things up. They don't make it a better judge. Post-training does.

The asset is the outcome data and the recipe, not one model. Base models improve every few months. A fine-tune of today's model will lose to tomorrow's generic model. So we re-apply the recipe to every new base. What compounds is our record of which experiments and strategies proved good or bad, and the method for learning to judge them.

A deterministic decision system is what makes an AI-run book safe enough to learn from at all. An AI-run book fails differently from a human one: confidently, quickly and at scale. Without a gate, its errors reach the market, and two things break. The losses can end the experiment before it teaches anything. And the grades are polluted: the market scores a stale price or a flipped sign, and the loop learns from its mistakes' noise instead of its judgement. A deterministic system blocks those errors before an order exists, and because its decisions are reproducible from logged inputs, every grade traces back to exactly what the system saw. That is what makes the record clean enough to learn from.

> **What would prove the thesis wrong.** Each claim has a test, a pass bar fixed in advance and a kill criterion (Section 13):

- **The loop doesn't learn.** Model versions show no gain over the first version on our scored forecasts, or the gain vanishes once the base model is held fixed (Month 6).
- **The record has no special value.** Our agents do no better than the same base model in a generic agent harness, without our graded record (Month 12).
- **Forecast skill doesn't become money.** Gains between model versions don't survive into net, after-cost portfolio returns (Month 12).
- **The decision system adds nothing.** Ideas it approves do no better on forward data than the rejects it tracks (Month 12).
- **Returns aren't real.** The live and paper book fails a Sharpe of 2 after deflating for every variant tried (Month 6, with a second attempt at Month 12).

If those tests fail, I stop and say so. If they pass, the thesis has survived the tests it could most easily have failed.

---

§ 03

The opportunity

QSI Labs trades in one of the largest pools of capital in the world, at the moment its research layer is being commoditised.

The arena is large. Global hedge fund capital reached a record $5.6 trillion at the end of Q2 2026, according to HFR [5]. Non-bank trading firms generated an estimated $114 billion of revenue in 2025, according to Crisil Coalition Greenwich [6]: $84.3 billion from proprietary trading and investing their own capital, and $30.2 billion from market making [7]. These figures describe the industry, not QSI Labs' market. Our size is bounded by strategy capacity (below). What they show is that trading on a firm's own capital is already a very large business.

Incumbents are adopting AI, inside their old structure. About 95% of Balyasny's roughly 180 investment teams use OpenAI's research platform [1]. Citadel's NXT program held its first AI hackathon [2]. Millennium and Anthropic will co-develop an AI risk analyst [3]. Coatue built an internal AI "Brain" with earnings, meeting, brainstorming and expert agents [4]. This confirms that AI now matters to returns. It also shows the pattern: agents are bolted onto structures built for humans. Analysts still analyse, portfolio managers still decide, and AI makes them faster. In my own conversations with firms including Citadel and Verition (first-hand, not published), interest centred on frontier models and internal agents, not on training finance-specific models.

Why the opening exists now. If research capability is a commodity (Section 2), a faster analyst is worth less every year, and the lasting asset is a graded decision record and a loop that improves itself from it. That loop is hard to retrofit. It needs every decision logged at the time it is made, gated by a system that no model can override, and fed back into strategies and a research harness the firm owns. No incumbent has yet built a firm where the training loop is the business. Two startups, Standard Signal and Prodigy, are going after the same opening and started first (Section 10). That is evidence the opening is real, and a risk to us (Section 14).

How QSI Labs makes money. By trading its own capital. It does not sell software or the decision system. Selling the loop would hand over the outcome data that makes it work.

What bounds it: strategy capacity, not market size. Any one strategy can only absorb so much capital before costs eat it. The firm's ceiling is:

independent strategies the loop keeps producing × capital each can absorb × net alpha.

None of these is measured yet. Capacity will be estimated per strategy, conservatively, by Month 12. The loop's job is to increase the first factor and protect the third. Growth comes from more independent strategies, not from bigger bets on the same ones.

How capital grows, step by step.

1. **Now:** QSI Labs trades only its own capital, on a small book. The plan proves the loop learns on forecasts, prediction markets and paper portfolios; live trading stays small.
2. **Next, once the tests pass:** QSI Labs trades the firm's own balance sheet, under a structure counsel confirms first (Section 14).
3. **Later:** outside capital only with an audited record and the right fund and adviser structure.

A possible second line, later. Once the live record proves the models, QSI Labs could offer its models to hedge funds. Not before: a model is only worth buying if its live record is real. The decision system and the outcome data would stay internal either way. This is an option, not part of the plan, and no revenue is assumed from it.

