andreessen horowitz

August 12, 2024

BY ELECTRONIC SUBMISSION

Director Moses Kim,
Office of Financial Institutions Policy
U.S. Department of Treasury
ATTN: Uses, Opportunities, and Risks of Artificial Intelligence in the Financial Services Sector
1500 Pennsylvania Avenue, NW
Washington, D.C. 20220

USES, OPPORTUNITIES, AND RISKS OF ARTIFICIAL INTELLIGENCE IN THE FINANCIAL SERVICES SECTOR

Dear Director Kim:

Andreessen Horowitz ("a16z") appreciates the opportunity to comment on the Department of Treasury's Request for Information on Uses, Opportunities, and Risks of Artificial Intelligence in the Financial Services Sector[1] ("RFI"). a16z was one of the earliest investors in many artificial intelligence ("AI")[2] companies and projects. Today, our investment vehicles hold positions in over 100 AI development firms, including as a leading investor in start-ups and open source AI developers. It has been the firm's long-held approach to work closely with such companies, using

[1] *Request for Information on Uses, Opportunities, and Risks of Artificial Intelligence in the Financial Services Sector, 89 Fed. Reg. 50048 (June 12, 2024).*

[2] We use the term "artificial intelligence" or "AI" in the same sense that Treasury does in the RFI. That is, it has the meaning in 15 U.S.C. 9401(3): a machine-based system that can, for a given set of human-defined objectives, make predictions, recommendations, or decisions influencing real or virtual environments. Artificial intelligence systems use machine and human--based inputs to perceive real and virtual environments; abstract such perceptions into models through analysis in an automated manner; and use model inference to formulate options for information or action.

our in-house technical understanding to support their business's growth. We also utilize myriad AI applications in our operations to enhance numerous work flows.

Our investments in AI include a broad and diverse range of projects, including:[3]

- A prospective broker-dealer building infrastructure that removes the complexities of embedding fixed income investments in any application;

- A wealth manager seeking to offer its customers access to private investment opportunities typically reserved for the most affluent and connected individuals, and combining these investment opportunities with services such as tax and estate planning;

- A company providing instant, risk-free ACH & card settlement to fintech and crypto companies, enabling them to stop payment fraud and increase conversion rates;

- An external vector database where developers can store and search relevant contextual data for large language model applications; and

- A platform for hosting and running machine-learning models in the cloud.

As one of the earliest and largest investors in many AI companies and projects, and as one of the largest venture capital investment advisers, we have keen insight into AI technology and we are well-positioned to comment on key aspects of the RFI, including:

- the challenges or barriers to access for small financial institutions seeking to use AI;

- the extent to which AI models and tools used by financial institutions are developed in-house, by third-parties, and/or based on open source code; and

[3] For further details on our involvement in AI, please see our firm's website "AI + a16z" at *https://a16z.com/ai/*.

- the actions necessary to promote responsible innovation and competition in the use of AI in financial services by the Treasury and by other regulators.

We begin our discussion by outlining several existing uses of AI in the financial markets, including in asset management and private funds, the capital markets, risk management and internal operations, trading, and in regulatory compliance. Beyond this descriptive introduction, however, we would like to take this opportunity to express significant concerns regarding the competitive landscape for AI, particularly for smaller firms and start-ups.

We believe there are major challenges around access to AI, particularly around cost and scalability. These challenges are, we strongly believe, best addressed by the continued development of AI itself, and through the sustained encouragement of development and innovation in this space. The most important task before regulators is how best to nurture AI development and innovation, without adding to regulatory costs for firms. We are particularly concerned about the risk of premature and inappropriate regulation, especially regulations that would disproportionately burden smaller firms.

One critical element towards lowering barriers around access to AI is the continued availability of open source AI — a significant resource that has the ability to erode the financial and other barriers that exist today. We therefore also provide an overview of how regulators can promote the conditions for AI innovation and entrepreneurship, both in the financial services sector and across AI more generally. We address these concerns and challenges in this response, but we begin by seeking to provide an overview of some prominent areas in which the use of AI has had a profound, often transformative impact on financial services.

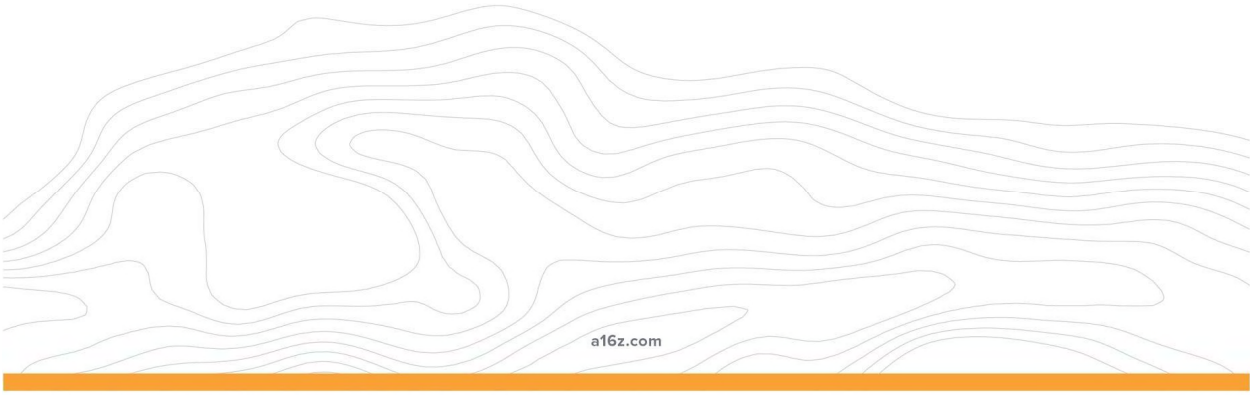

I. THE USE OF AI IN FINANCIAL SERVICES: AN OVERVIEW

The RFI seeks information on the opportunities and risks presented by AI developments and applications in financial services, including potential obstacles to responsible use of AI, the effect on impacted entities through use of AI by financial institutions, and recommendations for enhancements to legislative, regulatory, and supervisory frameworks applicable to AI in financial services. We therefore consider below some of the prominent ways in which AI is increasingly used in the financial services sector. We focus, as illustrative examples, on asset management and private funds, capital markets, trading, regulatory technology, and risk management.

