Should I buy? is the wrong question.

Asking an AI "Should I buy X?" collapses research, challenge, learning, and execution into one step.

You get a generic consensus answer from whichever "specialist" steps forward first.

You do not earn the right to trust the recommendation.

Your AI is not an oracle. It's four distinct roles: Analyst, Red Team, Tutor, and Executor. The way you frame your request determines who answers.

The Red Team: Attack Before the Market Does.

Don't ask for risks (which activates the generic Analyst).

Tell the AI: "I believe X. You are a short seller. Build the strongest case that my thesis will fail."

It targets your specific assumptions, not generic market risks.

The Iterate-Until-Silence Pattern

Present thesis → Red Team attacks → Revise thesis

Repeat until the attacks stop being substantive.

The Tutor: The Explain-Back Loop

Don't use AI as a crutch; use it as a training tool. When you hit a complex concept, switch to Tutor mode. The goal is comprehension, not just recognition.

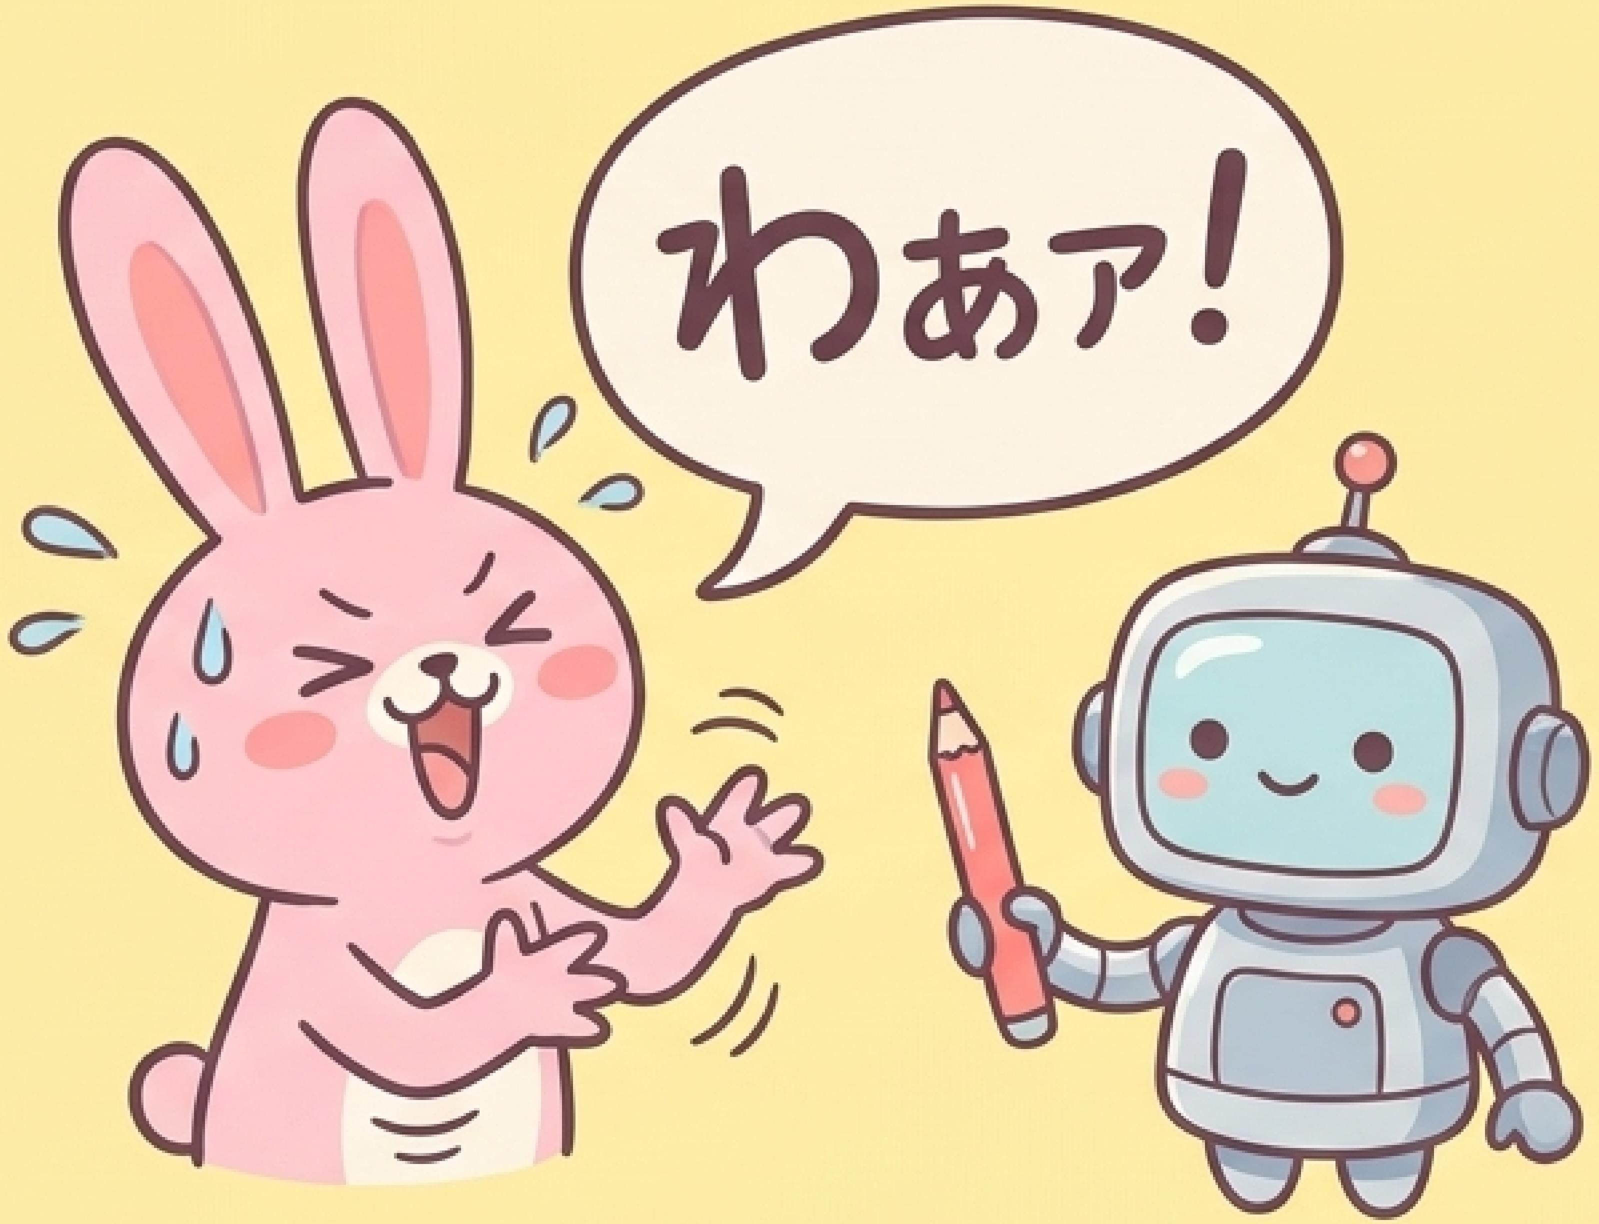

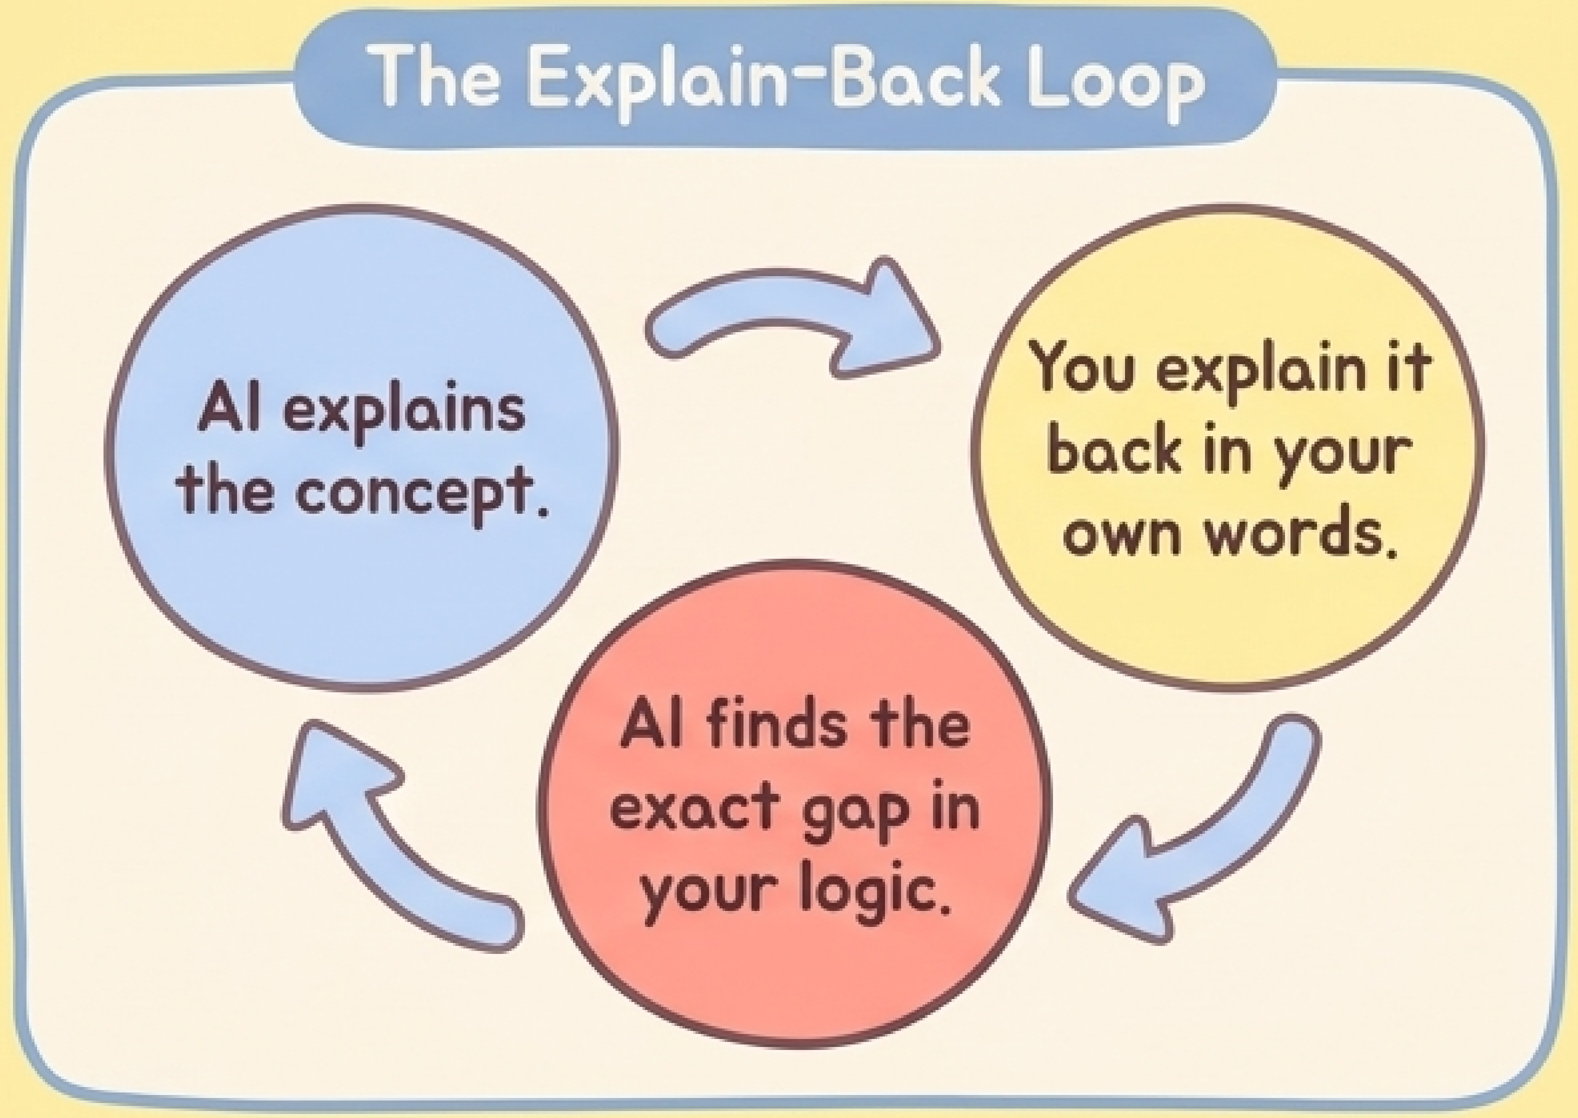

What Agents Cannot Do.

An agent can compute optimal position sizing mathematically, but it cannot simulate the visceral feeling of watching real money evaporate. Conviction under drawdown is built through experience, not analysis.

Cannot judge a management team's true character.

Cannot define your psychological relationship with money.

Cannot distinguish unprecedented paradigm shifts from historical noise.

The Executor & The Art of Doing Nothing.

Once your workflow is validated, use the **Executor** to run stable, repetitive data pulls across companies. But remember: in most market environments, the correct action is **nothing**.

Analyst: Information aggregation

Red Team: Adversarial challenge

Tutor: Calibrated learning

Executor: Stable workflows

The agent's job is to help you think clearly. The discipline to close the laptop is yours.