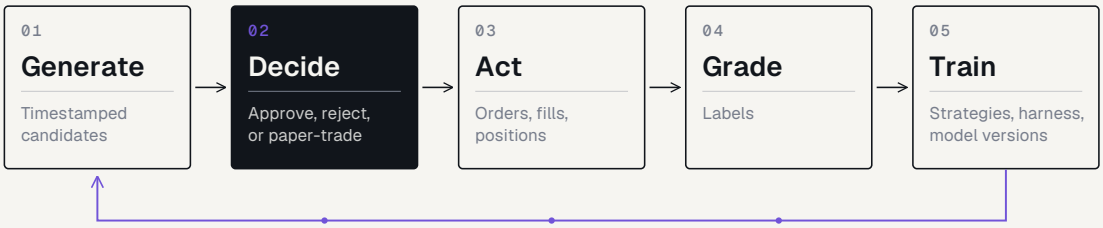
Targets, labelled as targets. The Month 6 headline (a deflated Sharpe of at least 2, beating the S&P 500, Section 13), the hypothesis that post-training will widen our agents' edge over the frontier models they run on today (Section 4) and the capacity estimates are all targets or hypotheses, not results. The only measured result today is the live record of about 2.5 months in Sections 1 and 7.

§ 04

The loop

STEP	WHAT HAPPENS	OUTPUT
Generate	Our agents produce forecasts, research ideas, strategies and proposed trades	Timestamped candidates
Decide	The deterministic decision system checks every candidate. No model approves a trade	Approve, reject, or paper-trade
Act	Approved ideas trade small or on paper, inside limits no model can change	Orders, fills, positions
Grade	The market resolves every decision, including a sample of rejected ones	Labels
Train	The grades retire, refine and replace strategies, and will drive improvements to the agents' research harness. Post-training separately makes the agents better researchers	Better strategies and harness; new model versions

FIG. 01
THE LOOP



What gets graded. Forecasts are scored against what actually happened. Prediction-market calls are scored against how the market settled. Strategies are scored on forward results over windows fixed before they start. Trades are scored on realised profit and loss, net of costs. The system, not the model, sets the questions, so a model cannot pick easy ones or decline to answer.

What is post-trained, and when. Today our agents run on frontier models (ChatGPT and Claude), and that system produced the live record. You can't post-train weights you don't own, so the next step is post-training an open-weight model to make the agents better researchers. It has not started yet and is not yet what trades. The question we will test is whether post-training makes the agents better than the frontier models they run on today, and whether bigger or better-trained models widen that edge. That is a hypothesis, not data; Section 13 sets out how it is measured. Scale only helps in the right direction: a bigger model counts only if it makes better decisions the market can grade.

Why better forecasts must become money. Better forecasts are not money on their own. Forecasts become signals, signals become candidate strategies, the decision system approves or rejects each, and approved strategies enter a portfolio on paper first, then small live. Every stage is graded, so we can see where skill is lost. If forecast gains don't survive into net, after-cost returns, they are academic.

FIG. 02
FROM FORECASTS TO MONEY



The loop is designed to improve itself. Section 5 sets out how, level by level: strategies, the research harness, post-training and, later, pretraining, each gated by the decision system.

§ 05

Recursive self-improvement, level by level

The loop is designed to improve itself at four levels. Only one of them is live today.

Recursive self-improvement means a system that gets better at getting better. For QSI Labs it is the core of the thesis: each turn of the loop (generate, decide, act, grade, train) should leave the next turn stronger. Two things keep it from optimising itself into risk:

- **An external grader.** The market grades every decision. No part of the loop marks its own work, and the market can't be persuaded.
- **A fixed safety gate.** Every self-modification, whether a new strategy, a harness change or a new model version, reaches capital only through the deterministic, AI-free decision system. The loop can propose changes; it can never change the gate. The gate's rules change only by reviewed code, written by people.

LEVEL	WHAT IMPROVES	WHAT DRIVES IT	STATUS
Strategy	Which strategies trade; decayed ones are retired, weak ones improved or replaced	Forward results, graded by the market	Live today
Harness	The agents' research harness: tools, prompts, memory, research workflow, evaluation scaffolding and, next, the RL environments agents will practise in	Graded outcomes of the research it produces	Built and improved by us today; self-improvement from grades is designed and next
Model: post-training	The agents' research ability, by post-training an open-weight model we control	Post-training runs; each version is judged on our market-graded forecasts	The next step; not started
Model: pretraining	What the base model knows about markets	Market and financial data at scale	Longest horizon; only if post-training proves out

Strategy level: live today. Our agents find strategies; QSI Labs' own trading system runs the ones the decision system approves; the market grades them. Strategies that decay are retired, and the loop's job is to keep finding new, independent ones. Gate: every strategy and every order passes the decision system. How we'd know: each strategy's forward net results, and the deflated Sharpe of the whole book (Section 13). What can fail: overfitting to the current regime; a short record mistaken for skill; as size grows, our own trades moving the prices we are graded against.

Harness level: designed, next. The harness is everything around the model: its tools, prompts, memory, research workflow and evaluation scaffolding. Today we build and improve it ourselves. The design is for the harness to improve from graded outcomes: changes that make the agents' research score better on future data are kept, and changes that don't are reverted. Gate: a harness change can alter how ideas are researched, never what is allowed to trade; every idea it produces still passes the decision system. How we'd know: a change is kept only after it beats the previous workflow on future data. What can fail: harness drift, where many small changes add up to behaviour nobody intended; tuning the harness to recent questions.