This discussion is not intended to be exhaustive. Rather, it intends to show that the use of AI in financial services is widespread and well-accepted and has had an overwhelmingly positive impact on the sector in terms of efficiency and resource allocation. It also illustrates three other general principles that we hope regulators will adopt as they continue to engage with AI. First, that AI development occurs organically across a large range of firms, each reacting to perceived needs within its own industry and market. Second, that AI development has occurred, and is continuing to occur within the existing financial regulatory framework, without requiring further bespoke regulation. Third, that open source AI has been of fundamental importance in the ongoing development of AI-based tools. These principles underpin much of the AI-related development we have seen in asset management and private funds, capital markets, trading, regulatory technology, and risk management.

A. Asset Management and Private Funds

i. Asset Management, Generally

Asset managers have been using AI for a significant period of time now, particularly to generate investment strategies and to guide portfolio allocation and optimization.[4] In recent times, deep learning models, including neural networks, have been frequently used in asset management applications.[5] AI assists risk managers in validating and backtesting risk models.[6] AI approaches can also extract information more efficiently than traditional techniques from various sources of structured or unstructured data and generate more accurate forecasts of bankruptcy and credit risk, market volatility, macroeconomic trends, and financial crises.[7]

The use of AI can drive down the costs of active management, reduce management fees, and can therefore open up participation in asset classes typically reserved for institutional investors. For example, robo-advisor programs have democratized investment advisory services by making them cheaper and more accessible to retail investors, particularly younger investors.[8] One prominent example in this space is Betterment, a registered investment adviser and broker-dealer which provides investment, retirement and cash management services. Betterment uses AI in a number of ways in its robo-advisory operations, including to reduce taxes on transactions where machine learning algorithms select the specific tax consequences of the transactions.[9] Similarly, Wealthfront, an automated investment adviser, uses a software solution that is partly AI enabled to implement an investment approach based on Modern

[4] OECD, *OECD Business and Finance Outlook 2021: AI in Business and Finance*, pp. 39-40.
[5] Söhnke M. Bartram, Jürgen Branke, and Mehrshad Motahari, "Artificial Intelligence in Asset Management," CFA Institute Research Foundation, 2020, p. 6, *https://www.cfainstitute.org/-/media/documents/book/rf-lit-review/2020/rflr-artificial-intelligence-in-asset-management.ashx*.
[6] *Id.* at 3.
[7] *Id.*
[8] *Id. See also* Chaman L. Sabharwal "The Rise of Machine Learning and Robo-Advisors in Banking." *Journal of Banking Technology* 2: 28–43.
[9] Ilker Koksal, "How AI is Expanding the Applications of Robo Advisory", Forbes (Aug. 18, 2020), available at *https://www.forbes.com/sites/ilkerkoksal/2020/04/18/how-ai-is-expanding-the-applications-of-robo-advisory/*.

Portfolio Theory. In Wealthfront's model, AI is used for several objectives, from determining the optimal asset mix that maximizes returns for given levels of risk, to adjusting portfolios to respond to market movements, cash flows (like dividends), and changes in client risk tolerance.[10]

Artificial neural networks, decision trees, and support vector machines are all examples of AI technologies that are used in forecasting by asset managers. Cluster analysis assists in asset classification, while the use of evolutionary algorithms refines portfolio optimization in ways that even classical optimization algorithms cannot.[11] To the extent that these managers provide external research, AI tools can and are being developed to educate end investors and market participants.

ii Private Funds

Beyond asset management in general, advisers to private funds have unique use cases for AI. Private investment strategies pursued on behalf of their clients (funds) demand a more hands-on approach to diligence than standardized or passive asset management approaches. AI tools that allow such diligence to be carried out more effectively and efficiently help private funds distinguish themselves in an extremely competitive landscape, while also contributing to a virtuous cycle of technological innovation. In this way, the use and constant refinement of AI tools by private funds not only promotes technological development and innovation, it also helps private fund advisers fulfill their fiduciary duties to their clients.[12]

The internal use of AI by private funds allows for data to be consumed and analyzed at high speeds, thus driving down costs for the funds' limited partners. The ability to distill large

[10] Wealthfront Automated Investment White Paper, available at *https://research.wealthfront.com/whitepapers/investment-methodology/*.
[11] Bartram *et al*, Artificial Intelligence in Asset Management at 6, 18.
[12] *See*, *generally*, Marc Andreessen, The Techno-Optimist Manifesto (Oct. 16, 2023), available at *https://a16z.com/the-techno-optimist-manifesto/*.

