Gurley Tells Regulators to Read Pinker on AI Safety; Collison Rereads Tocqueville for the AI Era

Recommended Reading from Tech Founders

2026-11-25

Gurley Tells Regulators to Read Pinker on AI Safety; Collison Rereads Tocqueville for the AI Era

By Recommended Reading from Tech Founders • November 25, 2026

Bill Gurley backs Steven Pinker's Quillette essay on AI safety versus doom, Patrick Collison uses Democracy in America to think about AI, and Sarah Guo points to Harvey's open-weight legal research and a Paul Graham essay.

Gurley: Pinker's essay is required reading for regulators

Bill Gurley called **Steven Pinker's** new Quillette article, *Will We All Go Together When We Go?*, "the most thorough and intelligent piece of writing on AI and AI safety." He went further: AI regulation "should be led by an impartial and thoughtful person like him," and "All regulators must read!" [1]. Pinker describes the essay as covering intelligence, consciousness, altruism, rationality, and "the difference between AI safety and AI doom" [2]. The prior period's picks from Altman and Schmidt argued against handing control to AI. Gurley's pick comes at the debate from the anti-doom side.

Collison: *Democracy in America* as an AI-era book

Patrick Collison wrote a long essay-post on **Alexis de Tocqueville's** *Democracy in America*. He calls it "a giant Substack from a young backpacker" [3], and in his view it is "a book about a new kind of society" with lessons as "we sit on the cusp of another upheaval" [3]. He takes three lessons from it:

- **Work and honor:** Tocqueville saw that slavery made work "dishonorable" for slaveowners. Collison asks whether AI that removes the need to

work could do the same [3].

- **Restraint:** Tocqueville saw religion as a check on what a self-governing society might do. That leads Collison to ask "which are the ideas, philosophies, religions, or doctrines that will help us wield AI well?" [3]
- **Soft despotism:** Tocqueville feared a "tutelary power" that keeps people in "perpetual childhood." Collison notes this looks a lot like the Iain Banks Culture novels that AI proponents often cite [3]. He ends by asking: if AI settles our disputes, "we might get better outcomes, but will we get better people?" [3]

Guo: Harvey's open-weight research and Paul Graham on identity

Sarah Guo called a Harvey–Baseten writeup by Niko Grupen, Gabe Pereyra and team "Important work" [4]. She says it shows the training recipe can be adapted across model families, and that US open-weight models are getting more competitive at agentic legal work [5].

In her own essay, *good beginner*, she links **Paul Graham's** *Keep Your Identity Small*. She is trying to apply it so she can enjoy math "with all the curiosity and competitiveness, and less angst about how much worse I am than " [6].

Khosla: try Ternary Bonsai 2.5

Vinod Khosla said PrismML's new model is "Worth everyone who uses models trying" [7]. PrismML says Ternary Bonsai 2.5 27B fits in 5.9 GB and scores 75.6 on SWE-Bench-Verified. It is released under Apache 2.0 [8].

O'Shaughnessy: Bret Taylor, and his Scott Wu episode

Patrick O'Shaughnessy also praised Colossus's Bret Taylor Q&A, which Sarah Guo recommended in the last brief. He said he "read this yearning for more conflict" but came away "feeling it's possible to be great without some giant trauma or trade off in your personality" [9, 10]. He also promoted his own interview with Cognition's Scott Wu, calling Wu "one of my favorite entrepreneurs" [11, 12].

Shorter picks

- **Reid Hoffman** recommends "to everybody" using ChatGPT as a tutor. Ask it to quiz you on what you know, then have it build lessons. He did this for quantum computing and called it "magically good." Next he plans to do the same for probabilistic computing and Bayesian math [13]. Video
- **Bing Gordon** (on Shane Parrish's channel) pointed to **David Ogilvy's** books. His favorite nugget for junior account execs: "Always know more than anybody in the room about your client's business" [14]. Video

- **Scott Belsky** suggested his family watch *You Can See Everything*, now out on digital, before Thanksgiving dinner "to up-level the convo" [15, 16].
- **Tobi Lütke** amplified a developer who heard Lütke's podcast tip: have AI find the implicit state machines in a codebase and build a TLA model of them. The developer turned the tip into a skill and found "10+ bugs." Lütke's caption: "My words can heal your codebase" [17, 18].

Sources

1. X post by @bgurley
2. X post by @sapinker
3. X post by @patrickc
4. X post by @saranormous
5. X post by @saranormous
6. X article by @saranormous
7. X post by @vkhosla
8. X post by @PrismML
9. X post by @patrick_oshag
10. X post by @colossusmag
11. X post by @patrick_oshag
12. X post by @patrick_oshag
13. In Silico Veritas: The Case for Strategic Optimism About AI | Reid Hoffman & Marko Ahtisaari
14. Ex-Amazon Board Member Reveals Everything Bezos Does to Win
15. X post by @scottbelsky
16. X post by @FilmUpdates
17. X post by @azolotov
18. X post by @tobi