Model level, post-training: the next step. Post-training has not started. When it does, we will post-train an open-weight model we control to make the agents better researchers, improving their judgement, not just what they can look up. Gate: a new model version only ever proposes; its outputs pass the same decision system, and it replaces the previous version only if it scores better on data it has never seen. Which records may be used for training is governed by strict rules on what data counts as clean (Section 8); this section does not change them. How we'd know: each version is scored on our market-graded forecasts: the Month 6 improvement curve and base-model controls, and the Month 12 test against the same base model in a generic agent harness (Section 13). What can fail: reward hacking, memorisation and leakage, model collapse from training on the models' own outputs, and compute cost.

Model level, pretraining: the longest horizon, and conditional. Only if post-training proves out, continued pretraining of a base model on market and financial data could deepen what the model knows before post-training begins. This is an option, not a plan. Domain text alone has lost to stronger general models before [18], while models trained on the task itself win (Section 9) [16]. Gate and how we'd know: the same as post-training; a continued-pretrained base is used only if it beats the same recipe on a standard base, on data neither has seen. What can fail: compute spent for no gain, and a base soon overtaken by a newer generic one.

RL environments: practice that keeps getting more realistic. We are building our own RL environments: simulated markets and research tasks where agents practise and are scored. Agents will hill-climb on them continuously. The environments will improve too: whenever live trading shows a gap between an environment and the real market, the environment is updated. Two loops compound: agents improve against the environments, and the environments improve against the market. The environments shape practice; the market stays the final judge, and no agent is promoted on its practice score alone.

Noise as markets make it. Market moves have an explained part, tied to drivers you can name, and an unexplained residual. That residual is not the well-behaved bell curve of textbooks. It is fat-tailed, as Mandelbrot showed in 1963 and decades of later research confirm [24][25], and much of it is impossible to model. So our environments are designed to build in fat tails, sudden jumps and changes of regime, so that agents never learn that tomorrow looks like an average day.

The next crisis is always outside the model. No environment can contain it, so environments are never trusted as tail protection, and we do not yet post-train models for tails; that job belongs to the decision system and hard limits (Section 6).

What self-improvement does not mean here. No level edits the decision system, its limits or its rules. No level counts as working until its measure says so. A level that doesn't improve is reported as such, and the firm keeps running on what does work: if post-training adds nothing, the agents keep running on the best available models (Section 14).

> The loop can improve everything except the gate.

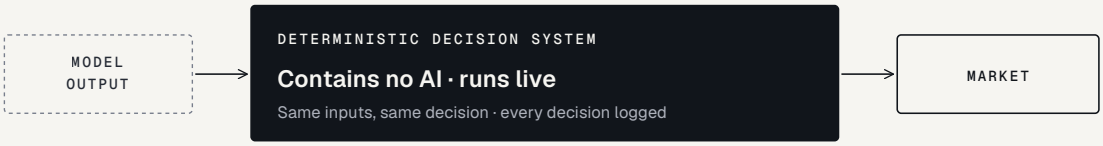
§ 06

The decision system: the differentiator

Models propose. A deterministic system decides.

Our agents will propose strategies by the thousand. Most will be false. Test enough and some look brilliant by luck. And some will be wrong in ways only an AI is wrong: a confident thesis built on a stale price, a flipped sign, a ticker that doesn't exist. Every candidate and every order passes through one deterministic system before it touches capital or our training data.

FIG. 03
THE DECISION SYSTEM



What can be said about it publicly:

- **It contains no AI.** It is rule-based. No model approves a trade, and no model can loosen a rule.
- **It is deterministic.** The same inputs always produce the same decision, and every decision can be replayed from its logged inputs.
- **It is built to catch AI's own mistakes**, as well as the usual market risks: an AI-run book fails confidently, quickly and at scale, so the system is designed around those failure modes.
- **It carries the tails.** Language models are weakest exactly where history is thin: crises and situations they have never seen. So tail protection sits in the system's hard limits, kill switches

and per-strategy loss limits, which hold whatever a model believes. Systematic firms that last work the same way: risk limits set and monitored independently, before the trade and per team, not forecasts [27][28].

- **It runs live today.** Every trade in the live record passed through it.
- **It is tested, not trusted.** It goes through formal, pre-registered calibration tests of the live system, and the ideas it rejects are tracked as well as the ones it approves, so we measure what it wrongly turns down as well as what it lets through.
- **It is internal and never sold.**

How it works, its rules and its limits are not published. They are part of what makes the record clean enough to learn from.

> The models propose. The decision system disposes. The market grades both.

§ 07 The live record

MEASURE	FIGURE
Period	About 2.5 months, August to October 2026
Return	+16.98%, net of fees and costs
Annualised	About 112% (compounded from 2.5 months; arithmetic, not a forecast)
Sharpe ratio	2.58
Maximum drawdown	About 5%
Trades	About 800
Strategies	3, found by our agents
Capital	About $10k of my own capital
Venues	Kalshi, Interactive Brokers and Public
Status	Internal, unaudited

Where it comes from. QSI Labs' own quant trading system, running strategies found by its agents, approved and run by the deterministic decision system. ChatGPT and Claude are only the agents' current underlying models. The strategies started on my own system in August 2026; QSI Labs was formed in September 2026.

