The Intelligence Revolution

When intelligence stops being scarce, value shifts from possession to deployment

LinkedIn and kjerside.com article | AI series, Part 1 of 5 | 29 July 2026

Series note: This is the first article in a five-part series on AI, economics and advisory work. The series moves from scarcity economics, to intelligence, to AI-native workflows, to what an AI-native advisory system looks like in practice, and finally to why generic intelligence makes domain-specific software more necessary, not less. Three companion essays provide the technical foundation beneath that arc: the AI capability stack, the distinction between models and systems, and model and provider choice.

Industrial revolutions do not merely raise output. They reprice human capabilities.

AI should be read through that lens. It is not an exception to economic history. It is the latest step in a long sequence where capital, technology and energy substitute for human effort whenever they become cheaper and more productive at a specific activity. Engines cheapened muscle. Software cheapened routine coordination. AI is now cheapening parts of cognitive production: reading, drafting, coding, searching, modelling, synthesis and, increasingly, structured reasoning.

The long-run effect is likely to be a larger economy. The transition will be more uneven. The Maddison benchmark data show world real GDP per capita rising from roughly $1,100 in 1820 to about $16,700 in 2022, in 2011 dollars. The pie gets larger. The distributional journey is another matter. Productivity is not a compensation scheme.

The uncomfortable question is not whether intelligence remains useful. It is what becomes scarce when a growing share of analytical work can be produced at near-zero marginal cost.

1. A scarcity shock to cognition

In economic terms, AI is a cost-performance shock. When capital outperforms labour in a task, labour loses value in that task even if the people doing it have not become less capable. That is not a moral judgement; it is production theory.

The magnitude of the shock is worth stating. The cost of querying a model at the level of GPT-3.5 fell from roughly twenty dollars per million tokens in late 2022 to seven cents by late 2024 - a decline of more than 280-fold in under two years, on the Stanford AI Index's numbers. The exact benchmark will age; the mechanism will not. The price of a fixed level of machine cognition has been falling by roughly an order of magnitude per year. No human input has ever repriced at that speed. Professional labour does not reprice on that curve.

What makes this wave different is the input being repriced. Previous waves industrialised muscle, movement and routine administration. This one industrialises parts of cognition. For professions that have built status and fee structures around intellectual throughput, that is not a small adjustment.

I do not approach this as a detached observer. After years in strategy and transaction work, I have seen how much professional value has been tied to structuring a problem, synthesising evidence, drafting a narrative and building confidence around a decision. AI does not make those things irrelevant. It changes which parts are scarce.

The diligence-room question is simple: which part of the value chain is being repriced, and who still captures the margin? In many advisory, legal, finance, audit, software and policy environments, the answer will not appear first as job titles disappearing. It will appear as fewer junior hires, faster first passes, smaller teams and less tolerance for routine analytical labour.

2. The labour-market impact will be uneven before it is obvious

Technological revolutions usually create more wealth over time. They also redistribute value before institutions have adapted. That is why the politics of transition so often lags the economics of transition.

AI will not affect entire occupations evenly. It will affect the production system inside them. Task-exposure estimates - the OpenAI-led "GPTs are GPTs" study, the OECD's employment analyses and Anthropic's Economic Index among them - point in the same direction: management analysts, lawyers, financial analysts, accountants and programmers sit on the higher-exposure side; electricians, plumbers and construction labourers much lower. Not because one group is more intelligent, but because AI is hitting different task bundles.

A lawyer, consultant, analyst or product manager may keep the same title while the economics of the role changes materially. First drafts arrive faster. Research gets compressed. Basic coding and modelling become easier. The scarce question becomes less "can someone produce a competent output?" and more "who knows what output should exist, how good it needs to be, and what decision it is meant to support?"

This is where many predictions become too crude. The issue is not a clean contest between humans and machines. It is redesign: firms will change what is automated, what is reviewed, what is signed off, and what must remain human because it carries judgement, trust or liability. Integration costs, data constraints, risk controls and legacy processes will make that transition slower than the headlines suggest - and once workflows are redesigned, the demand for routine throughput can fall quickly, which will make it more painful than the optimists admit.

