AI / KNOWLEDGE & JUDGEMENT

When summaries are free, *reality gets expensive*

Better work is one thing. Becoming better at the work is another.

Use AI to expand what people can attempt, while protecting the work that teaches them how to judge the result.

28 October 2026 · About 12 min read



The argument in 30 seconds

1. **AI can improve the quality and range of our work.** It can also weaken the practice through which we learn to judge it.
2. **Good research combines access with testing.** Source documents and local conversations matter when we use them to challenge attractive explanations.
3. **Firms should reinvest some of the time AI saves.** More feedback and independent judgement matter as well as faster, polished output.

Executive summary

AI makes it easier to produce convincing work and, used well, to explore ideas we might never otherwise attempt. But fluency can disguise a missing apprenticeship: summarising, explaining and defending a view are ways analysts learn, not merely ways they report conclusions. The investment edge increasingly comes from combining the archive with evidence that has not yet been written down, then testing both. Sri Lanka's recovery and Senegal's hydrocarbon-led growth show why aggregate numbers and local conditions must be read together. First-hand accounts can mislead too; they are prompts for investigation, not substitutes for verification. My preference is to use AI for exploration, synthesis and challenge while preserving independent first attempts, conversations with colleagues and explicit tests of what would change a conclusion. The goal is not just better output. It is a better analyst.

The promise and the problem

A strange thing is happening to knowledge work. It is getting easier to sound informed without becoming much better informed. I use AI constantly, and some of the best things it gives me are creative. It helps me develop half-formed ideas, connect thoughts across economics, markets, history and psychology, and turn them into research that is more interesting and more fun to produce. At times, it feels like a small personal renaissance. More ideas seem worth exploring because the distance between having a thought and doing something with it has shortened. It feels like a thinking partner that expands what I can attempt.

That is also why I take its risks seriously. A tool this useful can change how we think, including which parts of thinking we stop practising. A clear summary used to suggest that someone had spent time with the material. Now it may just mean they prompted well. The summary can still be valuable. It simply tells you less about how much its author understands.

There is a second problem, and I think it matters more. Summarising was not only how you demonstrated understanding. Sometimes it was how you developed it. It is the same reason teaching can be such a good way to learn. You think you understand something until you have to explain it clearly to someone else. Then the gaps appear. You cannot quite connect one step to the next, or answer what seemed like an obvious question. Working through that difficulty is part of the learning.

The work behind the note

That is what makes me uneasy about being a junior analyst now. You can produce work that looks experienced before you have had enough experience to judge it. Consider a budget release on a trading desk. The documents arrive, prices are moving, and someone wants to know what matters.

You work through the revenue assumptions, spending plans and financing tables. Is the deficit falling because policy has improved, because the economy is larger, or because payments have been pushed into the future? Does the borrowing plan imply more bond supply than the desk expected? You bounce a view off a colleague. They challenge a definition or an assumption. A trader asks what is already priced. You revise the argument and explain which part you are still unsure about. The note is the visible output. Along the way, you are learning how a budget works, how a forecast fails and how to turn an observation into a decision.

Figure 1

The same note can contain a different apprenticeship

Follow the work that happens before a finished answer.

01 Read the budget	02 Compare the assumptions	03 Form a first view
04 Ask a colleague	05 Check the cash flows	06 Challenge the explanation
07 Revise the argument	08 Ask what is already priced	09 **Write the note**

With practice: form a first view, test it, then revise. AI can help at each step.

Answer first: a finished interpretation can bypass those opportunities. A better-looking note does not establish better unaided judgement.

Conceptual task map adapted from the original essay. Steps are not measured time or learning gains. Both paths can produce a note; the question is what the analyst learns along the way.

Some of that process is inefficient. Some of it is the apprenticeship. AI can help with almost every step. But if you receive the interpretation before trying to form one, you have skipped an opportunity to practise. Reading a convincing answer does not necessarily teach you how to reach it, or recognise when it is wrong. The result can be a junior analyst who writes better notes without becoming proportionately better at analysis.

That outcome is not inevitable. Research has found both weaker unaided learning after students used an answer-giving AI and evidence of durable learning among workers using AI assistance. The settings differ, but they point to an important distinction: a tool can help people practise, or help them bypass practice. The responsibility extends beyond the junior. If a desk rewards speed and polish while removing opportunities to work through problems with colleagues, independent judgment may have less room to develop.[1][2]

> ### What the learning evidence can establish
>
> The mathematics experiment and customer-support study concern different settings and outcomes. They do not establish a universal effect on investment analysts. The second study is a substantive counterargument: AI assistance can also support learning. [1] [2]

Outside the archive

A firm can improve this quarter's output while underinvesting in the people it will depend on five years from now. "Humans will provide the judgment" is an incomplete answer if we automate the work through which humans acquire it. That judgment matters partly because the information needed for a good decision is not always sitting in a document.

