



THE SIGNAL BRIEF / DECISION QUALITY / JULY 13, 2026

Subject: AI had an answer. Did it have enough?

Preheader: A recommendation can be reasonable in general and still arrive before the question that matters.

AI Wasn't Wrong. It Answered Too Soon.

A mortgage choice and a 401(k) question surfaced the same gap: an answer can sound settled before the system has learned which facts could change it.



Illustration: a hand offers the missing paper piece before the recommendation proceeds.

Don Long · Founder, SummitMark · Helping people and organizations become capable with AI

The recommendation arrived before the qualifying question did.

I asked an AI system a familiar financial question: pay down a low-rate mortgage, or invest the money instead?

The answer sounded considered. It compared a fixed borrowing cost with possible long-term returns, laid out caveats, and leaned toward investing. It was easy to start weighing that recommendation before noticing what the system did not know: whether the money was spare cash or the household's emergency reserve.



That is the moment worth examining—not when AI says something obviously wrong, but when a plausible answer sounds finished before it is.

The detail that changed the 401(k) answer

I also asked about choosing between a 401(k) loan and a hardship withdrawal. The system initially leaned toward a loan, noting that a compliant plan loan is generally repaid to the account and is not taxed as a distribution while the rules are followed. Later in the response, it named job stability as a risk, but it had not asked about mine. (IRS guidance on plan loans; hardship distributions.)

So I added one detail: the person might leave the employer soon.

That matters. A plan may require repayment when someone leaves; if the loan is not repaid, the outstanding amount can become a taxable distribution, with additional tax potentially applying depending on age and exceptions. The specific plan terms matter. With that detail, the answer became more cautious and conditional, and pointed more strongly toward checking the plan's rules before deciding.

It did not turn into “never take a loan.” The important change was that the recommendation had to account for a fact that could alter the consequences.

Accuracy is not enough

Neither a loan nor a hardship withdrawal is the right choice in general. The same is true of paying down a mortgage versus investing. Personal circumstances and plan rules shape what is sensible.

So there are two questions to ask of an AI answer:

1. Is the information accurate?
2. Has the system gathered enough about this situation to recommend anything yet?

Sometimes the more useful response is a question back. Once a recommendation is already on the screen, we tend to evaluate it instead of asking what the system still does not know.

The product decision comes earlier

This matters beyond personal finance. A lending assistant, benefits tool, insurance comparison, or financial-wellness product may use sound general logic and still recommend a course before surfacing the detail that changes what it means for this person.

A disclaimer does not fill in missing context. A more articulate explanation does not either. The design choice happens earlier: should the product answer now, ask another question, or involve a person?

That is a product judgment, not just a writing choice. Teams need to identify which missing facts could change an outcome, and decide how the system should respond when those facts are unknown. Sometimes a provisional answer is useful. Sometimes the next step is to ask. In a consequential case, it may be to direct someone to the plan administrator or a qualified professional rather than imply that general information settles an individual decision.



One habit

The next time an AI answer feels complete, pause and ask:

What doesn't it know yet?

Then name the missing fact most likely to change the answer. "Is this cash also my emergency fund?" "Could I leave this employer before the loan is repaid?" A concrete question is more useful than a general request to "be careful."

Mostly right and right for you are not the same thing. The difference may be one question asked at the right time.

Where has one added detail changed an AI recommendation you received? Reply with the detail—and what it changed.

M.R.

Source and scope: Adapted from Don Long's first-person essay (July 13, 2026). The interactions and changed response are author-reported; prompts, outputs, and plan documents were not supplied for independent review. IRS links above support the general plan-rule statements. Not individualized financial, tax, or investment advice. Send date and author approval are pending.