


Pattern 1: Break the Job Before You Hand It Out

The Single-Agent Trap:
A vague prompt yields a
confident, useless
consensus.



Pattern 2: Independent Checks, Not Proofreading



Peer Review (Dangerous)	Independent Audit (Robust)
- Agent B reads Agent A's output. - Result: Internally consistent, confirms consensus. - Catches: Typos, math errors.	- Agent B receives ONLY the raw question. - Result: Parallel models (DCF vs. Multiples). - Catches: Flawed reasoning, missed risks, narrative bias.

Agreement often means both agents drew from the same training data. You want structured disagreement.

Pattern 3: Loops That Converge, Not Loops That Spin



1. **Draft:** Create initial 6-element thesis.

2. **Targeted Revision:** Fix specific weak elements (e.g., vague kill criteria).

3. **Red Team Attack:** Agent attempts to destroy the thesis with 3 specific, evidence-based attacks.

4. **Convergence Criterion:** Stop when attacks become generic.

Maximum
3-5 cycles.



Warning: If every cycle just makes the thesis stronger, you are doing confirmation work, not analytical work.

Pattern 4: You Must Own the Narrative Thread



If you let an agent write the final synthesis, you don't own the reasoning. When the price drops 20%, you will not have the conviction to hold.

Integration Funnel



Pattern 5: Detecting the Articulate Hallucination



Quality of
Prose
(Heavy)