Some of it is held by the contractor waiting for payment, the importer trying to obtain dollars, or the official deciding which expenditure can be postponed. It may emerge in a local meeting, an expert webinar or a conversation that only becomes useful after the third follow-up question. AI can help you prepare for those conversations and make sense of what you learn. But a model working from the material available to it cannot know an unreported payment delay merely because you ask the question well. A plausible account of what usually happens is not evidence of what is happening now.

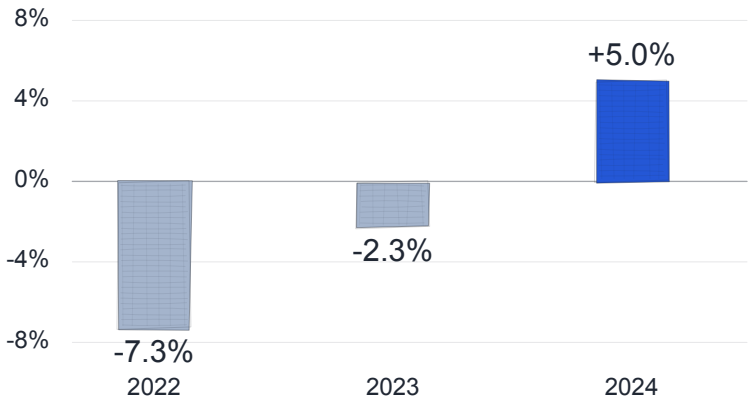
This is where relationships and access matter. Knowing whom to ask, earning a candid answer and recognising what deserves another question are part of research. A contact list alone is not enough. The relationship has to produce information you can evaluate.

Two different tests

Take Sri Lanka. The domestic economy was rebounding while bond negotiations were still going on. Real GDP grew by 5 per cent in 2024. The pace and durability of that recovery mattered for the securities emerging from restructuring. Sri Lanka's macro-linked bonds tied future payment adjustments to economic outcomes, including average dollar GDP over 2025–27 and a real-growth condition. A stronger recovery could affect both the amounts bondholders would receive and the perceived risk of receiving them. [3][4]

Figure 2

A recovery is not a bond return



Real GDP growth · annual observations

Read the contract as well as the recovery

What does this mean for a bondholder? A recovery can affect both repayment capacity and payments on macro-linked bonds. The contract, currency path and price still determine the investment return. GDP growth is not a bond return. Contract discussion [4].

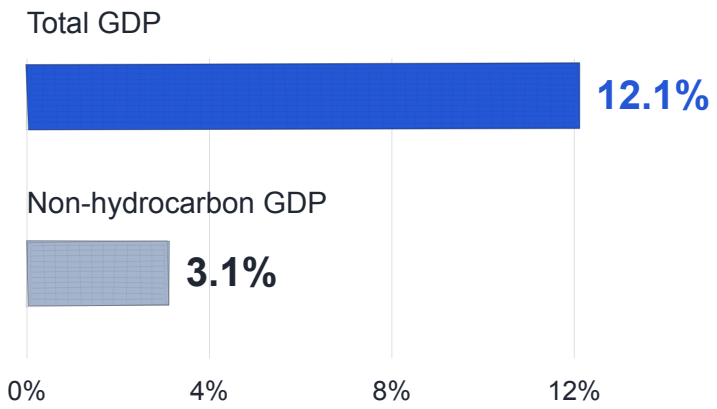
Sri Lanka · annual real GDP growth, 2022–24: −7.3%, −2.3%, +5.0%. Historical observations, not forecasts. Source: Sri Lanka DCS [3].

The question was how the growth outlook translated into the different payout scenarios, and how much was already reflected in price. Local knowledge could help assess the recovery. The bond contract determined how that view translated into value. Senegal illustrates the opposite tension. Signs of slowing capital expenditure and the authorities' efforts to restrain spending called for a closer look beneath the oil-and-gas growth story.

The IMF reported 12.1 per cent headline growth in the first quarter of 2025, but only 3.1 per cent outside hydrocarbons, with payment arrears weighing on construction. A new oil project can lift GDP sharply while domestic businesses remain under pressure. As production levels off, the initial boost to growth fades; government revenues follow their own path, depending on prices, costs and the state's share.[5]

Figure 3

Headline growth can leave domestic businesses behind



Ask who is getting paid

Total real GDP grew 12.1%; non-hydrocarbon GDP grew 3.1%. These are two growth rates, not additive sector contributions.

Senegal · Q1 2025, year-on-year. Source: IMF, 26 August 2025 [5]. The IMF linked construction weakness to payment arrears. Hydrocarbon growth did not establish a broad domestic recovery.

To understand the consequences, you need to know which contractors are being paid, which projects are slowing and where spending restraint is actually falling. Published data can establish part of the picture. Local contacts and subject experts can help establish what the aggregate figures leave out. None of this replaces the data. But it changes how you read the data.