What it does not show. About 2.5 months on small capital is not the Month 6 test. Annualising a short window overstates how certain the figure is. The Sharpe still has to be deflated for every variant tried [12][13]. The record covers one stretch of market conditions and says nothing yet about a different regime, a tail event, or a larger book. It was produced before any post-training or automated harness improvement, so it says nothing yet about whether the loop makes the agents better; that is what the Month 6 and Month 12 tests measure.

§ 08 What can't be copied, and why it isn't memorised

Memorisation is controlled by discipline, not argument. Models are only ever scored on data after what they have seen; only data after a model's training cut-off counts as clean; and every claimed gain has to survive controls that separate our loop from improvements in the base model itself. If our gains turn out to be just better base models, those controls expose it. A model scoring history it was trained on can already know what happened next [23]. So the evidence we rely on is live, forward results: a decision made before its outcome existed cannot have been memorised.

Replays are not a moat. Anyone can replay history, and replays are contaminated by what base models already know about that history.

What cannot be backfilled: decisions made in real time, before the outcome existed, with real information sets, real costs and real fills, after the base model's cut-off. Every day the loop runs, it logs decisions nobody can recreate later. That dataset starts small and grows only with time.

This cuts both ways. Standard Signal and Prodigy started before us and already trade live [8][9], so on this logic they are ahead in logged decisions. Time alone is not our moat. What we can claim is how our log is built: every trade passes a deterministic decision system, rejects are tracked, and every gain has to survive base-model controls. Whether that makes our data more useful per decision is what the plan tests (Section 13).

Dense labels at small capital. A small book makes too few trades to judge strategies by profit and loss alone. So the loop runs mainly on forecasts, prediction markets and strategy forward results, which produce far more graded decisions than trades do. Trades give execution realism, not label volume. Proving a strategy from its own returns takes a long time, so we never rely on one: each strategy runs small, on paper first and then small live under its own loss limit, alongside the others, and new ones join as the loop finds them.

---

§ 09

Why now

- Agentic AI now writes code, uses tools and runs long research workflows.
- Open-weight models are good enough to post-train for a domain at startup scale.
- The industry is moving, inside its old structure: Balyasny, Citadel, Millennium and Coatue have all put frontier AI into their existing teams [1][2][3][4].
- **Small trained models already win narrow tasks.** Most of this evidence is company claims, not independent benchmarks. Read together, it is consistent:
  - Harvey and OpenAI built a custom-trained case-law model that lawyers at 10 large law firms preferred to GPT-4 97% of the time [14]. Harvey later reported that seven newer models outperform its original system on its own benchmark [15]. An edge built once decays, which is why we re-apply our recipe to every new base (Section 2).
  - Cursor reports that Composer 2, trained with continued pretraining and large-scale RL on an open base model, is competitive with the strongest frontier models at significantly lower inference cost [16].
  - Checkr reports that a fine-tuned 8-billion-parameter Llama-3 model reached 97% accuracy on background checks against GPT-4's 88% [17].
- **The counterexample is BloombergGPT.** A model pretrained on finance text lost to GPT-4 on financial sentiment benchmarks in a 2023 study [18]. Domain text alone loses. The winners above were trained on the task itself. We post-train for the research task, and the market grades the results.
- Every day a firm doesn't run a loop like this is a day of graded decisions it can never get back.

---

§ 10

Competition

- **Standard Signal (YC Spring 2026)** [8] says its AI researches and executes every trade end to end, learns from market outcomes, logs every decision and enforces hard limits in infrastructure. Its thesis is closest to ours, and it started first.
- **Prodigy Research (YC Summer 2026)** [9][10] says it trains its own quantitative-finance foundation model and trades live with it. Its founders come from DeepMind, Jane Street and Apple. It also started first, and it is funded.

- **KelAI (YC Spring 2026)** [11] sells an autonomous research engine to funds; PMs stay in control.
- **Incumbents:** Two Sigma, Citadel, Balyasny, Millennium, Coatue.

Incumbents could do this. Technically they can, and if one commits fully it will out-resource us. Their public moves add assistants and agents to existing teams.

What QSI Labs builds in from day one. We don't know competitors' internals, and we don't claim they lack anything their public pages simply don't mention. What we build in from day one: a deterministic, AI-free decision system that no model can override; risk built to catch AI's own errors; tracking of the ideas we reject, not just the ones we take; strict discipline about what data counts as clean for each model; and controls that show whether a gain came from our loop or from a better base model. None of these is secret in principle; the bet is that building them in from day one makes our record more trustworthy and our log better to learn from.

---

§ 11

Founder and team

Shubham Singh, Founder.