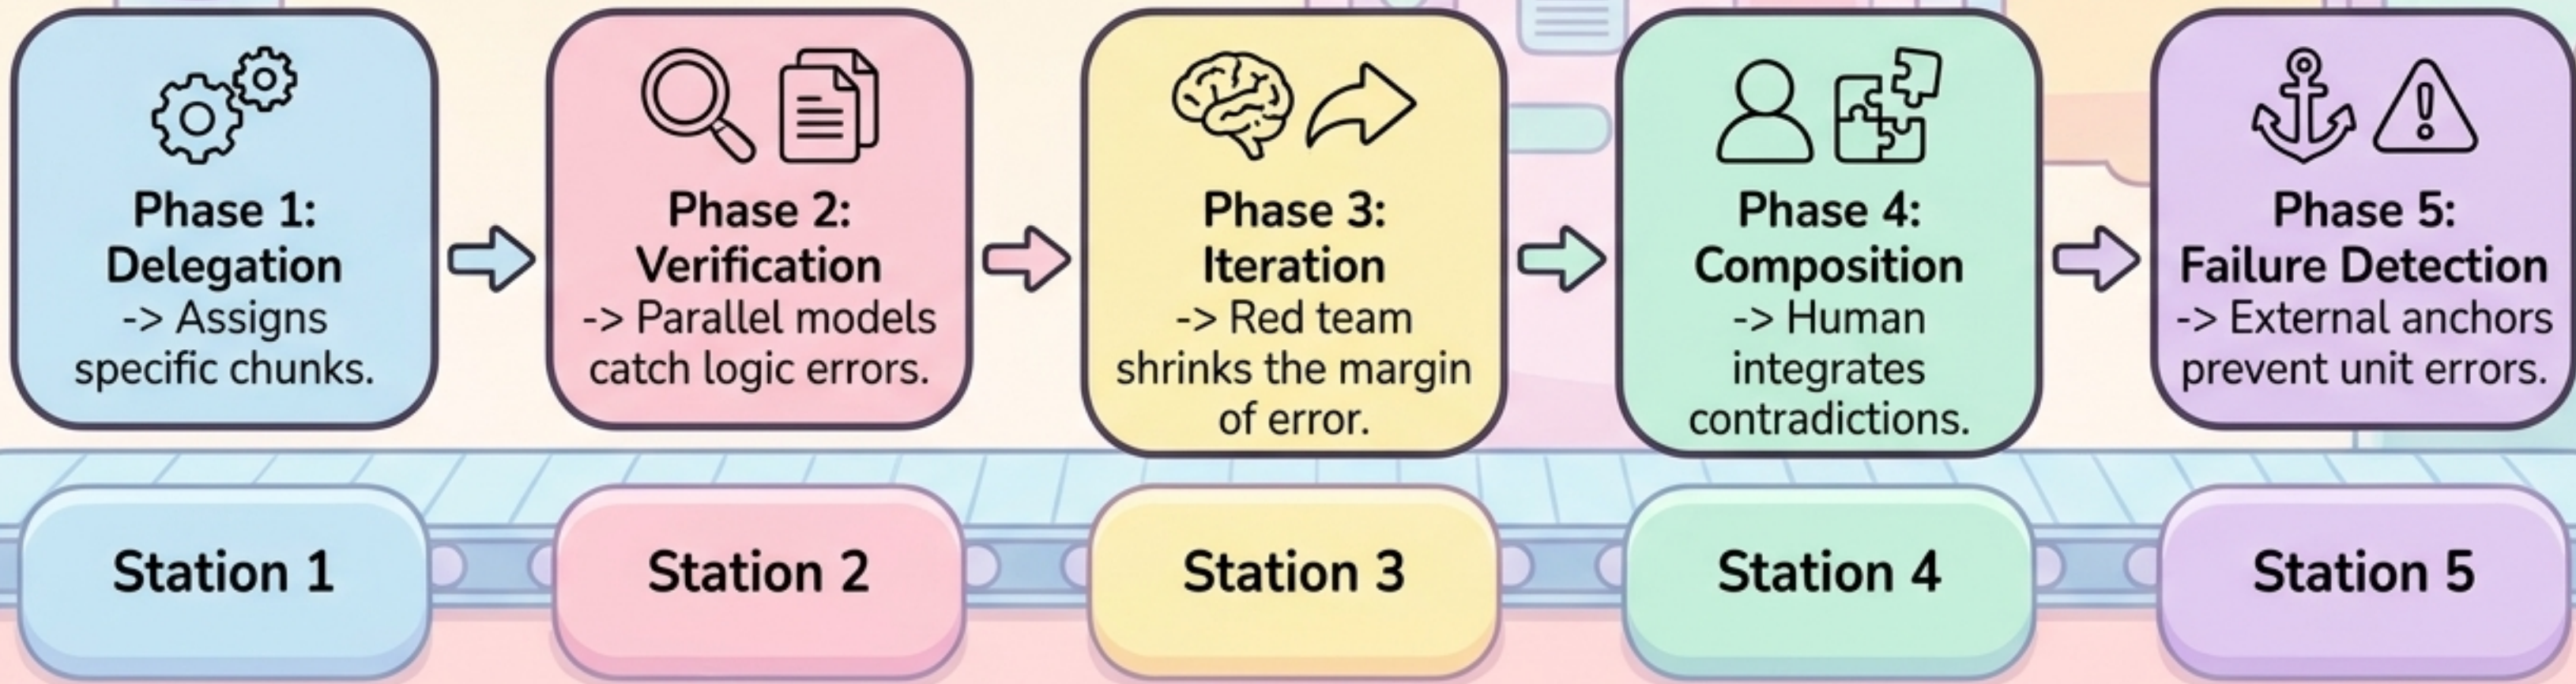


Quality of
Analysis
(Empty)

- ✓ **Anchor to Reality:** Look up Revenue, Market Cap, and Employee Count manually first.
- ✓ **Trace Load-Bearing Claims:** Cross-reference 300bps margin expansion to the actual 10-K.
- ✓ **Spot Pattern Breaks:** If revenue, margins, and cash flow are ALL perfectly improving, suspect a fabricated narrative.