amounts of data aids quantitative decision making, again benefiting limited partners and enabling advisers to better discharge their fiduciary duties to their clients. Consider for example, WorldQuant, a global quantitative asset management firm that incorporates AI in its investment process. WorldQuant develops and implements trading and investment strategies (including in liquid global asset markets) that may utilize AI, such as in connection with the development of trading signals derived from mathematical models in conjunction with historical data, seeking to predict future movements of various financial instruments. Or consider Numerai — a hedge fund in which an AI system chooses all the trades. Numerai is effectively an open source project built by several thousand anonymous data scientists competing to create the best trading algorithms.[13] These are examples of technological innovation that simultaneously advance the technological frontier and enable better adviser outcomes — all within the existing framework of the Investment Advisers' Act.[14]

Investment advisers frequently also go beyond developing their own platforms' use of AI. Advisers to venture capital funds often seek to bolster early-stage companies by providing them resources to pursue and develop technology. This helps the growth of portfolio companies as well as the broader tech ecosystem. For example, our firm develops AI tools for the benefit of both our internal portfolio companies as well as the broader public. While these tools have nothing to do with the investment process, we believe the development and provision of these tools foster collaboration, increase competition and technological agility, as well as increase public education and discourse around AI. As one example, and as we discuss elsewhere in this response,[15] our firm runs the Open Source AI Grant program, which seeks to provide start-ups with graphics processing unit computing resources that are crucial for generative AI.

[13] Cade Metz, "An AI Hedge fund Created a New Currency to Make Wall Street Work Like Open Source," *Wired* (Feb. 21, 2017), available at *https://www.wired.com/2017/02/ai-hedge-fund-created-new-currency-make-wall-street-work-likeopen-source/*.

[14] U.S. Senate Committee on Homeland Security and Governmental Affairs, "AI in the Real World: Hedge Funds' Use of Artificial Intelligence in Trading", (June 2024) at 33, available at *https://www.hsgac.senate.gov/wp-content/uploads/2024.06.11-Hedge-Fund-Use-of-AI-Report.pdf*.

[15] Please see Section III.B.

B. Capital Markets:

Modern capital markets and market participants also make extensive use of AI technologies — and in many cases, the same AI technologies that asset managers use. However, the subset of AI technologies used extensively in capital markets transactions have unique features that deserve specific mention.

Capital markets benefit from much of the same process automation that many other segments of the financial markets use. For example, in capital markets, as in other segments of the financial markets, AI enables or enhances processes through the trade lifecycle, such as customer onboarding, clearance and settlement, and regulatory compliance. There are also, however, AI technologies that specifically serve the capital markets, or are of particular use in the capital markets. For example, sentiment analysis through "natural language processing" can help market participants distinguish meaningful data points from background "noise" in a practice referred to as signal processing.[16]

Sentiment analysis is the analysis of financial sector news to forecast individual stock or market directions.[17] AI allows analysts and investors to process large volumes of both structured and unstructured data to determine trends and "sentiment" — for example, by surveilling social media, analyzing satellite imagery and GPS data, and reviewing other large data sets. Sentiment analysis enables efficient capital allocation and strategy formulation, with large language models ("LLMs") often proving particularly adept at forecasting market movements.[18]

[16] Alessio Azzutti, Wolf-Georg Ringe, and H. Siegfried Stiehl, "Machine Learning, Market Manipulation, and Collusion on Capital Markets: Why the 'Black Box, Matters," *University of Pennsylvania Journal of International Law,* vol. 43, no. 1 (2021), at 85, available at *https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3788872*.
[17] Gartner, "Sentiment Analysis," available at *https://www.gartner.com/en/finance/glossary/sentimentanalysis*.
[18] *Id.*

A second major area where AI has helped capital markets move toward greater efficiency and productivity is around credit underwriting. Predictive models have been used in lending operations for several decades now, for example through the creation and use of credit scores. More recently, however, lender firms have begun making greater use of machine learning measures to analyze large volumes of transaction data to analyze patterns and relationships that may not be visible in traditional lending models.[19] Such machine learning-based lending models may be more adaptive and "evolutionary" than traditional lending models, as they have been pioneered by nonbank firms, many of which are subject to far fewer restrictions than traditional banks.[20] However, banks too have made use of AI models for credit underwriting, both to identify useful variables and to refine traditional statistical models used in the credit underwriting process.[21]

In the capital markets and asset management contexts, AI appears to be creating positive feedback loops. For example, market participants who make use of AI, such as hedge funds or other trading firms, respond more swiftly to emerging data, particularly machine-readable data, and trade accordingly. There is also evidence to indicate that many companies that are actively traded by hedge funds and others are increasingly beginning to account for trading patterns and trading activity. Such companies are effectively "learning how to talk when a machine is listening", i.e., tailoring their disclosures in response to trading activity.[22] The net result of this ongoing symbiosis has been better market liquidity and price discovery for investors.

---

[19] FinRegLab, "Explainability and Fairness in Machine Learning for Credit Underwriting," at 7, available at *https://finreglab.org/wp-content/uploads/2023/12/FinRegLab_2023-12-07_Research-Report_Explainability-and-Fairness-in-Machine-Learning-for-Credit-Underwriting_Policy-Analysis.pdf.*

[20] John Adams, "Buy Now/Pay Later Fintechs Lean on AI to Survive the Banking Crisis," *American Banker,* (Mar. 27, 2023), available at *https://www.americanbanker.com/payments/news/buy-now-pay-later-fintechs-lean-on-ai-to-survive-the-bankingcrisis*.

[21] *See, for example*, Federal Reserve Governor Christopher J. Waller, "Innovation and the Future of Finance," speech at the Cryptocurrency and the Future of Global Finance, (Apr. 20, 2023), available at *https://www.federalreserve.gov/newsevents/speech/waller20230420a.htm*.