- **SpaceXAI, Post-Training AI Quant (Nov 2025 – present):** evaluation pipelines and stress tests for long-horizon financial reasoning in post-trained models, with automated grading and failure analysis.
- **Blockhouse, Quantitative Researcher (Sep – Nov 2025):** worked on trading strategies and trading systems.
- **GoQuant, Quantitative Research Lead (Sep 2024 – Sep 2025):** built production ML systems on financial data and agentic workflows that analyse financial data.
- **Research:** AI Researcher, Artificial Intelligence Institute of South Carolina (Sep 2024 – Sep 2025). Co-author of the Alignment Quality Index paper (EMNLP 2025) [19].
- **Exituity:** my previous startup, where I built AI agents for financial analysis. I shut it down in September 2026 and started QSI Labs the same month.
- **QSI Labs:** built the trading system from scratch; it is live today. I have built QSI Labs part-time since September 2026 alongside my SpaceXAI role.
- **Education:** M.S. Computer Engineering, NYU.

How I approach the risks. I have traded systematic strategies and AI models long enough to know where they break: overfit backtests, crowded signals, tails, and models that are confidently wrong outside the data they learned from. None of these goes away with a bigger model. We work on them from first principles and bottom up, starting at the lowest level: agents good enough to run one pod on their own, inside limits they cannot change, and only then more pods (Section 12).

Solo founder, key person, plainly. I am a solo founder today. If I leave, QSI Labs ends today. The mitigation is the first hire: a founding engineer who owns the pipeline, so the system doesn't live in one head, planned for Month 2, followed by an ML research engineer. I am also open to a co-founder and haven't found the right one yet.

§ 12

Roadmap

1. **AI-native** (today): one founder manages strategies; AI agents automate research and run three live strategies.
2. **Agentic workflows** (building): agents run research, directing other agents to generate, test and monitor, inside the loop.
3. **Agentic pods** (target): one loop per strategy or asset class, each run by an agentic PM, all under the same decision system. Humans set objectives and limits. Multi-manager firms already scale through many small, separately limited teams; Millennium runs more than 370 [26].

An agent moves up a level only after it beats the previous workflow on future data.

What has to be true before the book grows.

- A deflated-Sharpe-tested track record over at least 12 months of live trading.
- A reconciled, audited P&L.
- The improvement curve holding when the recipe moves to a new base model.
- Capacity estimated per strategy (Section 3), so capital never exceeds what the strategies can absorb.

§ 13

Milestones and kill criteria

> **Headline milestone, Month 6: risk-adjusted returns.** By Month 6 of the plan, the combined live and paper book clears a Sharpe of 2, with returns that beat the S&P 500, net of costs, and the Sharpe survives deflation for every strategy variant we tried [12]. That is the result to judge QSI Labs by. The Day 90 tests make it trustworthy; the Month 12 tests show whether the loop caused it. The August–October 2026 live record (Section 7) is the starting point for this test, not yet a pass.

Every test is pre-registered: the metric, the minimum sample and the pass bar are fixed and written down before the data exists. These are targets, not results.

BY	WHAT IS TESTED	WHY IT MATTERS
Day 90	The data pipeline has no look-ahead leaks	No model is ever scored on data it has seen
Day 90	Formal calibration tests of the live decision system	Shows it blocks what it should and measures what it wrongly turns down
Day 90	Live P&L reconciled to broker statements, every day	The record is real
Month 6	**Headline: risk-adjusted returns**	Sharpe of 2, above the S&P 500, net of costs, after deflation
Month 6	Improvement curve and base-model controls	The loop makes the agents better, beyond the base model
Month 12	Our agents versus the same base model in a generic agent harness	Our record and loop have special value
Month 12	Approved ideas versus tracked rejects	The decision system's selection adds value
Month 12	Forecast gains versus net returns	Better forecasts become money

Kill criteria, one per claim.

- **Headline missed at Month 6:** a second attempt runs alongside the Month 12 tests. Missed again at Month 12: the firm's own balance sheet does not trade; the book stays small, on my own capital.
- **The models don't improve, or the gains are just a better base model:** stop the loop thesis.

- **Our agents do no better than the same base in a generic harness:** our record and loop add nothing; stop building on them.
- **Approved ideas do no better than rejects:** rebuild the decision system's selection.
- **Forecast gains don't reach net returns:** fix the bridge before any capital scales.

> Every result gets reported, including the ones that go against me.

§ 14

Risks: what can fail

Most of what could kill QSI Labs is listed here, by where it would fail: what failure looks like, how we'd detect it, and what we do.

The order matters. Thesis risk comes first because, if it fires, nothing else matters: the firm would be a small AI-assisted trading book, not a training loop. Where a pre-registered test decides a risk, the test and its kill criterion are named. Where no test can, the system is built to fail safe.

Thesis risk: the idea is wrong.