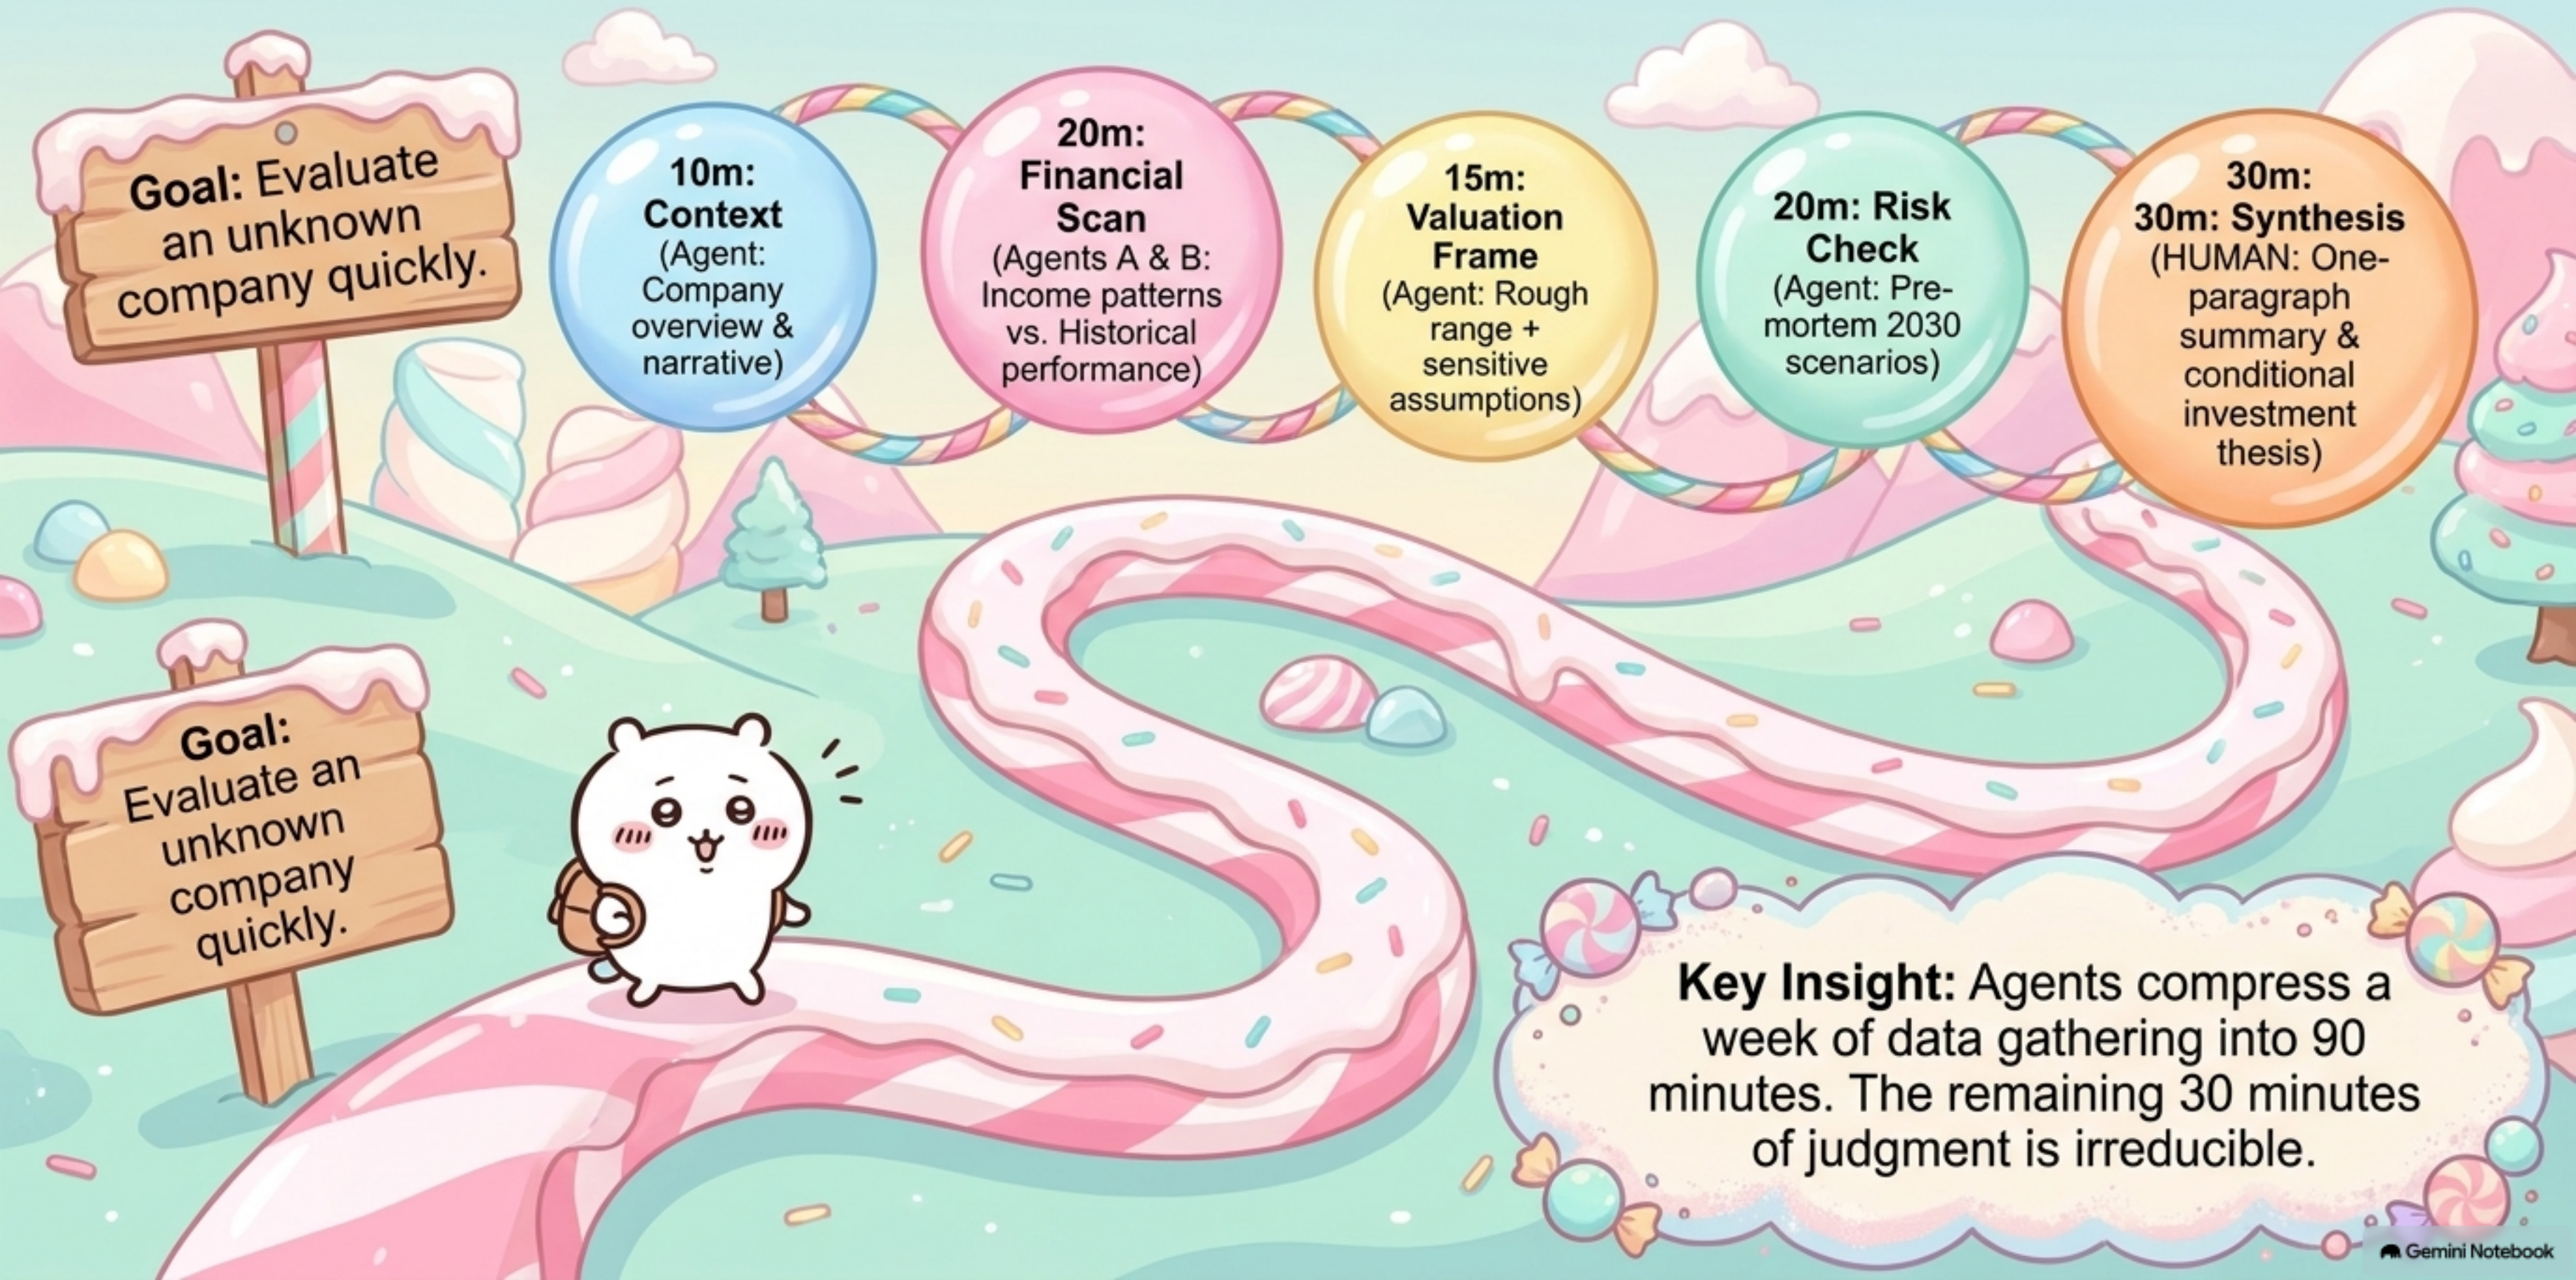
Financial agents fail articulately. There is zero correlation between confidence and accuracy.