[22] Sean Cao et al., *How to Talk When a Machine Is Listening? Corporate Disclosure in the Age of AI,* National Bureau of Economic Research, Working Paper no. 27950, (Oct. 2020), pp. 18-19, available at *https://www.nber.org/papers/w27950*.

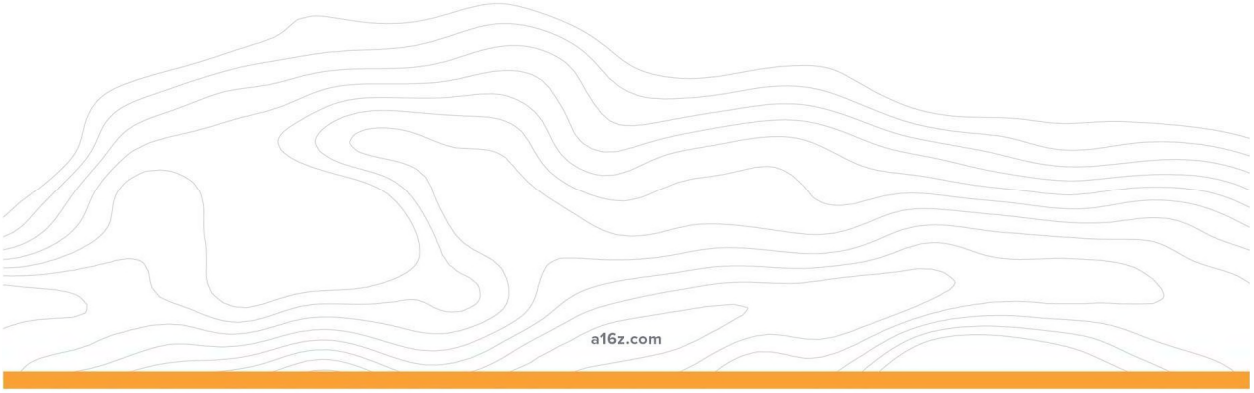

Notably, again, in the capital markets context, open source resources have been, and will continue to be of vital importance. This includes both open source financial data, such as large volumes of index or market performance data, as well as the development of open source AI tools such as Llama or Gemma. The continued availability of these resources will be vital in levelling the barriers to entry for firms that seek to access the capital markets.

C. Trading

Few areas in finance have been as significantly improved by the use of AI as trading. Many firms use AI at all stages of the trading process, from pre-trade analysis to trade execution to post-trade analysis.[23] Pre-trade analysis entails using data to analyze and forecast the performance, risks and costs of trading financial assets. In high frequency trading, such pre-trade analysis typically does not involve human intervention, and is almost entirely performed by algorithms, as is the actual execution of the trade.[24] Post-trade analysis involves the evaluation of trading outcomes, and this may involve some degree of human supervision.[25]

AI can shape trade execution models by actively learning from real market microstructure data to create optimal execution strategies. Algorithms can receive various combinations of input variables and can map these on to corresponding trading actions to optimize execution, and minimize transaction costs. Unlike traditional trading strategies, AI-based trading relies on data to determine market impacts and price movements. Such AI-based trading is therefore dynamic and flexible in its response to changing market conditions and the availability of new data. Conversely, however, and as we discuss below,[26] these trading models are difficult to develop and "train", and this is especially the case for larger portfolios.

[23] Bartram *et al*, Artificial Intelligence in Asset Management at 14-15.
[24] *Id.*
[25] *Id.*
[26] *Id.*

Notably, the development of these tools points to the centrality of technologists, rather than traders, in developing the modern trading environment. A case in point is Renaissance Technologies ("RenTec"), a firm that specializes in investment management strategies that employ mathematics and statistics in its trading systems. Today, RenTec manages approximately $69 billion in securities and has over $106 billion in assets under management,[27] and its founders have their backgrounds in mathematics and technology, rather than in trading. Drawing from a diverse pool of academic disciplines—ranging from mathematics and physics to computer science and statistics—RenTec's founders laid the foundation for a data-driven revolution that has helped shape the current financial landscape.[28]

Technologists developing the trading landscape are, of course, significantly advancing the tools on which markets rely. However, through their technological advances, they are also increasingly refining and giving shape to the legal duty of best execution, which binds both investment advisers and broker-dealers. Improved quantitative tools and transparency increasingly inform traders' choices, helping clarify where best execution lies. Conversely, and because of the expanded availability of trading data, investors are better equipped to discern flaws in execution, and to understand execution quality. In trading, as in many other areas, the growth of AI has served both to improve outcomes as well as to give shape and focus to what was, until very recently, a relatively amorphous legal duty.

D. RegTech

The use of AI in regulatory technology, or RegTech, is by now a relatively settled and stable field. The term RegTech may, however, span a broad range of functions, from the onboarding of customers (in which AI has played a role for over a decade), to newer use cases including

[27] AI in the Real World: Hedge Funds' Use of Artificial Intelligence in Trading, at 32.
[28] Lee Bailey, "The Rise of Renaissance Technologies: A Deep Dive into RenTec's Success", (July 26, 2023), available at *https://grizzlybulls.com/blog/rise-of-renaissance-technologies*.

fraud prevention, stress testing, and micro and macroprudential reporting.[29] Financial institutions are key users of RegTech, and the RegTech function can, depending upon the sensitivity, cost and complexity of the function, be outsourced to third-party providers. Below, we consider the role of AI in RegTech functions for a range of financial institutions.

