

Regulators Must Avert Overreach When Targeting AI

By **Jack Solowey** (September 13, 2023)

The splashy introduction of generative artificial intelligence tools in recent months has been met with a flurry of government attempts to grapple with the newest new thing.[1]

While many efforts look to tackle AI generally, U.S. financial regulators have been busy addressing AI's role in capital markets, consumer finance and banking.[2]

Though covering wide ground, financial regulators' 2023 statements, guidance, and proposed rules touching AI have coalesced around five core themes: competence, data protection, bias mitigation, duties to customers and transparency.



Jack Solowey

When both setting and pursuing AI priorities, policymakers generally should be wary of counterproductive interventions that hamstring the very innovations that could help to improve predictions, expand access to financial services and manage longstanding risks.

Avoiding unintended consequences means validating assumptions about AI risks, surgically targeting mitigation strategies and studiously accounting for the potential gains lost to regulatory encroachments.

Below is a high-level summary of financial regulators' AI policy goals, followed by a discussion of typical dangers of regulatory backfires, as well as principles for averting overreach.

Financial Regulators' Five AI Goals

In 2023, the U.S. Securities and Exchange Commission, Commodity Futures Trading Commission, Consumer Financial Protection Bureau, U.S. Department of the Treasury, Federal Reserve Board of Governors, Federal Deposit Insurance Corporation, National Credit Union Administration and Federal Housing Finance Agency — either at the agency level or through statements by their officers — have collectively homed in on the following five AI policy issues.

While not every regulator addressed every issue, every issue received attention from multiple financial regulators.

1. Competence

Financial regulators generally want AI tools to be fit for purpose.

Specifically, they are concerned about whether one can be confident in models' estimates.[3] They fear chatbots can give inaccurate, unreliable or insufficient information, and can fail to offer meaningful customer assistance.[4]

And regulators wish to see automated systems developed with appropriate regard for the contexts in which they'll be used and the nature of their users.[5]

2. Data Protection

Given AI models' voracious appetites for data, financial regulators are prioritizing cybersecurity and data privacy when addressing AI.

Cybersecurity-wise, they fear hackers can disrupt AI-enabled services and reveal personal information in data breaches.[6]

As for privacy, financial regulators are concerned with safeguarding personal information, both in terms of the data used to train AI models and the data consumers provide to applications.[7]

3. Bias Mitigation

The policy priority receiving perhaps the greatest attention from financial regulators is the risk that machine learning and AI will "perpetuate unlawful bias" and "automate unlawful discrimination," as a group of officials including CFPB Director Rohit Chopra put it.[8]

This typically takes the form of concerns that models' outputs will be prejudiced by datasets that are unrepresentative or that incorporate historical bias.[9]

These biases could manifest in unlawful activity as a result of disparate treatment, i.e., treating an individual differently on the basis of a prohibited factor, or disparate impact, i.e., a neutral practice disproportionately excludes certain people on a prohibited basis.[10]

4. Duties to Customers

Regulated financial entities often have special legal responsibilities to their customers. These include fiduciary duties, such as loyalty and care, with roots in common law, as well as regulatory obligations regarding fair conduct.[11]

Financial regulators fear that AI tools may confound financial institutions' ability to honor their duties to customers.

For example, regulators are concerned that certain models may embed conflicts of interest where model training prioritizes the regulated entity's profits over a client's interests.[12] And regulators warn that a chatbot's failure to provide clear answers could constitute a form of unfair dealing.[13]

5. Transparency

Financial regulators are wary of certain AI models being black boxes, meaning that no one, including the models' developers, may understand exactly why a model made a particular prediction or why that prediction is likely to succeed.[14]

Regulators consider this black-box problem to be both something of a risk in itself — potentially complicating disclosures, auditability and risk management — as well as a challenge when dealing with many of the other risks they identified.[15]

For example, they point to difficulty in effectively rooting out bias or conflicts of interest where the bases of decisions remain obscure.[16]

Confronting the Risks of Overreach

The primary risk of policy overreach on financial AI is the loss of potential gains from innovation, such as reduced costs, novel strategies and opportunities, and improved access to financial services.

And a primary driver of such counterproductive interventions is the mistaking of existing, pre-AI risks for new, AI-created ones.

Avoiding these missteps is a matter of determining whether supposed risks warrant new interventions — or any at all — pursuing minimally invasive mitigation strategies, and always considering the costs of lost innovation.