RISK	WHAT FAILURE LOOKS LIKE	HOW WE'D DETECT IT	MITIGATION, OR WHAT WE DO
The loop doesn't learn	Newer model versions score no better than the first on new data; or the gains are just a better base model	Month 6 improvement curve and base-model controls (Section 13)	Kill criterion: stop the loop thesis. The firm does not keep claiming it
The record has no special value	Our agents do no better than the same base model in a generic agent harness, without our graded record	Month 12 loop versus generic test	Kill criterion: stop building on it
Forecast skill doesn't become money	Forecast scores improve, net portfolio returns don't	Month 12 forecasts-to-money test; per-stage grading shows where skill is lost	Kill criterion: fix the bridge before any capital scales
Post-training adds nothing	Our post-trained open-weight models never improve on the frontier models our agents run on today	Paired comparison on our scored forecasts; Month 6 controls; Month 12 loop versus generic	Keep running the agents on the best available models; the value then sits in the record, the harnesses and the decision system, not in our weights, and we report it

Record risk: the early numbers mislead.

RISK	WHAT FAILURE LOOKS LIKE	HOW WE'D DETECT IT	MITIGATION, OR WHAT WE DO
A short window, possibly luck	The +16.98% over about 2.5 months, about 800 trades across three strategies, reflects one favourable stretch of market conditions rather than skill	Month 6 deflated Sharpe over 120 trading days, counting every variant tried [12]	The record is reported as early and small; it is not a pass until it survives deflation
Annualisation overstates	About 112% annualised is read as a forecast	Not a test: a presentation rule	Always shown as arithmetic compounded from 2.5 months, never as a forecast
Unaudited figures	Internal figures differ from broker statements or from an audit	Day 90 daily P&L reconciliation to broker statements	Every gap explained; an audited P&L is required before the book grows (Section 12)

Model risk: the models learn the wrong thing.

RISK	WHAT FAILURE LOOKS LIKE	HOW WE'D DETECT IT	MITIGATION, OR WHAT WE DO
Overfitting and leakage	Strong scores on history that vanish on future data	Day 90 leak test; deflated Sharpe counting every variant [12][13]	Strict walk-forward scoring; only data after each model's cut-off counts as clean; point-in-time labels
Reward hacking	Models chase easy questions, decline to answer, or hedge every call	Calibration checked for every model version	The system, not the model, sets the questions; scoring rules reward honest probabilities
Too few effective labels	Correlated decisions masquerade as many labels; tests are underpowered	Effective label counts measured by Day 90	A shortfall lengthens the Month 6 window and is reported as a warning; an underpowered test never counts as a pass
Base-model drift	A newer generic base beats our fine-tune of the old one, as happened to Harvey's custom model [15]	Base-model control on every new base	Re-apply the recipe to each new base; the asset is the graded record, the loop and the recipe, not one model

Self-improvement risk: the loop optimises the wrong thing.

RISK	WHAT FAILURE LOOKS LIKE	HOW WE'D DETECT IT	MITIGATION, OR WHAT WE DO
Reward hacking and overfitting to recent regimes	The loop learns to score well on the latest market conditions, or on quirks of how it is graded, and the gains vanish when the regime changes	Calibration checked for every version; scoring only on data the models haven't seen; deflated Sharpe counting every variant [12]	The system, not the model, sets the questions; gains count only on future data; hard limits hold regardless of regime
Leakage and memorisation	A newer model or harness looks better only because it has seen the answers	Day 90 leak test; base-model controls (Section 13)	Only data after each model's cut-off counts as clean; point-in-time labels (Section 8)
Feedback loops from our own trades	As size grows, our own orders move the prices the loop is then graded against	Realised costs and impact split out of every trade and compared with modelled costs	Small size; capacity estimated per strategy; our own impact is measured as a cost, not learned as a signal
Model collapse	Training on the models' own generated reasoning narrows them or degrades them version after version	Every version scored on the same fixed questions against the previous one	Every version is judged against external market outcomes before it is used; a version that scores worse is not used
Compute cost	Training and evaluation cost more than the improvement is worth	Cost of each version against its measured gain	Post-training starts small; pretraining only if post-training proves out; scale only on measured gains
Environment-reality gap	Agents get good at a practice environment that differs from the real market, or exploit its quirks	Practice scores compared with live forward results for the same strategies	Live results decide promotion, never practice scores alone; environments are versioned and updated when the market shows a gap; tails stay with hard limits
Harness drift	Many small harness changes add up to behaviour nobody intended, or to results nobody can reproduce	Every decision logged with the harness version that produced it	Every change versioned, and kept only if it beats the previous workflow on future data; otherwise reverted
The gate becomes the bottleneck	As the loop proposes more, the decision system turns down good ideas or limits how fast the firm can learn	Tracked rejects show what the gate wrongly turns down (Month 12 approved versus rejects)	Gate rules change only by reviewed code, informed by the tracked rejects; the loop never loosens them

Execution and AI-error risk: the system acts on a mistake.

