

Sarah Guo's "Friction" essay heads a set of agent-economy reads, alongside Chamath and Levie on model harnesses

Recommended Reading from Tech Founders

2026-10-13

Sarah Guo's "Friction" essay heads a set of agent-economy reads, alongside Chamath and Levie on model harnesses

By Recommended Reading from Tech Founders • October 13, 2026

This brief covers what founders and investors recommended over the past day. Sarah Guo shared her essay on personal agents with reading to go with it, Chamath Palihapitiya and Aaron Levie shared work on harnesses, and Garry Tan, Demis Hassabis and Aravind Srinivas added shorter picks.

Sarah Guo: "Friction," with her own reading list

Guo shared her essay *Friction*. She summed it up as "personal agents, the end of friction, economic consequences of finally getting around to it" [1]. The essay argues that many businesses profit from customers' inattention: "Somewhere, a CRM still records your exhaustion as loyal recurring revenue" [2]. Once agents do the comparison shopping, she writes, "A bank could keep me perfectly happy and still have to win over my personal agent every day" [2]. She also expects public benefit payouts to rise as agents make it easier for people to claim them [2].

She replied to her own post with the sources behind the essay: - **Personal Agent Protocol**, via Bret Taylor's post [3]. Sierra has published a first draft at personalagentprotocol.org and is asking for feedback. The goal is that "companies know when they're dealing with an agent and who it represents" [4]. - **Patrick O'Shaughnessy's interview with Instinct founder Noah Shinn** (post) [5]. The essay draws on it for Shinn's picture of an agent that arranges dinner while you're on the way home from the airport [2].

Chamath Palihapitiya and Aaron Levie: the harness matters as much as the model

Chamath shared an 8090 Labs analysis. He says he asked the 8090 Labs team to run it after he stopped trusting leaderboards, because “improvements in a ranking don’t show up in the actual work” [6]. The study ran 7 models through 5 harnesses. Claude Fable 5.1 scored 41% in Claude Code and 61% in Codex, while GPT-5.6 Sol and Grok 4.6 “barely cared which harness they were in” [6]. His takeaway: choosing from a model-only leaderboard means deciding “based on only half the information” [6].

Levie pointed to an arXiv paper that Pedro Domingos had shared with the line “the model swallows the harness” [7]. Levie called it “quite cool research” and “the eventual architecture of agents,” arguing that AI “will obviously be better at determining its own harness once it sees the task at hand” [8].

Garry Tan: three picks

- **DeepRSI** by @jparkjmc (post), a benchmark that has agents work 48 hours on ML tasks taken from ongoing proprietary research [9]. Tan: “Next 18 months is the Kitty Hawk moment for RSI. Evals will be how you can tell” [10].
- **ThePeelPod’s clip on pro rata** (post), about his own Initialized Fund 1. According to the clip, Coinbase, where he kept taking his pro rata, produced the final 53x [11]. Tan: “pro rata matters and is powerful for investors” [12].
- **A City Journal piece on Walker Percy’s *Love in the Ruins***, shared by Rob Henderson (article). Henderson quoted the line that things fell apart because “things stopped working and nobody wanted to be a repairman” [13]. Tan’s comment: “We need repairmen” [14].

Demis Hassabis: Google’s AMIE paper in *The Lancet*

Hassabis highlighted the study Sundar Pichai announced (post) and said it was “Amazing to see research move from the lab to real-world clinical utility” [15]. According to Pichai, it is the first prospective study of patient-facing conversational diagnostics in a real clinic. Clinicians said AMIE’s summaries helped them prepare for 75% of visits, and its differential diagnoses matched doctors’ final diagnoses 90% of the time [16].

Shorter mentions

- **Aravind Srinivas** shared Perplexity’s engineering write-up on the shared filesystem for agents behind Perplexity Projects. Perplexity is his own company [17]. His highlight: plain **ls** and **grep** beat semantic search 63% to 37% on answer accuracy for financial document analysis [18].

- **Alexandr Wang** pointed readers to Meta’s Muse Wonder research blog, promoting his own lab’s model [19]. He described it as built to power personal agents, with active memory and domain expertise in finance and law [20].
- **Tim Ferriss** posted a clip of Kevin Rose “on the coming ‘headless’ era of software” [21].
- **Sahil Lavingia** posted pages from “the new Michael Lewis book.” He didn’t name the title or say why he shared it [22].

Sources

1. X post by @saranormous
2. X article by @saranormous
3. X post by @saranormous
4. X post by @btaylor
5. X post by @saranormous
6. X post by @chamath
7. X post by @pmddomingos
8. X post by @levie
9. X post by @jparkjmc
10. X post by @garrytan
11. X post by @ThePeelPod
12. X post by @garrytan
13. X post by @robkhenderson
14. X post by @garrytan
15. X post by @demishassabis
16. X post by @sundarpichai
17. X post by @perplexity_ai
18. X post by @AravSrinivas
19. X post by @alexandr_wang
20. X post by @alexandr_wang
21. X post by @tferriss
22. X post by @shl