While banks are prominent in deploying AI-based RegTech solutions, they are by no means the only financial institutions to do so. Broker-dealers, investment advisers, trading platforms, and currency exchanges, among many others, also routinely use AI-based RegTech solutions in their operations. For example, given the increase in the number of channels marketers use to reach customers and prospects since the COVID pandemic, RegTech has been vital in effective marketing reviews. Smart platforms can help compliance teams prioritize and manage their work, while providing visibility into planned, in-progress, and completed assignments. These platforms, built with machine learning and workflow management capabilities, provide an operational backbone that can help compliance leaders collaborate with multiple teams and work more efficiently.[30]

Beyond marketing reviews, AI-based RegTech has myriad and expanding use cases. AI systems are used for trade surveillance, to detect collusive behavior and price manipulation in the securities market.[31] Advanced analytics and technologies—and their monitoring and surveillance applications—can play an important role in combating unfair trading practices, misuse of privileged information, and other schemes aimed at creating unfair market conditions.[32] In fact, it is difficult to overstate the increasing role of AI in ensuring compliance more generally — not just in terms of securing compliance with applicable regulation, but even

[29] Tobias Adrian, "AI and Regtech," International Monetary Fund, (Oct. 29, 2021), available at *https://www.imf.org/en/News/Articles/2021/10/29/sp102921-ai-and-regtech*.
[30] Allison Lagosh, "Compliance concerns: Workloads, knowledge gaps, and separate systems" (Saifr), available at *https://saifr.ai/blog/compliance-concerns-workloads-knowledge-gaps-and-separate-systems*.
[31] AI and Regtech.
[32] Elia Alonso, Roy Ben-Hur, Niv Bodur (Deloitte), *RegTech's Hunt for Financial Market Abuses*, The Wall Street Journal (Apr. 29, 2021), available at *https://deloitte.wsj.com/cio/regtechs-hunt-for-financial-market-abuses-01619722931*.

in ensuring that firms adhere to their internal policies, written supervisory procedures, and codes of ethics.

AI-based RegTech can help significantly improve both compliance and business outcomes. For example, many banks are increasingly using AI analytics to improve balance sheet quality and create stress-testing models to meet regulatory requirements.[33] Conversely, bank regulators are also making use of AI through supervisory technology ("SupTech"), with some regulatory authorities, in the U.S. and abroad, using or actively considering the use of AI in the supervision process.[34] As the International Monetary Fund notes in a recent report, "AI in microprudential supervision could improve the assessment of risks, including credit and liquidity risks, and governance and risk culture in financial institutions, allowing supervisors to focus on risk analysis and on forward looking assessments."[35]

AI-based analysis of unstructured data and consumer behavior is also used extensively in compliance with anti-money laundering/countering the financing of terrorism ("AML/CFT") requirements. AI can, in particular, be very effective in reducing the number of false positives, which constitute the vast majority of AML/CFT alerts. This could allow financial institutions to devote more resources to cases that are likely to pose more serious concerns.[36] However, one significant challenge to financial institutions' use of AI arises from concerns regarding regulatory scrutiny and regulators second-guessing firms' compliance decisions, especially the accuracy of needing human decision intervention in the process, as opposed to using AI. While financial institutions must comply with AML/CFT requirements, regulators have a role to play in encouraging the regulated financial sector to use more technology towards effective compliance. Such an approach would encourage market efficiencies, improve efficacy and protect customers. By contrast, an environment where "regulation by enforcement" becomes

[33] AI and Regtech.
[34] *Id.*
[35] *Id.*
[36] *Id.*

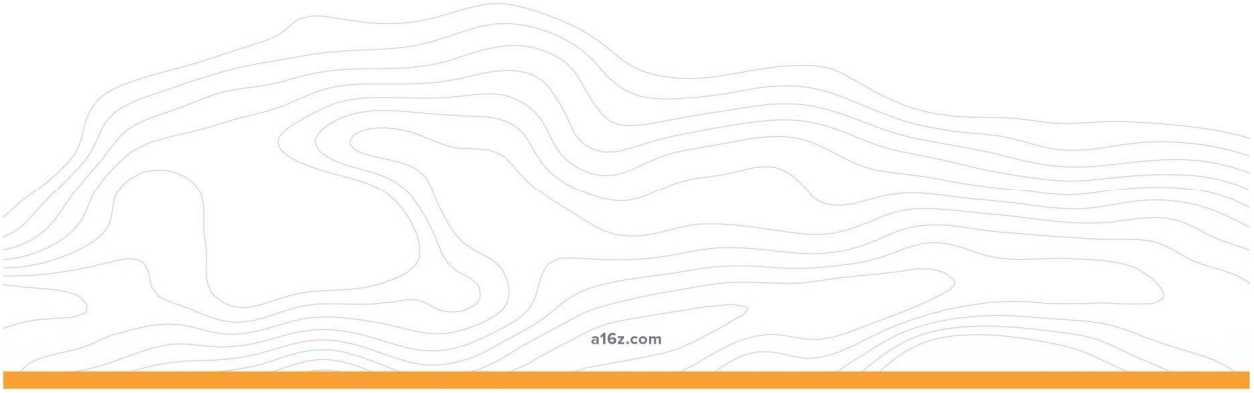

the norm will chill the adoption of technology by financial institutions at a pivotal point of adoption.

E. Risk Management and Internal Operations

AI tools are routinely used in portfolio risk management, around the management of both market risk and credit risk. Among other things, AI is used to manage market risk by using qualitative data for risk modeling, validating and backtesting risk models, and more accurately forecasting economic or financial variables. Artificial neural networks are particularly important in this context, because of their success in forecasting certain types of variables (such as interest rates), as well as their ability to formulate risk factors, including market volatility and systemic risks. Credit risk modeling, which we have touched on earlier in the context of credit underwriting, extensively uses AI tools, particularly artificial neural networks and "support vector machines", i.e., supervised algorithms that are typically used for classification purposes.

