

Bring Your Own Numbers

2026-05-17 / 00:23:52

“The agents do the typing. The humans do the specification.”

— Lenar Kess, today's narration

A Sunday show about doing your own arithmetic. Mustafa Suleyman gives the white-collar tier eighteen months, in a piece whose own counter-data sits two paragraphs down. The State of Brand argues every AI subscription is a subsidized loss-leader two weeks away from a forcing function. William Angel runs the tokens-per-hour math on an M5 MacBook Pro and finds OpenRouter cheaper. Frederick Vanbrabant uses The Goal to explain why agents move the bottleneck rather than break it. Marlene Mhangami's Playwright talk shows the cleanest pattern for tests AI should write. Calif's public M5 / MIE exploit write-up lands. Artem Loenko explains why every chat UI keeps ending up on a browser engine. And Luke Lanchester's MCP hello page is the small fix I most enjoyed this week.

- [Suleyman's eighteen-month claim, and the article that knocks it down](#)
- [The subsidy that ends on June 1](#)
- [Apple Silicon costs more than OpenRouter](#)
- [Frederick Vanbrabant on why AI doesn't speed up your process](#)
- [Marlene Mhangami: tests AI writes vs tests AI verifies](#)
- [Calif on the five-day M5 / MIE exploit, in their own words](#)
- [Artem Loenko: native all the way, until you need text](#)
- [Luke Lanchester's MCP hello page](#)

CHAPTERS

- 00:00:04 Eighteen months, and the article that knocks it down
- 00:03:01 The subsidy that ends on June 1
- 00:06:50 Apple Silicon costs more than OpenRouter
- 00:09:28 AI doesn't speed up your process — it moves the bottleneck
- 00:12:18 Tests AI writes, tests AI verifies
- 00:15:26 The five-day exploit, in their own words
- 00:18:19 Native all the way, until you need text
- 00:21:18 The MCP hello page