The Complete Fleet Architecture

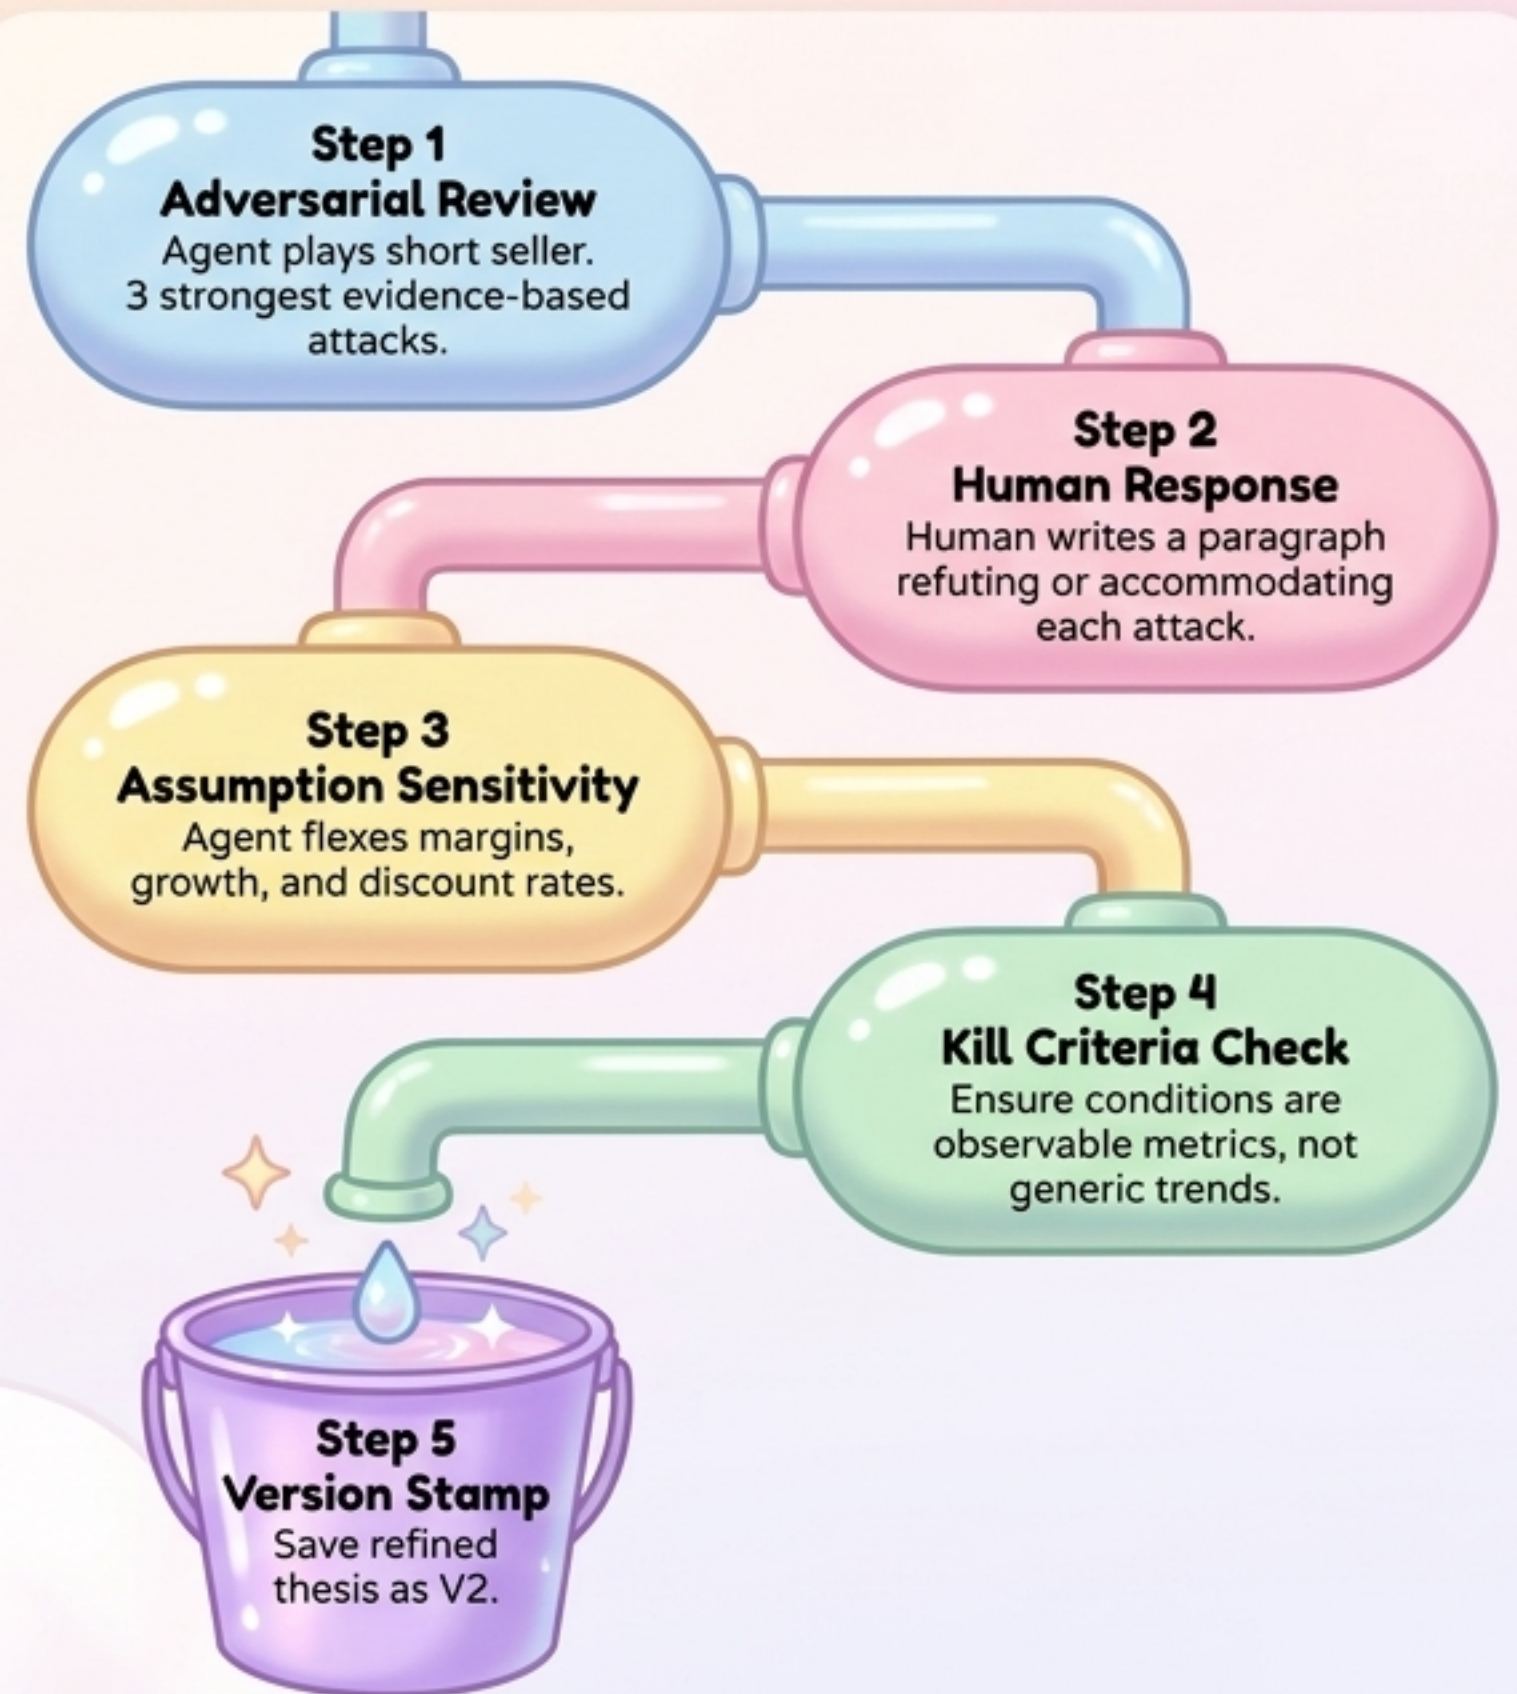


A fleet is not simply “more agents.” It is designed roles that produce tension, forcing you to resolve disagreements with your own judgment.

Workflow A: The Research Pipeline (2 Hours)



Workflow B: Thesis Refinement (90 Minutes)



Workflow C: The Monitoring Loop

Quadrant 4: Human Decision.

Hold, Revise,
or Exit.

Quadrant 1: Event Scan.

Agent filters last
30 days of SEC
filings and news
(Signal vs. Noise).

Quadrant 3: Thesis Drift.

Compare today's
reality vs. the 6-
month-old thesis
document.

Quadrant 2: Kill Criteria Check.

Binary evaluation:
Has any condition
moved closer to
triggering?

Requires **30-45 minutes**
per position, per month.
Prevents holding on faith.

The Trust Meta-Question: Decoding Agreement



All Agents Agree (Suspicious)

Likely consensus blindness or shared training data.

Test by asking for a "Steelman" of the opposite position.