RISK	WHAT FAILURE LOOKS LIKE	HOW WE'D DETECT IT	MITIGATION, OR WHAT WE DO
AI's own errors	A confident order built on a stale price, a flipped sign, a wrong unit or an instrument that doesn't exist	Day 90 calibration tests of the live decision system with planted errors	The deterministic decision system blocks them before an order exists. Agents never send orders; models can't breach the limits
Data faults	Missing, stale or inconsistent feeds; figures revised after the fact	Data checks in the decision system; daily P&L reconciled to broker statements	No trade on bad data: the default is to do nothing. A source that fails a point-in-time check is replaced before training
The decision system errs	It passes noise, or wrongly kills good ideas	Day 90 calibration tests of the live system; Month 12 approved versus rejects	Results published; if selection adds nothing, rebuild it
Market abuse	An agent's order pattern looks manipulative, or an idea rests on non-public information	Order-level controls; provenance checks on every input	Manipulative patterns blocked at order level; non-public information kept out

Unknown unknowns: the biggest risk. The failure that hurts most is the one nobody anticipated: a market event, a data fault, a model behaviour not seen before. We can't list it, so the system is built to fail safe without knowing what it is:

- hard limits on position, exposure and order size that no model can change;
- drawdown stops, per strategy and firm-wide, that cut risk automatically;
- kill switches, automatic and manual, that freeze or flatten the book;
- no trade on missing, stale or inconsistent data: the default is to do nothing;
- small size, own capital only, until the record earns more;
- every decision logged and replayable, so each new failure becomes a new rule.

Market risk: the edge goes away.

RISK	WHAT FAILURE LOOKS LIKE	HOW WE'D DETECT IT	MITIGATION, OR WHAT WE DO
Strategy decay	A strategy that worked stops working; markets never finish (Section 2)	Each strategy's forward net results; per-strategy drawdown stops	Per-strategy loss limits; the loop's job is to keep producing new independent strategies as old ones decay
Crowding	Our ideas are the ones every frontier-model agent finds, and the trade is gone or reverses	Month 12 crowding check against what a generic frontier-model agent finds	Crowding is a risk the decision system measures before every trade; many small, independent strategies limit what one crowded exit can cost; the edge is the rate of learning, not one idea
Capacity	Costs and impact eat the edge at the size we want to trade	Capacity estimated per strategy by Month 12	Size never exceeds estimated capacity
Regime shifts	A model calibrated in one regime is confidently wrong in the next	Calibration checked every version; walk-forward scoring; drawdown stops	Re-scoring every version on the newest outcomes; hard limits hold regardless of regime
Tail events	A sudden move larger than history suggests; the next crisis is always outside the model	Tail exposure measured before every trade	Hard limits, drawdown stops and kill switches; small size until the record earns more. Models and environments are never relied on to foresee tails

The live record of about 2.5 months (Section 7) covers one stretch of market conditions. It says nothing yet about how the strategies behave in a different regime or a tail event.

Business risk: the firm fails around a working loop.

RISK	WHAT FAILURE LOOKS LIKE	HOW WE'D DETECT IT	MITIGATION, OR WHAT WE DO
Key person	I am a solo founder; if I leave, QSI Labs ends	Not a test: a fact today	The first hire is a founding engineer who owns the pipeline. I am open to a co-founder
Model-provider dependency	The agents run on ChatGPT and Claude models today; a provider changes pricing, access, terms or model behaviour	Agent results tracked on our scored forecasts and through the decision system's checks	The decision system contains no AI, so no model change can loosen a limit; the open-weight post-training path is the step that removes the dependency for the models we train
Competition	Standard Signal and Prodigy started first, are funded and are live [8][9][10]; Standard Signal's thesis is close to ours; incumbents can out-resource us	Their public records and claims; our own record against the pre-registered bars	Compete on how the log is built and proven (Section 10), not on time
Data licences	Some data licences forbid model training, or cover individual use only	Licence review before any dataset enters training	Every dataset's business licence confirmed in writing before its data enters training
Venues	A venue restricts access, or the book outgrows retail-grade venues	Venue terms reviewed as the book grows	Trade only where QSI Labs is permitted to; move to more professional venues as the book grows
Structure and regulation	If holding securities becomes the firm's main business, Investment Company Act questions follow [20]; managing anyone else's money brings fund and adviser rules [21]	Counsel's review before any change in what capital is traded	The firm's own balance sheet trades only under a structure counsel has confirmed; outside capital only with an audited record and the right fund and adviser structure (Section 3)

Scaling risk: net alpha shrinks as capital grows.

RISK	WHAT FAILURE LOOKS LIKE	HOW WE'D DETECT IT	MITIGATION, OR WHAT WE DO
Net alpha shrinks with size	A strategy that earns its edge on about $10k loses it at a larger size, because impact, fees and financing rise with size; a short, small record says nothing about this	Realised costs split out of every trade and compared with modelled costs as size steps up; capacity estimates per strategy	Grow by adding independent strategies, not by enlarging one; size each to its estimated capacity; step up only on an audited, deflated record of at least 12 months (Section 12)

> None of these risks is solved by saying it out loud. Each is either decided by a test with a kill criterion, or contained by limits that no model can change.

§ 15

References

All accessed between 25 and 28 October 2026; [22]–[28] accessed 30 October 2026.

