

Speech

"Isaac Newton to AI" Remarks before the National Press Club



Chair Gary Gensler

July 17, 2023

Thank you, Mark, for that kind introduction. As is customary, I'd like to note that my views are my own as Chair of the Securities and Exchange Commission, and I'm not speaking on behalf of my fellow Commissioners or the SEC staff. Nor for or by a generative AI model.

As the pandemic swept across Europe and the UK, students packed up their books and went home. One student at Trinity College in Cambridge left for his remote family farm. While in isolation, he continued his work in physics and math.

Now, I'm talking about the 1660s and the bubonic plague.[1] It is said that while in confinement Isaac Newton contemplated gravity among other profound insights, including a new form of math—calculus.[2]

That math innovation—perhaps the bane of some of your school experiences—is embedded in so much of life, from pharmacology to finance.

In the beginning months of a more recent pandemic, when many 21st century students were packing up their books and going home, there was a rollout of a new version of a generative artificial intelligence (AI) model, GPT-3.

Given the health crisis, that release may not have gotten as much attention as subsequent model releases. ChatGPT-4 was unveiled this year on Pi Day, March 14, the same day as two other competitor releases.[3] Just two days later, Baidu released Ernie Bot, a Chinese language competitor.[4]

A lot of the recent buzz has been about such generative AI models, particularly large language models. AI, though, is much broader. I believe it's the most transformative technology of our time, on par with the internet and mass production of automobiles.

Just like calculus, the math underlying AI is built on the groundwork of many others. Newton built upon many earlier mathematicians' work, including René Descartes's 17th century mathematical achievements—Cartesian planes and the famous formula for linear transformation, $y = mx + b$. [5] Yes, as you can tell, I'm a bit keen on math.

Similarly, AI and Newton's work on gravity both are based upon data and computational power. Newton built upon Galileo's work and gathered data from other scientists' research. His computational power was his own mind and quill.

Discussions about AI go back to at least the mid-20th century. You might remember Alan Turing from "The Imitation Game" movie and cracking of the Enigma code. In 1950, he wrote a seminal paper, "Computing Machinery and Intelligence," opening with, "I propose to consider the question, 'Can machines think?'"[6]

In the 1980s, world chess champion Garry Kasparov claimed a computer would never beat a grandmaster. In 1996, he beat IBM's Deep Blue only to be defeated by it a year later.[7]

Fast forward 14 years, IBM's Watson used brute force computational power to win "Jeopardy!" outpacing two human competitors. Watson had been fed more than 200 million pages of documents, from dictionaries to novels. [8] Ken Jennings, a "Jeopardy!" champ, wrote on his screen during final "Jeopardy!": "I, for one, welcome our new computer overlords." [9]

By 2016, AlphaGo, using new forms of AI, defeated a human world champion in the complex game of Go.[10]

Though perhaps not catching as many headlines as these computer versus human stories, or ChatGPT, we've already seen a lot of adoption of AI.

Text prediction in our mobile devices and emails has been commonplace for years. The Postal Service has been using it to predict addresses from human handwriting. It's being used for natural language processing, translation software, recommender systems, radiology, robotics, and your virtual assistant.

In finance, AI already is being used for call centers, account openings, compliance programs, trading algorithms, sentiment analysis, and more. It's also fueled a rapid change in the field of robo-advisers and brokerage apps.

Math, Data, and Computational Power

Just as Newton's work on gravity was based on math, data, and computational power, advances in all three are fueling AI's progress today.

The math underlying AI has moved from brute force computer power to neural networks to deep learning. Neural networks, inspired in part by the brain's architecture, have layers of individual nodes or neurons. The parameters (coefficients related to the connections between nodes) no longer number in the hundreds; they can be in the thousands, millions, or even billions.[11] Descartes would be impressed. Collectively, the models are being optimized to discern patterns in data and make predictions.

We also are living in an era of exponential data growth. Data are being drawn from our social media footprints, shopping, and spending patterns. They're being drawn from the explosion in sensors—cell phones, fitness devices, telematics in cars, appliances, cameras, and other so-called IoT sensors.[12] There are bound to be more sensors than humans in the United States. We've also seen the early stages of multimodal AI systems, which combine multiple types of data sets: video, audio, speech, images, and text to make even more accurate determinations.

Computational power has been increasing exponentially for decades. This has been true since Gordon Moore, co-founder of Intel, predicted so in 1965.[13]

Opportunities and Challenges

AI opens up tremendous opportunities for humanity, from healthcare to science to finance. As machines take on pattern recognition, particularly when done at scale, this can create great efficiencies across the economy.