The purpose of such risk management is not confined to more accurate forecasting, though that is an important objective. Machine learning, and the use of support vector machines can also improve portfolio composition through better classification and segmentation.

We believe that the financial services sector is poised to use AI, including generative AI for five goals: (1) personalized consumer experiences, (2) cost-efficient operations, (3) better compliance, (4) improved risk management, and (5) dynamic forecasting and reporting.[37] Our discussion of the various AI use cases above — asset management, capital markets, trading, private funds, RegTech, and risk management — indicate to us that in financial services generally, the use of AI appears to create at least three types of broad benefits. First, the use of AI promotes democratization by increasing investors' access to various asset classes and management techniques. Second, the availability of technology has reduced reliance on certain

[37] Angela Strange et al, *Financial Services Will Embrace Generative AI Faster Than You Think*, (Apr. 19, 2023) available at *https://a16z.com/financial-services-will-embrace-generative-ai-faster-than-you-think/*.

types of key intermediaries such as market-leading investment advisers. Third, AI use has likely resulted in better resource allocation across the sector.

Yet, these benefits are not equally distributed across the financial services sector, and this unequal distribution is largely attributable to many of the barriers around access to AI technologies. In the next section of our response, we turn to an examination of the principal barriers to the use of AI, particularly by small businesses.

II BARRIERS TO AI USE BY SMALL BUSINESSES

The RFI asks about the challenges or barriers to access for small financial institutions seeking to use AI.[38] These barriers exist, and are very significant. The significance of these barriers is enhanced by the fact that AI integration is a notable trend among small businesses — one survey showed that 57% of small business owners and executives were incorporating AI to some extent in their operations, while nearly a quarter stated that they interacted with AI "very often" within their companies.[39] Below, we identify three major barriers to AI startups and other small AI firms — the cost of compute, the challenges of AI integration, and the uncertainties associated with data regulation.

A. The Cost of Compute

Unsurprisingly, cost is the single largest barrier to AI adoption by small businesses. One major deterrent, particularly for firms that train AI models, is the cost of "compute"[40] (i.e. the computational resources required for AI systems to perform tasks, such as processing data, training machine learning models, and making predictions and unique proprietary large

[38] *See* Question 4, 89 Fed. Reg. at 50053.
[39] Michelle Kumar & Justis Antonioli, *Small Businesses Matter; Navigating the AI Frontier*, (Apr. 29, 2024), available at *https://bipartisanpolicy.org/report/small-businesses-matter-navigating-the-ai-frontier/*.
[40] Nur Ahmed & Muntasir Wahed*, The De-democratization of AI: Deep Learning and the Compute Divide in Artificial Intelligence Research*, available at *https://doi.org/10.48550/arXiv.2010.15581*.

datasets.) Demand for compute far outstrips supply, and access to compute resources—at the lowest total cost—has become a determining factor for the success of AI companies.[41] In our own experience, many companies spend more than 80% of their total capital raised on compute resources.[42] AI startup financing rounds are large, and AI valuations have been on the rise[43] — even so, AI firms tend to overwhelmingly spend much of the money they raise on compute. And even after such bloated expenditures, it is very difficult for startups to produce models that could match models put out by much larger firms such as Google or OpenAI.

Also conspicuous by their absence are state-of-the art models released by academic institutions. To us, it is notable that even the largest and best resourced academic institutions, many of which boast very large endowments and strong sources of funding, cannot produce models that can effectively compete with the models put out by the largest firms, such as Google. This absence, again, speaks to the intimidatingly high cost of compute, a factor with which every AI firm that seeks to train a model must contend.

AI infrastructure is expensive, and looks set to remain that way in the medium term. That does not mean, however, that these expenses are necessarily an insurmountable barrier. We first provide an overview of some of the reasons that drive the cost of compute, and then discuss some measures that companies, particularly start-ups and small businesses, can take to mitigate these costs.

[41] In our view, supply may be starting to open up, but training AI models remains fundamentally much more expensive than building normal software.
[42] Guido Appenzeller, Matt Bornstein & Martin Casado, Navigating the High Cost of AI Compute, (Apr. 27, 2023), available at *https://a16z.com/navigating-the-high-cost-of-ai-compute/.*
[43] Consider, for example, Mistral, a French AI startup, that recently raised $645 million at a $ 6 billion valuation. *See* Arjun Kharpal, "Microsoft-backed Mistral AI raises $645 million at a $6 billion valuation", (June 12, 2024), available at *https://www.cnbc.com/2024/06/12/mistral-ai-raises-645-million-at-a-6-billion-valuation.html*.

