CHAPTER 14 · ESSENTIALS · PART IV · WHAT MARGIN IS FOR

14. The Judgment Dividend

Generation got cheap. Evaluation did not. The value of a human mind migrates to judgment — and judgment runs on margin.

The chapter in brief

Cara and Dev were given the same AI tools on the same Monday, and a year later they were not in the same place. The parable is invented and flagged as such; the chapter then earns it with a field experiment.

Underneath is a reversal most people are living through without having named it. For essentially all of history, generation was expensive and evaluation cheap: writing the report took days, reading it took an hour. That has flipped, and when one input becomes abundant the value flows to the one still scarce. Judgment is not intelligence, critical thinking or experience, though it borrows from all three. It is integrative, which makes it the most expensive operation you own — and the first to go at zero.

What to hold on to

- **Two ways judgment fails at zero.** *Rubber-stamping*: you approve because evaluating properly is expensive and you have nothing to spend. *Blanket distrust*: you reject everything and pay full price for the tool while getting none of its help. Opposite-looking, same failure.
- **The jagged frontier.** The machine's competence is uneven in ways invisible from outside, so knowing where the edge runs is not a fact you can look up. It is a judgment you keep making.
- **Judgment is trainable.** Tetlock found the average expert barely beat chance — and also found a learnable set of practices that produced measurably better forecasters.

The evidence, honestly

- The anchor is a preregistered field experiment (Dell'Acqua and colleagues, 2023) with the Boston Consulting Group: 758 consultants, randomly assigned to work with or without GPT-4. Inside the frontier of what the AI did well, consultants with AI finished 12.2 per cent more tasks, 25.1 per cent faster, at markedly higher rated quality.
- On a task deliberately placed *outside* that frontier — built so the output would be fluent, plausible and wrong — consultants using AI were 19 percentage points *less* likely to reach the correct answer than colleagues with no AI at all.
- The least-quoted finding is the one the book most wants you to see. A group that also received prompt-engineering training did better inside the frontier and worse outside it — falling 24 points against 13 for the untrained — because the training increased how much output people kept verbatim. Becoming better at using the tool made them worse at doubting it.

THE AI THREAD This chapter is the AI thread, whole. Every sidebar in this book was the same argument arriving in installments: the value of a human mind migrates to what the machine cannot own, and that runs on margin.

Your Margin Move — The Decision Journal, AI Edition

For every consequential AI-assisted decision this month, record five things — the first **before you look**. Your read, the machine's proposal, what you changed and why, your confidence as a probability, a review date. That comparison is the engine of calibrated trust.

*Printable worksheet: **ThoughtMargin.com/downloads/the-decision-journal-ai-edition.pdf***

Carry forward. *Judgment is what margin buys at work. The same reserve buys something with no professional value at all. · From Thought Margin: Reclaiming the Capacity to Think in the Age of AI by Rachael-Renée Rychling. A companion, not a substitute — the caveats that make the claims trustworthy live in the book.*