In finance, there are potential benefits of greater financial inclusion and enhanced user experience.

When Newton pondered the laws of gravity, he thought of the micro, the apple, and the macro, the cosmos.

In that vein, I'm going to discuss what these AI advances might mean in the micro, to each of us as individuals, and in the macro, to our broader economy and society.

Micro: Narrowcasting

Today's AI-based models provide an increasing ability to make predictions about each of us as individuals. This growing capability facilitates being able to differentially communicate to each of us—and do so efficiently at scale. How might we respond to individualized communications? How might we respond to individualized product offerings? How might we respond to individualized pricing? In other words, narrowcasting. Quite a difference from Turing's day at the dawn of network broadcasting.[14]

On a daily basis already, we receive messages from AI recommender systems that are considering how we might as individuals respond to their prompts.

Models have been developed to assist in making decisions about who gets jobs, loans, credit, entry to schools, and healthcare, to name a few.[15]

This raises a host of issues that are not necessarily new to AI but are accentuated by it.[16]

Explainability, Bias, Robustness

AI models' decisions and outcomes often are unexplainable. Part of this is inherent to the models themselves. The math is nonlinear and hyper-dimensional, from thousands to potentially billions of parameters. It's dynamic with the ability to change and learn from new data and from the model's use.

Thus, the insights that come out of such models by design are inherently challenging to interpret in terms of accessibility to humans. Why does a model qualify one person for a loan while rejecting another?

AI also may make it more difficult to ensure for fairness. The outcomes of its predictive algorithms may be based on data reflecting historical biases as well as latent features that may inadvertently be proxies for protected characteristics. Further, the challenges of explainability may mask underlying systemic racism and bias in AI predictive models.

The ability of these predictive models to predict doesn't mean they are always accurate or robust. They might be predicting based upon some latent feature that leads to a false prediction. Does the model look at a photo of dog and say it's a dog or a cat?

These challenges of data analytics are not new. In the late 1960s and early 1970s, the Fair Housing Act, Fair Credit Reporting Act, and Equal Credit Opportunity Act were, in part, driven by similar issues.

As advisers and brokers incorporate these technologies in their services, the advice and recommendations they offer—whether or not based on AI—must be in the best interests of the clients and retail customer and not place their interests ahead of investors' interests.[17]

Rent Extractions and Conflicts

When communications, product offerings, and pricing can be narrowly targeted efficiently to each of us, producers are more able to find each individual's maximum willingness to pay a price or purchase a product. With such narrowcasting, there is a greater chance to shift consumer welfare to producers.

If the optimization function in the AI system is taking the interest of the platform into consideration as well as the interest of the customer, this can lead to conflicts of interest. In finance, conflicts may arise to the extent that advisers or brokers are optimizing to place their interests ahead of their investors' interests. That's why I've asked SEC staff to make recommendations for rule proposals for the Commission's consideration regarding how best to address such potential conflicts across the range of investor interactions.

Deception

Since antiquity, bad actors have found new ways to deceive the public. With AI, fraudsters have a new tool to exploit. They may try to do it in a narrowcasting way, zeroing in on our personal vulnerabilities. We used to all get similar spam. Now, communications can be efficiently individualized.

As a former faculty member, I also ponder how generative AI will lead to changes in teaching methods, academic integrity, and guarding against plagiarism.

Deception also can be broadcasted. Earlier this month, my communications director contacted me. "Lots of rumors on the internet today that you resigned. You'd tell me if you did, right?" he asked. The false rumor was apparently from AI-generated text on a website.[18]

More seriously, bad actors may seek to use AI to influence elections, the capital markets, or spook the public, potentially making Orson Welles 1938 "The War of the Worlds" radio broadcast look tame.[19]

Make no mistake, though, under the securities laws, fraud is fraud. The SEC is focused on identifying and prosecuting any form of fraud that might threaten investors, capital formation, or the markets more broadly.

Lastly, public companies making statements on AI opportunities and risks need to take care to ensure that material disclosures are accurate and don't deceive investors.

Macro: System-wide Risk

Now, to the macro. Just as with historic transformative times of moving to more automation of the farm, factory, and services, there will be macro challenges for society in general.

To the extent that generative AI and other AI technologies have the potential to automate a significant percentage of how workers spend their time,[20] it could lead to significant changes in jobs and the job market. AI, particularly given its transformative nature, also is shaping up as yet another way in which China and the U.S. compete economically and technologically.[21] There also are macro and geopolitical challenges from state actors and militaries' potential use of AI.[22] Further, though Ken Jennings of "Jeopardy!" fame may have been joking when he wrote it, there now are leading AI technologists who are quite serious and fearful of AI becoming our "overlords." [23]

While there are important policy debates around those macro issues, I'm going to focus on two other sets of macro issues.

