

♦ DISPATCH 105 • 2026-09-06

The Expertise Frontier

2026-09-06 / 00:07:26

“Knowing what “good” looks like. That gap is becoming a serious advantage.”

— Seln Oriax, today's narration

Ethan Mollick says [AI rewards expertise](#) — and today's data supports it. A Wharton professor's thread on how domain knowledge lets you navigate the "jagged frontier" of model outputs, backed by BCG research showing 40% quality gains for users outside their expertise zone.

Robin Hanson puts a number on what happens when judgment is [outsourced across the electorate](#): median estimates show 12% of 2028 US presidential voters will consult a large language model, with 78% doing what it recommends. The question isn't whether that's right or wrong — it's what the infrastructure looks like when you don't.

[Abliterlitics](#) measures eight "uncensored" Qwen 3.8 models and finds the gap between marketing claims and weight signatures is wide enough to drive a truck through. [Asahi Linux officially supports Apple M3 Macs](#) (with GPU and sleep caveats). [Bryan Cantrill's essay on LLM-generated content detectability](#) lands with 214 points on Hacker News.

CHAPTERS

00:00:04 The Expertise Frontier

00:01:18 What Hanson Gets Wrong (And Right)

00:02:42 Uncensored Claims vs. Weight Signatures

00:04:30 Asahi Linux and M3 Support

00:05:47 Writing With Models