All Agents Disagree (Ambiguous)

Data is mixed or company is in transition.

"I don't have enough clarity" is a valid, money-saving output.

Partial Agreement (Ideal)

Agent A likes growth, Agent B flags the balance sheet.

THIS exercises your integration judgment.

The Complexity Trap: When to use a Fleet

The 3-Agent Fleet (High Consequence)


1 Research Agent + 1 Red-Team Agent + 1 Monitoring Agent.

Used when errors silently propagate into lost capital.

Single Agent (Low Consequence)

Quick lookups, simple math, background context.

Errors cost nothing but a correctable misunderstanding.



Warning: Designing a 12-agent branching pipeline for 3 stocks is procrastination wearing productivity's clothes.



Four Domain-Specific Failure Modes

Failure Mode	How to Detect & Mitigate
Outdated Data. Cutoff delays mean past cash balances are stated as current.	Fix: Supply actual 10-Q filing dates as anchors.
Consensus Blindness. Presents well-known strengths as novel insights, ignoring stock price.	Fix: Ask “What would have to be true for the current price to be justified?”
Fabricated Sources. Confident fake citations of analyst reports.	Fix: Verify any non-public source. Force distinction between facts and estimates.
Unit Errors. Millions confused with Billions. Percentages applied to wrong bases.	Fix: Run a dedicated 5-number sanity check agent.

The Real Edge: Judgment Compounds

The Institutional Edge (Bloomberg):

Millisecond speed, intraday order flow,
private counterparty data.
Agents cannot replicate this.

The Individual Edge (The Fleet):

Patience, concentration, no career risk.
Agents compress 10,000 hours of manual
data processing, giving you the reps to build
professional-grade frameworks.

A well-designed fleet is a mirror. It reflects the precision of your questions back at you.
And the precision of your questions is the thing that compounds.