We've seen in our economy how one or a small number of tech platforms can come to dominate a field. There are a number of reasons we may see the same with AI platforms.

Given that today's AI relies on an insatiable demand for data and computational power, there can be economies of scale and data network effects at play. We've already seen companies, both incumbents and startups, relying on base or foundation AI models and building applications on top of them.[24]

The downstream applications' use of the base models provides those models with greater data. The base models can be improved in the process, refining their parameters based upon the downstream applications' uses. This not only will enhance the base models but give them greater competitive advantage, adding to the possibility of having but one or a small number of platforms.[25]

I believe that multimodal AI systems, particularly those including video, are likely to further augment this trend, given the need for even more data and computation.

Once again, this raises a host of issues that are not new to AI but may be accentuated by it.

Privacy, Intellectual Property, and Rent Extractions

The debate around privacy and AI initially had been about who owns an individual's data. Are individuals in control of their own data? There also have been important debates about the use of remote biometric identification and

surveillance.[26]

Today's AI, though, reveals related issues around privacy and intellectual property rights that is not just about any one individual. Through the data being collected on each of us, we're all helping train the parameters of AI models. Further, if base models are training off of downstream applications, they are able to gain significant value.

So, whose data is it?

This debate is playing out right now in Hollywood,[27] with software developers,[28] and with social media companies.[29]

This raises questions about economic rents at the macro level as well. The platforms may garner greater economic rents if they become dominant.

For the SEC, the challenge here is to promote competitive, efficient markets in the face of what could be dominant base layers at the center of the capital markets. I believe we closely have to assess this so that we can continue to promote competition, transparency, and fair access to markets.

Financial Stability

The possibility of one or even a small number of AI platforms dominating raises issues with regard to financial stability. While at MIT, Lily Bailey and I wrote a paper about some of these issues called "Deep Learning and Financial Stability." [30] The recent advances in generative AI models make these challenges more likely.

AI may heighten financial fragility as it could promote herding with individual actors making similar decisions because they are getting the same signal from a base model or data aggregator. This could encourage monocultures. It also could exacerbate the inherent network interconnectedness of the global financial system.

Thus, AI may play a central role in the after-action reports of a future financial crisis.

While current model risk management guidance—generally written prior to this new wave of data analytics—will need to be updated, it will not be sufficient. Model risk management tools, while lowering overall risk, primarily address firm-level, or so-called micro-prudential, risks. Many of the challenges to financial stability that AI may pose in the future, though, will require new thinking on system-wide or macro-prudential policy interventions.

Conclusion

We've come a long way from Newton contemplating an apple and the cosmos. He would marvel at the math, data, and computational power of our times. He'd probably be intrigued by the fact that AI is helping refine data and detect gravitational waves [31] as well as its use in everyday life.

We at the SEC are technology neutral. Just like with the treatment of calculus—we focus on the outcomes, rather than the tool itself. Securities laws, though, may be implicated depending upon how AI technology is used. Within our current authorities, we're focused on protecting against both the micro and macro challenges that I've discussed.

While recognizing the challenges, we at the SEC also could benefit from staff making greater use of AI in their market surveillance, disclosure review, exams, enforcement, and economic analysis.

I think AI is going to continue significantly transforming science, technology, and commerce. I suspect Newton, though, would agree that it is important to focus on the micro and macro challenges of AI. Given that we're dealing with automation of human intelligence, the gravity of these challenges is real.

Lastly, a disclaimer. Though these remarks likely will be analyzed by AI models, they are not the product of any generative AI. They're all human.

[1] See Thomas Levenson, "The Truth About Isaac Newton's Productive Plague" (April 6, 2020), *available at* <https://www.newyorker.com/culture/cultural-comment/the-truth-about-isaac-newtons-productive-plague>.

[2] See Sebastian Castro Ramos, "The Discovery of Calculus: Leibniz vs. Newton" *available at* <https://stmuscholars.org/the-discovery-of-calculus-leibniz-vs-newton/>.

