Warning: AI does not fix systemic errors in venture investment decisions.

By Ravi Chalaka, Founder and CEO, Knoh.AI

I keep falling for the same trick.

I ask an AI something. It pauses for a beat. Then it comes back with an answer so smooth, so composed, that some part of my brain goes "oh, it actually thought about that."

It didn't.

And if you're a VC using LLMs to move faster through decks, meeting notes, and data rooms, you could be introducing a real risk to how you make decisions. Here's why.

The Brutal Funnel Problem

Venture capital investments have a brutal funnel problem to begin with:

Only a small fraction of the opportunities a firm sees ever convert into a deal pipeline, and a small number of those convert into an investment, and an even smaller number convert to an expected return.

A third of them arrive through warm referrals rather than independent sourcing, which means much of the funnel may be shaped by who you already know.

Layer an unstructured cognitive bias on top of that, and the picture gets worse.

Researchers have cataloged over a dozen distinct biases that quietly steer VC decisions. For example:

- **Confirmation bias,** where you notice the data that supports your thesis and skip past the data that doesn't
- **Recency bias,** where a hot quarter overrides the long-term picture
- **Herding,** where firms pile into a deal mainly because other firms already have
- And many more. We'll cover an entire edition on this next.

SoftBank's continued funding of WeWork even as the business model came under serious doubt is a well-known example of confirmation bias playing out at scale.

None of this is news to anyone who's spent time in a partner meeting. What is new is that AI tools are now sitting inside that same process, and most of them don't fix the bias problem.

They accelerate it.

Why AI Makes This Worse Before It Makes It Better

Your Brain vs. Artificial Brain

Your brain generates questions when something doesn't add up. It runs a constant prediction against reality, and the moment reality breaks the pattern, it flags a "why" for you to chase down.

That’s how attention and memory circuits actually behave when they detect a mismatch.

LLMs don't do this. They complete questions. They don't generate them.

Ask an LLM "why is this a great deal?" and it will build you a confident case. Ask it "why is this a bad deal?" using the same data, and it will build an equally confident case the other way.

The model hasn't changed its mind because it never had one.

It's just completing whatever frame you handed it.

That's the dangerous part, specifically for a VC. If confirmation bias means you already lean toward liking a deal, an LLM won't catch that. It will hand you back a well-written version of your own bias, dressed up as analysis.

The Best Thinkers Interrogate the Problem Before It's Even a Problem

A VC's process runs through four stages: opportunity sourcing, thesis fit analysis, due diligence, and the investment decision.

Thesis fit is a filter, not an interrogation. Its job is to pull the handful of deals that match a firm's stated strategy out of a much noisier opportunity pipeline. It answers "Does this belong in our funnel?" not "Should we fund this?"

The real testing happens in the due diligence pipeline. That is where a team pressure-tests everything against evidence, across every stage, market, team, product, financials, legal, competitive, until the assumptions baked into the thesis have actually been tested, not just accepted once and carried forward.

To de-bias decision-making, VC firms have to shift from intuitive gut-picking to structured evaluation systems and evidence-based judgment. Managing capital relies on minimizing systematic errors across the entire investment workflow, not just catching them once at the start.

Why We Built “Knoh AI” Around This Idea

Earlier, I said AI doesn't generate questions, it only completes the one you hand it. That's true of a generic LLM sitting in a chat window. It's not true of AI agents.

An AI agent carries a set of structured steps and questions to probe at every stage of the diligence process that an experienced VC partner thought to ask. That's the actual gap we built Knoh.AI to close.

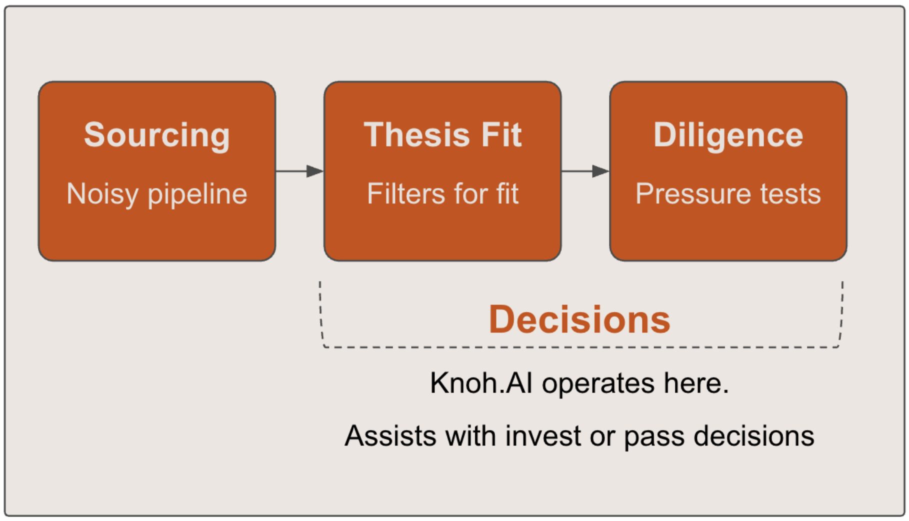

Knoh.AI isn't a chat window you throw a prompt into and get a paragraph back.

It automatically syncs the relevant data into the deal-specific context library. It's a set of specialized agents, each one focused on a specific area of the decision, market, team, product, financials, legal and competitive landscape, that work through the material a VC firm has in front of them. Pitch decks, data rooms, call notes, financials, contracts, whatever exists.

Each agent's job isn't just to summarize. It's to produce analysis and insights for its area: an executive summary, the implications, the risks, gaps and the next steps worth taking.

Prior to the structured due diligence workflow, Knoh also scores every inbound opportunity against the firm's thesis and strategy for each active fund, based on the firm's criteria.

To de-bias your VC firm's decision-making, you must shift from intuitive "gut" picking to structured evaluation systems. Managing capital relies on minimizing systematic errors during the investment workflow.

Those pieces get stitched together across the stages a firm actually works through, intake, thesis fit, and due diligence, so the firm isn't starting from a blank page at each stage. They're starting from structured, specific work that's already been done and can be checked, challenged, and built on.

So it's less "ask Knoh a question, get an answer" and more "hand Knoh the material, get back the analysis a sharp associate would have produced, across every area that actually matters to the investment committee decision, consistently, every time."

That consistency is the actual point. Not novelty. Not a smarter chatbot.

The Real Frontier Isn't Better Answers

I don't think the next real leap in AI comes from squeezing out slightly better answers from an LLM to the same old questions.

And I want to be careful here, because there are plenty of investors who already ask great questions and make great decisions. This isn't about them being wrong. It's that even the best process is inconsistent when it runs on human bandwidth. A great partner asks the sharp thesis question on Monday and misses it on Friday, because they're tired, or the deal came in through a warm intro, or there wasn't time before the next meeting.

What we're trying to do at Knoh.AI is take the best of what investors already do, the habit of interrogating the frame before accepting it, and make it something a firm can rely on every single time, not just on their best days.

Because the breakthroughs, in investing or anywhere else, rarely come from answering old questions more precisely. They come from consistently asking the right question without bias before the decision is made, rather than after.

That's the process we’ve built into Knoh to help VC partners make the best decisions.