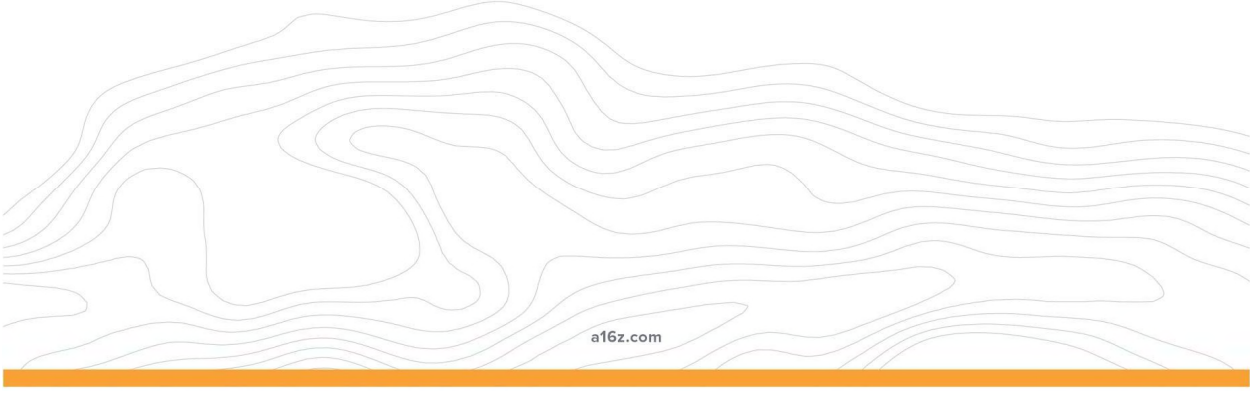

i. The Drivers of the Cost of Compute

There is a wide variety of generative AI models, and inference and training costs depend on the size and type of the model. AI infrastructure is expensive because the underlying algorithmic problems are computationally extremely difficult. The algorithmic complexity of sorting a database table with a million entries is insignificant compared with the complexity of generating a single word using standard AI means. As a result, AI users have an incentive to pick the smallest possible model that will satisfy the use case.[44] Generative AI requires massive investments in AI infrastructure today. There is no reason to believe that this will change in the near future.

Training most of today's state-of-the-art AI models involves manually cleaning and labeling large datasets. This process is laborious, expensive, and among the biggest barriers to more widespread adoption of AI. Even after a model is deployed, to maintain accuracy, new training data needs to be continually captured, labeled, and fed back into the system. Although techniques like drift detection and active learning can reduce the burden, our experience, based on a study conducted approximately four years ago, indicates that many companies may spend up to 10-15% of revenue on this process.[45] We have not recently reviewed the percentage of revenue spent on cleaning and labeling data sets, but we have no reason to believe the figure is any smaller.

The costs of developing and deploying AI systems are further increased by the continuing need for human intervention. For many tasks, especially those requiring greater cognitive reasoning, humans are often plugged into AI systems in real time. Social media companies, for example, employ thousands of human reviewers to augment AI-based moderation systems. Many autonomous vehicle systems include remote human operators, and most AI-based medical devices interface with physicians as joint decision makers. A number of AI companies,

[44] Navigating the High Cost of AI Compute.

[45] Martin Casado & Matt Bornstein, *The New Business of AI (and How It's Different From Traditional Software)*, Feb. 16, 2020, available at *https://a16z.com/the-new-business-of-ai-and-how-its-different-from-traditional-software/*.

including start-ups, that planned to sell pure software products are increasingly bringing services capability in-house and booking the associated costs.[46]

The need for human intervention will likely decline as the performance of AI models improves. It is unlikely, however, that humans will be entirely eliminated from these processes. Many problems – like self-driving cars – are too complex to be fully automated with current AI technology. Issues of safety, fairness, and trust also demand meaningful human oversight – a fact that many developing regulatory regimes are increasingly seeking to enshrine into law.[47]

ii. Mitigating Measures for Small Businesses

Small businesses and start-ups can take certain decisions in order to minimize the costs of building AI infrastructure. First, such businesses do not need to build their own AI infrastructure, they can simply contract with external providers. Hosted model services like OpenAI or Hugging Face or Replicate allow business owners for suitable products without the need to manage the underlying infrastructure or models. Pricing for these services is consumption-based, so contracting for them is often cheaper than building and running proprietary infrastructure.[48]

Second, and relatedly, start-ups and smaller businesses may choose to locate their AI infrastructure in the cloud, instead of building their own data centers, for example. Using the cloud will typically entail lower up-front costs, the ability to scale up and down, and regional availability — these are usually compelling considerations for both start-ups and larger companies. Note, however, that even using cloud providers can be very expensive, and may sometimes be infeasible. Larger public cloud providers typically charge a premium based on brand reputation, proven reliability, and the need to manage a wide range of workloads. Smaller specialty AI providers may offer lower prices, either by running purpose-built data

[46] *Id.*
[47] *Id.*
[48] *Id.*

centers or by arbitraging other clouds. (Consider, for example, Coreweave or Lambda Labs – each a purpose-built cloud provider that offers access to a broad range of compute solutions for smaller firms.) Larger buyers are typically able to negotiate prices directly with cloud providers, typically by committing to minimum expenditure and time commitments (usually between 1-3 years, in our experience). However, start-ups and smaller companies can sometimes get competitive pricing from specialty cloud providers without large spending commitments.[49]

Third, start-ups and smaller businesses can make significant savings by choosing the right graphics processing units ("GPUs") for their specific application. GPUs are foundational for today's generative AI era, largely because GPUs employ parallel processing, scale up to supercomputing heights, and because the GPU software stack for AI is broad and deep. As a result, GPUs perform technical calculations faster and with greater energy efficiency than central processing unit.[50]

As this discussion suggests, the costs of compute are considerable and can skew towards larger firms with significantly higher negotiating power. We also strongly believe that both compute capacity and the demand for compute capacity will continue to increase significantly — in other words, the shortage of compute and the shortage of GPUs that we see today will likely persist into the foreseeable future. Will the continued high cost of AI infrastructure make it impossible for small businesses and new entrants to catch up with well-funded incumbents? Not necessarily — while the cost of training an LLM may today seem prohibitive, the proliferation of open source models such as Alpaca and Stable Diffusion have shown that competitive advantages can quickly disappear. In other words, continued technological innovation can itself help erode the costs of technology, without requiring regulatory intervention around costs or access to products and services.

