

A FINRA for AI? Be Careful What You Wish For

By Evan LeFlore | *Director of AI, Innovation, and Economic Opportunity* and Benjamin Schiffrin | *Director of Securities Policy*

October 8, 2026

Demis Hassabis has a plan to govern one of the world's most powerful technologies, and it's modeled on one of Wall Street's most contested institutions.

[In an essay](#) published this summer, the Google DeepMind CEO called for a U.S.-initiated frontier artificial intelligence standards body that is explicitly “modelled on a federally overseen public-private partnership or self-regulatory organization, much like the Financial Industry Regulatory Authority (FINRA).” FINRA is the self-regulatory organization for securities brokers. That means it is essentially run by the same brokers that it regulates. Securities firms fund FINRA, and although FINRA is a non-profit corporation, [its members are for-profit entities](#).

The need for AI oversight is more clearly becoming a necessity, even to the model builders most likely to profit from limited regulation. Hassabis's call for oversight puts him ahead of an administration that calls for stripping away AI regulation and that threatens to withhold federal funding from states that do not gut their own AI rules.

But before Washington adopts a FINRA-like oversight model, it's worth asking an important question: does FINRA actually work?

Our assessment is: not well enough. And the reasons it falls short are the reasons AI oversight should be built differently.

The Cautionary Tale: Conflicts, Capture, and Weak Deterrence


FINRA writes and enforces the rules for almost 3,500 brokerage firms. But it is not a government agency. Instead, it is [composed of securities brokers themselves](#). These firms fund FINRA, which means they control FINRA's budget. The fact that FINRA's budget is “substantially sourced from the brokers it regulates, combined with the fact that FINRA executives are paid large salaries,” raises [an obvious concern](#), described concisely by state-level securities law enforcers: if executives “are paid relatively opulent salaries by the brokers whom they police, the motivation to regulate can easily be dampened.” Broker-dealer firms also elect representatives to FINRA's Board of Governors, including both “industry” and “public” members, many in the latter camp with prior professional associations with broker-dealers. This structure allows the industry to [regulate itself with limited public oversight](#).

Indeed, [the risk of self-regulatory organizations](#) is, as noted by financial regulation scholars, the “potential for conflicts of interest inherent in a grant of regulatory powers to an organization representing, essentially, the regulated entities themselves.” Critics of self-regulatory organizations [describe the model](#) as “self-serving, self-interested, lacking in sanctions, beset with free rider problems, and simply a sham.”

With respect to FINRA specifically, there is “an inherent [conflict of interest](#)” between FINRA’s regulatory objectives and the private for-profit motives of its members. One [view](#) of the system of self-regulation in the securities industry is that it is “no longer effective” at meeting the public policy goals of necessary regulation. For these reasons, [other jurisdictions with securities markets](#) “have abandoned the self-regulatory system, concluding finally that the conflicts of interest make self-regulation unworkable.”


The decline in FINRA enforcement actions over the last decade suggests a weakness of a self-regulatory model. In 2025, FINRA [brought the fewest disciplinary cases in a decade](#). It brought [22% fewer cases](#) than it brought in 2024, which continued the downward trend of enforcement actions since 2015. In 2015, FINRA pursued more than 1,300 cases, which was [more than three times](#) the number of cases it brought last year.

FINRA’s own statistics bear this out. Although FINRA tracks its enforcement cases using multiple methods, its enforcement activity has hit rock bottom regardless of the metric used:



Source: [Eversheds Sutherland](#)

Even when FINRA brings enforcement actions, it is going easier on firms. The fines and restitution FINRA forces firms to pay are also down from their peak a decade ago.



Source: [Eversheds Sutherland](#)

Although there is no way to be certain of the reason for the decline, several possible explanations are attributable to the fact that the organization is not a government agency but an association of member firms. One possible explanation is that, with the [passage of time since the 2008 financial crisis](#), FINRA has become less focused on enforcement. Its member firms may no longer see the need to appear tough on white-collar crime. Another possible explanation is that the firms FINRA oversees, and that comprise FINRA's membership, have prevailed on FINRA to adjust their exam and surveillance practices. As FINRA's former head of enforcement [said](#), the statistics indicate that FINRA is more "focused on member relations than retail investor protection." A third and related explanation is that FINRA appointed a new CEO in 2016, who prioritized [engagement with brokers](#). A self-regulatory organization may be more likely than a government agency to be concerned with how its enforcement is perceived by the entities it is supposed to regulate.

Whatever the reason, the lack of recent enforcement activity at FINRA should give policymakers pause before exporting the self-regulatory organization model to other industries. Indeed, FINRA's enforcement activity and ability to deter misconduct was already considered [weak](#) even before this precipitous drop in enforcement activity over the last decade. So there is little evidence that, as a self-regulatory organization, FINRA has proven capable of adequately policing the securities brokers it exists to regulate.

FINRA, for all of its flaws, still answers to the SEC and to the public by extension, at least on paper. Its enforcement program [is subject to SEC oversight](#). Rule changes regarding broker

conduct are also subject to public notice and comment, and must be approved by the SEC.

So, theoretically, the SEC could sue FINRA if it believes FINRA is not sufficiently rigorous [in enforcing its own rules](#). The SEC has not often done so, however. And although the SEC has the authority to regulate other parts of FINRA as well, including the organization's oversight, [historically](#) the SEC "has only barely supervised FINRA's governance."