> **Read the layers together**
>
> Published data, primary contact, contract and price each answer a different question. Each layer can challenge the others; none is sufficient alone. These are historical cases, not current trade recommendations. No bond valuation or executable price is supplied.

Keep testing the evidence

These are also examples of how human research and AI can strengthen each other. A conversation changes your working assumption. AI helps you trace its implications, compare it with other cases and identify what to check next. You return to the source with a better question. AI can increase the value of a good conversation by helping you do more with what you learn. There is no permanent human monopoly on this process. Once a conversation is recorded or fresh information becomes accessible, AI can help analyse it too. The advantage comes from access to relevant evidence, the questions that draw it out and the judgment used to assess it. It has to be renewed.

Nor is first-hand experience automatically reliable. A memorable meeting can create fake conviction as easily as a polished summary. An official has incentives. A business owner sees part of the economy. A busy restaurant may tell you very little about the country around it. Texture tells you where to investigate. It does not relieve you of the need to investigate.

The same applies to primary documents. A filing is not valuable simply because it is original. The useful detail may be a footnote nobody has connected to the financing plan, or a contradiction between what management says and what customers do. Sometimes the information is new. Sometimes the interpretation is. AI can help with both. The discipline is to keep the explanation connected to evidence that could change it.

That is what worries me about fluency. It can make borrowed opinions sound owned. It can make the difficult part feel finished before you have done it. The danger is not that the model is dumb. The danger is that it is good enough to make you stop checking.

What I would watch

The answer is to be deliberate about what we ask it to do. Use it to explore unfamiliar fields, find surprising connections and develop ideas you might otherwise abandon. Use it to explain the difficult passage. Then try explaining the argument yourself. Ask it to challenge your view, and take the strongest objection back to the source. Keep the conversations with colleagues. Keep track of the assumptions behind your conclusions and revisit them when the evidence changes. In markets, a profitable trade does not by itself prove that the reasoning was right.

Four checks to take back to the desk

1. **Try before you ask.** Write your explanation in two sentences before asking AI for one. Where does your chain of reasoning become uncertain?
2. **Find the missing observation.** What fact would distinguish your explanation from the strongest alternative? Name a document, a person or an observable outcome that could establish it.
3. **Make the tool disagree usefully.** Ask AI to identify the weakest link in your reasoning. Request evidence that would contradict it, not just a more polished version of your own case.
4. **Explain it without the answer.** Close the generated response. Can you explain the mechanism, defend the key assumption and say what would change your mind? Return to the source with the gap you found.

The useful test is whether better assisted work leaves better unaided judgement behind. As an editorial interpretation of the essay's learning argument, I would be less concerned if juniors became better at explaining, defending and revising their own conclusions after using AI. I would be more concerned if their notes improved while those abilities did not. These are proposed checks, not results established by the studies above.

The investment in judgement

For firms, some of the time saved should become more discussion, more feedback and more opportunities for juniors to make and defend a judgment. Automating tedious work is progress. Preserving the learning it once contained requires thought. The best version of this feels expansive. You can explore more ideas, bring more disciplines into the same conversation and produce work that would previously have been out of reach.

But the quality of that partnership still depends on what you bring to it: knowledge you have worked to understand, people who can tell you something new, and a willingness to discover that your attractive explanation is wrong. The output should improve. So should the person producing it. Use AI to compress the archive. Spend some of the time you save expanding it.

Source notes

This is the research essay dated 28 October 2026. Its evidence has not been refreshed. The original PDF marked its references as requiring primary-document verification; the supplied HTML added exact links and source descriptions. The historical claims and their qualifications are preserved here. Link-access checks do not amount to a fresh research clearance.

1. Bastani et al. · PNAS, 2025 ↗

   Generative AI without guardrails can harm learning: Evidence from high school mathematics. A school mathematics experiment, not a study of investment analysts. Open-access copy at PMC.

2. Brynjolfsson, Li & Raymond · QJE, 2025 ↗

   Generative AI at Work. Evidence from customer-support agents, including learning observed during AI outages. Not directly comparable to the classroom experiment.

3. Sri Lanka Department of Census and Statistics · Statistical Abstract 2025 ↗

   Reported real GDP growth: −7.3% in 2022, −2.3% in 2023 and +5.0% in 2024. Historical observations, not forecasts.

4. IMF · Working Paper 25/175, September 2025 ↗

   Sri Lanka's Sovereign Debt Restructuring: Lessons from Complex Processes. Macro-linked payments depend on average dollar GDP in 2025–27 and a cumulative real-growth condition.

5. IMF · Staff visit to Senegal, 26 August 2025 ↗

   Q1 2025 real growth of 12.1% overall and 3.1% excluding hydrocarbons; the release links construction weakness to payment arrears. These are growth rates, not sector contributions.