[1] OpenAI, "Balyasny Asset Management" (customer case study, 6 March 2026). https://openai.com/index/balyasny-asset-management/

[2] Citadel, "NXT offsite's AI hackathon", 24 September 2026. https://www.citadel.com/careers/career-perspectives/nxt-offsites-ai-hackathon/

[3] Millennium, "Millennium and Anthropic to co-develop AI-powered digital risk analyst", August 2026, https://www.mlp.com/life-at-millennium/millennium-and-anthropic-to-co-develop-ai-powered-digital-risk-analyst/ ; date from Bloomberg, "Millennium Partners With Anthropic to Develop AI Risk Analyst", 6 August 2026, https://www.bloomberg.com/news/articles/2026-08-06/millennium-partners-with-anthropic-to-develop-ai-risk-analyst

[4] Coatue, CTEK Q2 2026 presentation. https://www.coatue.com/ctek/documents/Coatue_CTEK_Main_Q2-2026.pdf

[5] HFR, "Hedge Fund Industry Asset Growth Shatters Records", 23 July 2026. https://www.hfr.com/media/market-commentary/hedge-fund-industry-asset-growth-shatters-records/

[6] Crisil Coalition Greenwich, "Revenues Surging For Nonbank Liquidity Providers", 3 June 2026. https://www.greenwich.com/press-release/revenues-surging-nonbank-liquidity-providers

[7] Hedgeweek, "Non-bank trading firms surge as Jane Street and Citadel Securities drive record $114bn revenue pool", 3 June 2026 (Crisil Coalition Greenwich data). https://hedgeweek.com/news/non-bank-trading-firms-surge-as-jane-street-and-citadel-securities-drive-record-114bn-revenue-pool

[8] Y Combinator, Standard Signal. https://www.ycombinator.com/companies/standard-signal

[9] Y Combinator, Prodigy Research. https://www.ycombinator.com/companies/prodigy-research

[10] 007 Venture Partners, Prodigy Research. https://www.007venturepartners.com/signals/prodigy-research-the-frontier-ai-trading-research-lab

[11] Y Combinator, KelAI. https://www.ycombinator.com/companies/kelai

[12] Bailey, D. H. and López de Prado, M. (2014), "The Deflated Sharpe Ratio", Journal of Portfolio Management 40(5). https://www.davidhbailey.com/dhbpapers/deflated-sharpe.pdf

[13] Bailey, D. H., Borwein, J., López de Prado, M. and Zhu, Q. J. (2017), "The Probability of Backtest Overfitting", Journal of Computational Finance 20(4). https://www.davidhbailey.com/dhbpapers/backtest-prob.pdf

[14] OpenAI, "Harvey" (customer story). https://openai.com/index/harvey/

[15] Harvey, "Expanding Harvey's Model Offerings", 13 May 2025. https://www.harvey.ai/blog/expanding-harveys-model-offerings

[16] Cursor, "Composer 2 technical report". https://cursor.com/blog/composer-2-technical-report

[17] Computerworld, Checkr interview, 30 September 2024. https://www.computerworld.com/article/3541362/checkr-ditches-gpt-4-for-a-smaller-genai-model-streamlines-background-checks.html

[18] Li, X. et al. (2023), "Are ChatGPT and GPT-4 General-Purpose Solvers for Financial Text Analytics?", EMNLP 2023. https://arxiv.org/abs/2305.05862

[19] Borah, A., ..., Singh, S., ..., Das, A. (2025), "Alignment Quality Index (AQI): Beyond Refusals", EMNLP 2025, pp. 2888–2947. https://aclanthology.org/2025.emnlp-main.145/

[20] Investment Company Act §3(a)(1)(C). https://www.law.cornell.edu/uscode/text/15/80a-3

[21] NFA, "Exemptions Available to CPOs". https://www.nfa.futures.org/members/cpo/cpo-exemptions.html

[22] McLean, R. D. and Pontiff, J. (2016), "Does Academic Research Destroy Stock Return Predictability?", Journal of Finance 71(1). https://doi.org/10.1111/jofi.12365

[23] Glasserman, P. and Lin, C. (2024), "Assessing Look-Ahead Bias in Stock Return Predictions Generated by GPT Sentiment Analysis", Journal of Financial Data Science 6(1). https://arxiv.org/abs/2309.17322

[24] Mandelbrot, B. (1963), "The Variation of Certain Speculative Prices", Journal of Business 36(4), 394–419. https://doi.org/10.1086/294632

[25] Cont, R. (2001), "Empirical properties of asset returns: stylized facts and statistical issues", Quantitative Finance 1, 223–236. https://doi.org/10.1080/713665670

[26] Millennium Management, firm overview. https://www.mlp.com/

[27] Flow Traders, Annual Report 2025, risk management section. https://flowtraders.com/media/2s2pdl01/flow-traders-annual-report-2025.pdf

[28] Point72 Europe (London) LLP, "MIFIDPRU 8 Disclosures" (period 1 January to 31 December 2024; published 2025). https://point72.com/wp-content/uploads/2025/08/Point72-Europe-London-LLP-MIFIDPRU-8-Public-Disclosure-2025.pdf