[3] On March 14, 2023, OpenAI unveiled Chat GPT-4, Google said chatbots would be integrated into Gmail and docs, and Anthropic announced its large language model, Claude. See James Vincent, "OpenAI announces GPT-4 — the next generation of its AI language model" (March 14, 2023), *available at* <https://www.theverge.com/2023/3/14/23638033/openai-gpt-4-chatgpt-multimodal-deep-learning>; and Jordan Novet, "Google is bringing A.I. chat to Gmail and Docs" (March 14, 2023), *available at* <https://www.cnn.com/2023/03/14/google-starts-testing-generative-ai-in-gmail-docs-coming-to-sheets.html>; and Emma Roth "Google-backed Anthropic launches Claude, an AI chatbot that's easier to talk to" (March 14, 2023), *available at* <https://www.theverge.com/2023/3/14/23640056/anthropic-ai-chatbot-claude-google-launch>.

[4] See Zeyi Yang, "Chinese tech giant Baidu just released its answer to ChatGPT" (March 16, 2023), *available at* <https://www.technologyreview.com/2023/03/16/1069919/baidu-ernie-bot-chatgpt-launch>.

[5] See Sebastian Castro Ramos, "The Discovery of Calculus: Leibniz vs. Newton" *available at* <https://stmuscholars.org/the-discovery-of-calculus-leibniz-vs-newton/>.

[6] See A.M. Turing, "Computing Machinery and Intelligence" (Oct. 1950), *available at* <https://phil415.pbworks.com/f/TuringComputing.pdf>.

[7] See Bharath Krishnamurthy, "The Evolution of Chess AI" (Oct. 22, 2022), *available at* <https://builtin.com/artificial-intelligence/chess-ai>.

[8] See Mac Schwerin, "America Forgot About IBM Watson. Is ChatGPT Next?" (May 5, 2023), *available at* <https://www.theatlantic.com/technology/archive/2023/05/ibm-watson-irrelevance-chatgpt-generative-ai-race/673965/>.

[9] See Bryan Walsh, "Looking back at Watson's 2011 'Jeopardy!' win" (Feb. 13, 2021), *available at* <https://www.axios.com/2021/02/13/ibm-watson-jeopardy-win-language-processing>

[10] See Cade Metz, "What the AI Behind AlphaGo Can Teach Us About Being Human" (May 19, 2016), *available at* <https://www.wired.com/2016/05/google-alpha-go-ai/>.

[11] See Aleks Farseev, "Is Bigger Better? Why the ChatGPT Vs. GPT-3 Vs. GPT-4 'Battle' Is Just a Family Chat" (Feb. 17, 2023), *available at* <https://www.forbes.com/sites/forbestechcouncil/2023/02/17/is-bigger-better-why-the-chatgpt-vs-gpt-3-vs-gpt-4-battle-is-just-a-family-chat/?sh=4153b0c05b65>.

[12] See Jacob Biba, "14 Types of IoT Sensors Available Today" (Feb. 21, 2023), *available at* <https://builtin.com/internet-things/iot-sensors>.

[13] Moore's Law was the prediction of Gordon Moore, co-founder of Intel, in 1965. He said computer power would double every two years while computer costs would be cut in half. See "Cramming More Components onto Circuits" *available at* <https://www.intel.com/content/www/us/en/history/virtual-vault/articles/moores-law.html>.

[14] For instance, Sen. Richard Nixon in 1952 gave a televised address to 60 million people, the largest audience to that point, advocating to stay on the ticket that year. See Richard Nixon Foundation, "How 'Checkers' Changed the Game of Television" (Sept. 23, 2016), *available at* <https://www.nixonfoundation.org/2016/09/how-checkers-changed-the-game-of-television/>.

[15] The EU AI Act's approach is to include heightened requirements for using AI to allocate many such resources. See European Parliament, "EU AI Act: first regulation on artificial intelligence" (June 8, 2023), *available at*

<https://www.europarl.europa.eu/news/en/headlines/society/20230601STO93804/eu-ai-act-first-regulation-on-artificial-intelligence>.

[16] A number of these issues have been written about by many others. In 2019, the Organization for Economic Cooperation and Development issued recommendations on AI. Last year, the White House released a Blueprint for an AI Bill of Rights. This year, the National Institute of Standards and Technology published an AI Risk Management Framework. In Europe, the AI Act, among other things, touches upon these issues. See Organization for Economic Cooperation and Development, "Recommendation of the Council on Artificial Intelligence" (May 21, 2019), *available at* <https://legalinstruments.oecd.org/en/instruments/OECD-LEGAL-0449>; and The White House, "Blueprint for an AI Bill of Rights" (October 2022), *available at* <https://www.whitehouse.gov/ostp/ai-bill-of-rights/>; and National Institute of Standards and Technology, "AI Risk Management Framework" (January 26, 2023), *available at* <https://www.nist.gov/itl/ai-risk-management-framework>; and European Parliament, "EU AI Act: first regulation on artificial intelligence" (June 8, 2023), *available at* <https://www.europarl.europa.eu/news/en/headlines/society/20230601STO93804/eu-ai-act-first-regulation-on-artificial-intelligence>.