Validating Assumptions About Risks

Perceived risks from AI in finance typically have a clear analog in, well, analog financial activity.

Before assuming that the use of AI in finance presents novel risks or heightens existing ones, one should ask, as my colleague Jennifer Schulp flags, whether the application of AI to a financial activity meaningfully changes the character of that activity.[17]

It is sometimes contended that the expanded scale or scope of activity enabled by a new technology like AI heightens risks, but it is important to parse in what specific sense the risk is then considered aggravated.[18]

For instance, the ability to reach more customers more often does not necessarily entail heightened or novel risk in any given customer interaction.

And even where one insists that the quantity of activity has a quality of risk all its own, there remains the matter of identifying why existing rules are inadequate for confronting these risks.[19]

Moreover, when it comes to known risks, one must consider whether faults attributed to AI are already tolerated in analog alternatives.

For example, where existing frameworks do not take zero-tolerance approaches to, say, inaccurate predictions or unsatisfactory assistance from human financial advisers or customer service agents, the presumption also should be against applying zero-tolerance standards to AI tools.

Surgically Targeting Risks

The black-box issue described above exemplifies a perceived challenge that, when handled maladroitly, can jeopardize potential AI benefits. In such cases, potential harms ought to be surgically targeted.

To understand how a blunt approach like severely restricting the use of black-box models can cause more harm than good, it's worth taking a brief, illustrative foray beyond financial services.

In 2020, black-box AI led Massachusetts Institute of Technology scientists to discover a new potent antibiotic, helicin, for treating antibiotic-resistant infections.[20]

After the initial in silico discovery, scientists could test whether the AI-recommended molecule was effective without needing to understand how the AI arrived at its prediction and without cutting off a promising investigative pathway for fear of its unfamiliar form of intelligence.[21]

Black-box AI predictions in financial markets similarly can produce testable outputs that reveal opportunities to harvest gains and hedge risks.[22]

Accordingly, outcome-oriented frameworks, unlike prescriptive restrictions on AI processes that are not yet fully explainable, can address risks without eliminating benefits.

Considering Lost Benefits

To avoid undermining potential benefits from AI in finance, the expected value of such upsides should always be given equal consideration alongside that of AI risks.

AI tools can both create new opportunities and mitigate existing risks. The ability to provide financial services at lower cost and greater scale can help to achieve the important policy goal of expanding access to the underserved.[23]

Relatedly, one regulatory fear for the use of certain chatbots in finance is that these tools can hinder access by those "with limited English proficiency," in the words of a CFPB report.[24]

But when it comes to cutting edge chatbots that leverage large language models and multilingual machine translation, there's good reason to expect that ultimately the opposite may be true. Not only can advanced AI tools help to translate concepts into more languages than one is likely to find at a typical human-staffed call center, but also those tools have the capacity to tailor explanations based on one's education level.[25]

In multiple domains, AI could address some of financial regulators' primary concerns.

Indeed, Federal Reserve Vice Chair for Supervision Michael Barr put it well when he recognized in a discussion of the risk of algorithmic bias in lending that "new artificial intelligence techniques such as machine learning have the potential to leverage [alternative data sources] at scale and at low cost to expand credit to people who otherwise can't access it." [26]

Regulators should heed this example and always give AI's potential benefits their due.

Conclusion

For decades, leading-edge financial firms have been among the early adopters of AI techniques, leveraging insights from advances in natural language processing and machine learning.[27]

For better or worse, understanding the future of AI policy likely will mean grappling with financial policy specifically.

While the potential benefits from AI in finance are by no means guaranteed, as financial regulators pursue AI policy, they should avoid stymieing technology that could enhance forecasts, better reach the historically underrepresented, and spot risks and opportunities.

Jack Solowey is a policy analyst at the Cato Institute's Center for Monetary and Financial Alternatives.

The opinions expressed are those of the author(s) and do not necessarily reflect the views of their employer, its clients, or Portfolio Media Inc., or any of its or their respective affiliates. This article is for general information purposes and is not intended to be and should not be taken as legal advice.

[1] See Michael Lewis, *The New New Thing: A Silicon Valley Story* (2000).