The Backstop that Isn't There

As we've established, the self-regulatory model for brokers has some real limitations. Unfortunately, Washington is signaling that should an AI-focused SRO be established, it would have even less teeth than what we're accustomed to. First, unlike FINRA, which has the SEC looking over its shoulder (however imperfectly), AI currently has no consolidated regulator to oversee the rulemaking or enforcement of an industry SRO. What's more, the administration is signaling a trend in the wrong direction, with its [AI Action Plan](#) directing federal agencies to strip away "onerous regulation" and elevate industry-led standards over binding rules. A [more recent executive order](#) goes further, creating a task force to challenge state AI laws and threatening federal funds for states that don't fall in line. Meanwhile, the SEC [withdrew its own proposed rule](#) on AI-driven conflicts of interest in predictive data analytics, and the main AI guidance broker-dealers have received to date, a [2024 FINRA notice](#), "does not create new legal or regulatory requirements," instead reminding firms that existing rules still apply.

Taken together, you get a preview of a worse version of governance for AI oversight: a self-regulatory organization model without even FINRA's imperfect SEC standing behind it.

It's important to note that Hassabis isn't proposing such an outcome. Hassabis's version of a FINRA for AI would have a board composed of independent technical experts. It would evaluate frontier AI models such as ChatGPT or Claude for risks related to cyber, biological, and deception capabilities, and start with voluntary model reviews before graduating to mandatory ones. When it comes to funding such an agency, Hassabis [recommends](#) that it "needs to be substantial and likely mostly come from industry." The problem is that Washington's current trajectory is likely to produce the worst version of a FINRA for AI if this framework were adopted without a strong public anchor.

Hassabis is not alone in advocating for a self-regulatory structure, either. Researchers like [Dean Ball](#), now at OpenAI, have proposed "private governance" bodies that would certify AI labs in exchange for liability protection. California floated "[independent verification organizations](#)" that would do something similar. Industry consortia like [Frontier Model Forum](#) already set voluntary technical standards.

Why This is a Finance Story, Not Just a Tech Story

This isn't an abstract governance debate. How AI is governed will have substantial downstream effects for our individual and collective safety. And it's also a financial stability question: AI is already reshaping the sector it was supposed to help police, and a handful of AI model and cloud providers now underpin a significant share of financial-sector infrastructure, a concentration risk the Government Accountability Office flagged as a threat to financial stability [in a May 2025](#)

[report](#). When many firms lean on the same small set of underlying models, correlated errors and herding behavior (that is, many trading algorithms reaching the same conclusion at the same time) can amplify volatility rather than dampen it. This is the same general dynamic that [regulators have long worried about](#) when it comes to automated trading. AI is also moving into credit underwriting, investment advice, and fraud detection, raising fair-lending and consumer-protection questions that current banking regulators have only begun to address.¹

An AI standards body that inherits FINRA's weaknesses—from weak enforcement to industry-tilted governance to a thin public backstop—would be poorly positioned to catch these risks before they undercut financial market integrity and households' finances.

A Better AI Oversight Model

Hassabis is asking the right question, but a FINRA-like model points towards the wrong, or at least incomplete, answer. If Washington is going to build something like a standards body for frontier AI, it should learn from what FINRA got wrong as well as what it does well:

- **Independent public oversight with teeth, not a paper backstop:** a genuinely resourced government regulator that reviews rules *and* outcomes, unlike the SEC's "limited or no oversight" of parts of FINRA.
- **No industry capture of governance:** a board comprised with a majority of *actually* independent members who represent public-interest and worker/consumer representation, not a body mostly governed and funded by the firms it polices.
- **Funding that doesn't buy leniency:** stable funding (a real FINRA strength) but structured so the biggest payers aren't the least policed.
- **Enforcement with real deterrence:** mandatory, transparent penalties large enough to matter and individual accountability, reversing the "[kinder, gentler](#)" drift.
- **Statutory authority, not delegation-by-executive-order:** Congress should define the mandate, rather than leaving AI oversight to voluntary industry compacts and liability shields.

Hassabis is right that AI needs oversight built by people who understand the technology and build it fast. But speed and expertise aren't the difficult part of self-regulation. Accountability is. FINRA's record shows what happens when an institution's design answers that question in the industry's favor. AI is too consequential and moving too quickly to make the same mistake twice.

¹ Banking regulators' April 2026 [model-risk-management guidance](#) notably excluded generative and agentic AI systems from its scope.



Better Banks | Better Businesses
Better Jobs | Better Economic Growth
Better Lives | Better Communities

Better Markets is a public interest 501(c)(3) non-profit based in Washington, DC that advocates for greater transparency, accountability, and oversight in the domestic and global capital and commodity markets, to protect the American Dream of homes, jobs, savings, education, a secure retirement, and a rising standard of living.

Better Markets fights for the economic security, opportunity, and prosperity of the American people by working to enact financial reform, to prevent another financial crash and the diversion of trillions of taxpayer dollars to bailing out the financial system.

By being a counterweight to Wall Street's biggest financial firms through the policymaking and rulemaking process, Better Markets is supporting pragmatic rules and a strong banking and financial system that enables stability, growth, and broad-based prosperity. Better Markets also fights to refocus finance on the real economy, empower the buy-side and protect investors and consumers.

For press inquiries, please contact us at press@bettermarkets.org or (202) 618-6430.



[SUBSCRIBE](#) to Our Monthly Newsletter

FOLLOW US ON SOCIAL



2000 Pennsylvania Avenue NW | Suite 4008 | Washington, DC 20006 | 202-618-6464 | www.bettermarkets.org
© 2026 Better Markets, Inc. All Rights reserved.