[17] See Regulation Best Interest: The Broker-Dealer Standard of Conduct, Exchange Act Release No. 86031, 84 FR 33318 (June 5, 2019) ("Reg BI Adopting Release"); Commission Interpretation Regarding Standard of Conduct for Investment Advisers, Investment Advisers Act Release No. 5248, 84 FR 33669 (June 5, 2019), ("Fiduciary Interpretation").

[18] See Tom Mitchelhill, "AI-generated fake news sparks rumors of Gary Gensler's resignation" (July 3, 2023), *available at* <https://cointelegraph.com/news/gary-gensler-resigns-ai-generated-fake-news>.

[19] See Brad Schwartz, "The Infamous 'War of the Worlds' Radio Broadcast Was a Magnificent Fluke" (May 6, 2015), *available at* <https://www.smithsonianmag.com/history/infamous-war-worlds-radio-broadcast-was-magnificent-fluke-180955180/>.

[20] See "The economic potential of generative AI: The next productivity frontier" (June 14, 2023), *available at* <https://www.mckinsey.com/capabilities/mckinsey-digital/our-insights/the-economic-potential-of-generative-ai-the-next-productivity-frontier#introduction>.

[21] See Derek Cai and Annabelle Liang, "ChatGPT: Can China overtake the US in the AI marathon?" (May 2023), *available at* <https://www.bbc.com/news/business-65034773>.

[22] See David E. Sanger, "The Next Fear on A.I.: Hollywood's Killer Robots Become the Military's Tools" (May 5, 2023) *available at* <https://www.nytimes.com/2023/05/05/us/politics/ai-military-war-nuclear-weapons-russia-china.html>.

[23] See David Brooks, "Human Beings Are Soon Going to Be Eclipsed" (July 13, 2023), *available at* <https://www.nytimes.com/2023/07/13/opinion/ai-chatgpt-consciousness-hofstadter.html>; and Cade Metz, "'The Godfather of A.I.' Leaves Google and Warns of Danger Ahead" (May 4, 2023), *available at* <https://www.nytimes.com/2023/05/01/technology/ai-google-chatbot-engineer-quits-hinton.html>.

[24] See Sarah H. Cen, Aspen Hopkins, et. al., "AI Supply chains (and why they matter)" (April 3, 2023), *available at* <https://aipolicy.substack.com/p/supply-chains-2>.

[25] See Aspen Hopkins, Andrew Ilyas, et. al., "Who will provide AI to the world?" (May 1, 2023), *available at* <https://aipolicy.substack.com/p/supply-chains-3>.

[26] See European Parliament, "AI Act: a step closer to the first rules on Artificial Intelligence" (May 11, 2023), *available at* <https://www.europarl.europa.eu/news/en/press-room/20230505IPR84904/ai-act-a-step-closer-to-the-first-rules-on-artificial-intelligence>.

[27] See Andrew Webster, "Actors say Hollywood studios want their AI replicas – for free, forever" (July 13, 2023), *available at* <https://www.theverge.com/2023/7/13/23794224/sag-aftra-actors-strike-ai-image-rights>.

[28] See Mana Ghaemmaghami and Stuart Levi, "Ruling on Motion To Dismiss Sheds Light on Intellectual Property Issues in Artificial Intelligence" (May 24, 2023), *available at* <https://www.jdsupra.com/legalnews/ruling-on-motion-to-dismiss-sheds-light-6984451/>.

[29] See Wes Davis and Richard Lawler, "Elon Musk blames data scraping by AI startups for his new paywalls on reading tweets" (July 1, 2023), *available at* <https://www.theverge.com/2023/7/1/23781198/twitter-daily-reading-limit-elon-musk-verified-paywall>; and Mike Isaac, "Reddit Wants to Get Paid for Helping to Teach Big A.I. Systems" (April 18, 2023), *available at* <https://www.nytimes.com/2023/04/18/technology/reddit-ai-openai-google.html>.

[30] See Gary Gensler and Lily Bailey, "Deep Learning and Financial Stability" (Nov. 13, 2020), *available at* https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3723132.

[31] See Jared Sagoff, "Scientists use artificial intelligence to detect gravitational waves" (July 6, 2021), *available at* <https://www.anl.gov/article/scientists-use-artificial-intelligence-to-detect-gravitational-waves>.