[2] Anna Lenhart, *Federal Legislative Proposals Pertaining to Generative AI* (118th), Institute for Data, Democracy & Politics, last updated August 15 2023. See *The White House, Blueprint for an AI Bill of Rights* (October 2022); *Securities and Exchange Commission Proposed Rule, Conflicts of Interest Associated with the Use of Predictive Data Analytics by Broker-Dealers and Investment Advisers* (2023); *Consumer Financial Protection Bureau Issue Spotlight, Chatbots in consumer finance*, June 6, 2023; and Michael S. Barr, *Furthering the Vision of the Fair Housing Act*, July 18 2023.

[3] Office of the Comptroller of the Currency et al. *Proposed Rule, Quality Control Standards for Automated Valuation Models* (2023), at 40638, 40648.

[4] CFPB Issue Spotlight, at 11-12.

[5] Rohit Chopra et al., *Joint Statement on Enforcement Efforts Against Discrimination and Bias in Automated Systems*, April 25, 2023, at 3.

[6] CFPB Issue Spotlight, at 17-18.

[7] Gary Gensler, "Isaac Newton to AI" Remarks before the National Press Club, July 17, 2023; and CFPB Issue Spotlight, at 7 and 18.

[8] The other officials issuing the statement were Kristen Clarke, Assistant Attorney General for the Justice Department's Civil Rights Division; Charlotte A. Burrows, Chair of the Equal Employment Opportunity Commission; and Lina M. Khan, Chair of the Federal Trade Commission. Chopra et al., *Joint Statement*, at 1.

[9] Chopra et al., *Joint Statement*, at 3.

[10] Barr, *Furthering the Vision of the Fair Housing Act*; *Federal Reserve Consumer Compliance Handbook, Federal Fair Lending Regulations and Statutes Overview*, last accessed August 23, 2023, at 2-3.

[11] Jack Solowey, "AI will revolutionise financial advice," *iai News*, May 10, 2023. See CFPB, *Final rules: fair conduct*, last accessed August 23, 2023.

[12] See Office of the Comptroller of the Currency et al. *Proposed Rule*, at 123 and SEC *Proposed Rule*, at 82.

[13] CFPB Issue Spotlight, at 19.

[14] See *Explainable artificial intelligence*, Wikipedia, last accessed August 23, 2023.

[15] See Christy Goldsmith Romero, Statement at the Technology Advisory Committee on Responsible AI in Financial Services, DAOs and Other DeFi & Cyber Resilience, July 18, 2023 and Barr, *Furthering the Vision of the Fair Housing Act*.

[16] See Chopra et al., Joint Statement, at 3 and SEC Proposed Rule, at 66 and 120.

[17] See Jennifer J. Schulp, Comment Letter Re: Request for Information and Comments on Broker-Dealer and Investment Adviser Digital Engagement Practices, October 1, 2021, at 2-3. See also Jack Solowey, "The Untested Assumptions in SEC Chair Gensler's Pivot to AI," *Cato at Liberty* (blog), July 31, 2023.

[18] See SEC Proposed Rule, at 144.

[19] There are instances where the introduction of new technology can require new guidance or rules to clarify the application, or lack thereof, of existing authorities. But a good rule of thumb is that such clarifying guidance is most appropriate where the new technology mitigates risks, making legacy rules inapt, not where technology merely expands the scale of a business activity carrying long-understood risks. See Jack Solowey and Jennifer J. Schulp, "We Need Regulatory Clarity to Keep Crypto Exchanges Onshore and DeFi Permissionless," *CoinDesk*, May 10, 2023.

[20] Anne Trafton, "Artificial intelligence yields new antibiotic," *MIT News*, February 20, 2020.

[21] Henry A. Kissinger, Eric Schmidt, and Daniel Huttenlocher, *The Age of A.I.* (2021), at 9-1.

[22] See Gregory Zuckerman, *The Man Who Solved the Market* (2019), at 224.

[23] See Schulp Comment Letter, at 2-3.

[24] CFPB Issue Spotlight, at 10, 20.

[25] See Angela Fan, *Introducing the First AI Model That Translates 100 Languages Without Relying on English*, *Meta*, October 19, 2020; Derek Gallimore, *Multilingual call centers: All you need to know*, *Outsource Accelerator*, last updated March 27, 2023; and Russ Roberts, Marc Andreessen on *Why AI Will Save the World*, *EconTalk* (July 10, 2023).

[26] Barr, *Furthering the Vision of the Fair Housing Act*.

[27] See Zuckerman, *The Man Who Solved the Market*, at 174, 224. But cf. Matt Levine, "AI Funds Don't Like AI Stocks," *Bloomberg*, August 22, 2023.