> *The scarce question becomes less "can someone produce a competent output?" and more "who knows what output should exist, how good it needs to be, and what decision it is meant to support?"*

3. Cognitive production is not the same as operational judgement

In most analytical work there is a recurring chain: choose the objective, frame the problem, produce the analysis, test the answer, decide and act. AI is strongest in the cognitive middle. It can generate hypotheses, interrogate text, draft documents, write code, build scenarios and reason through bounded problems at speed when placed inside a well-designed process.

The shift from large language models to reasoning models expands that middle. It makes AI more deliberative: better at multi-step reasoning, debugging, planning and self-correction. But it does not erase the ends of the chain. A model can reason about a deployment or an investment thesis; it still does not naturally know the true objective, the binding constraints, the live operating context, or who should bear the downside if the answer is wrong.

That distinction matters. Markets do not pay only for answers. They pay for converting ambiguity into action. In transaction work, the value is rarely just the spreadsheet or the slide. It is the judgement behind the assumption, the confidence to recommend a path, the ability to persuade a board, and the accountability to stand behind the decision.

> *In transaction work, the value is rarely just the spreadsheet or the slide. It is the judgement behind the assumption, the confidence to recommend a path, the ability to persuade a board, and the accountability to stand behind the decision.*

As analysis gets cheaper, strategy becomes less about possessing problem-solving horsepower and more about deploying it well. The premium belongs less to the person who can produce the twentieth competent memo and more to the person who can decide which memo should exist, why, for whom and with what consequences.

There is a second premium too: verification and liability. When GPS became universal, navigation got cheaper. Responsibility for

safety did not. Something similar will happen in professional services. Cognitive production will become cheaper; operational judgement, governance, verification and sign-off will become more valuable.

4. What leaders should do

"Embrace AI" is too vague to be useful. The serious response has three parts.

First, map exposed work rather than exposed job titles. The relevant unit is not "consultant", "lawyer" or "analyst". It is research, first drafting, modelling, coding, checking, synthesis, client interaction, judgement and sign-off.

Second, redesign the workflow. Move machines into low-risk cognitive production, structured reasoning and challenge tasks where they improve speed, breadth and consistency. Move humans deliberately towards problem selection, context, client trust, decision rights, verification and accountability.

Third, treat transition costs as real economics. People cannot all move capability, sector, geography or seniority at the same speed. Firms will need clearer role design and deliberate reskilling. Policymakers will need to talk about adjustment before politics talks through it.

Before deciding what to do, leaders need to accept one fact: AI is not a referendum. It is a fact. We did not refuse steam power, electricity, computing or the internet because they disrupted work; we reorganised work around them. Fear is understandable, especially where roles are exposed. But fear is not a strategy, and refusal is not an operating model. The relevant questions are where to use AI, how to govern it, what not to automate, what humans must still own, and how quickly institutions can adapt.

The right response is therefore not panic and not worship. It is acceptance, experimentation, redesign and governance.

Closing thought

We have built tractors for the mind. The long-run result is likely to be a larger economy. The short- and medium-run result will be uneven, slower than the hype suggests, and uncomfortable for some groups that assumed they were safest.

The real question is not whether intelligence will remain useful. It is what kind of intelligence remains scarce enough to command a premium. As models move from fluency to reasoning, memory and tool use, the repriced zone will expand. My bet is still that value shifts from possession of analytical capability to deployment of it: choosing what work should exist, to what standard, and with what consequence.

In the next article, I will go one level deeper: what do we actually mean by intelligence, and where will human judgement, social context and creative

tension still matter as machine intelligence moves from fluency to reasoning, memory and bounded agency?

Disclosure

This article was AI-assisted. The argument, judgement and conclusions are my own.

Jesper Kjerside is co-founder of *taonta*, an AI-native advisory firm, and *fastland*, the software company behind it.

AI series, Part 1 of 5 | 29 July 2026