[49] *Id.*
[50] *Id.*

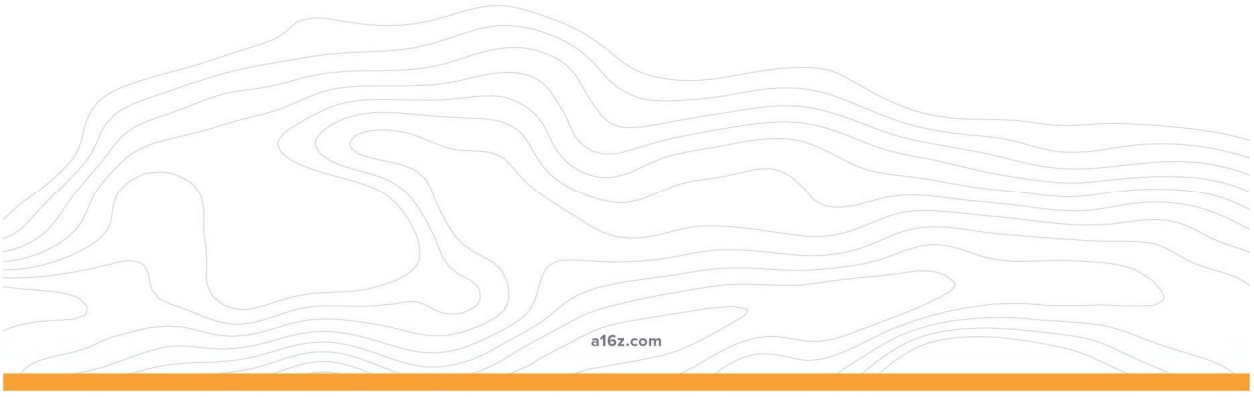

B. AI Accessibility and Integration

The costs of compute that we discuss above indicate that infrastructure vendors are likely to capture most AI-related revenues thus far. Companies developing applications, by contrast, may grow topline revenues quickly but often struggle with retention, product differentiation, and gross margins.[51] Model providers — i.e. the companies training generative AI models and applying them in new apps — likely create the most value in AI, but are unable to capture much by way of revenue.[52]

Generative AI relies on foundational research and engineering work done at a handful of entities such as Google, OpenAI, and Stability. Novel model architectures and training create the capabilities of current LLMs and image-generation models. Yet the revenue associated with these companies is still relatively small compared with the usage and publicity they have garnered. We acknowledge that this may be a temporary phenomenon, and given the huge usage of these models, large-scale revenues may not be far behind.[53] At this time, however, AI enterprise solutions remain largely bespoke. AI solutions are usually not easy to integrate with any given IT infrastructure and workflows. As we have previously discussed, AI model creation and deployment is a costly, time-consuming process, tailor-made to the needs of each business.[54]

For smaller businesses to benefit from AI, AI solution-providers must create solutions that can be integrated and used by multiple end users. While some AI vendors have started down this path, creating operating systems that allow businesses to integrate AI models into their information technology infrastructure, enterprise-scale AI adoption remains expensive and inaccessible for most businesses, and small businesses in particular. As with the cost of

[51] Appenzeller, Bornstein & Casado, *Navigating the High Cost of AI Compute*.
[52] *Id.*
[53] *Id.*
[54] Terence Tse, *Four Reasons Why Your Business Isn't Using AI*, California Management Review, (Jul. 22, 2019) available at https://cmr.berkeley.edu/2019/07/why-your-business-isnt-using-ai/.

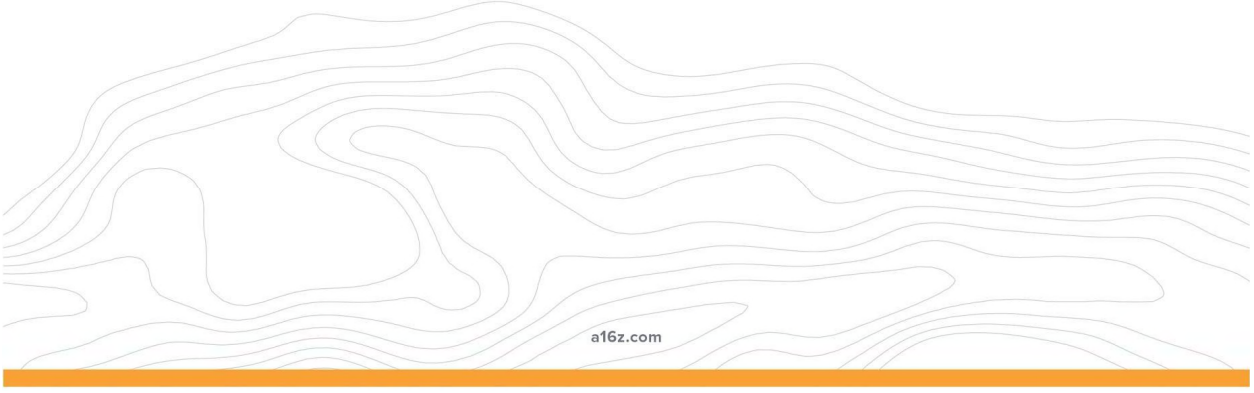

compute, however, there is no reason why this must remain indefinitely the case. What is required is the maintenance of competitive conditions in the AI ecosystem to allow for rapid innovation, and to provide opportunities and rewards for entrepreneurial founders. If these conditions can be created and maintained, a second wave of AI innovation may very well resolve the structural problems created by the first.

C. Data Access and Regulation

A third significant barrier for small businesses and start-ups is the proliferation of multiple regulatory regimes around data protection and data access. A survey from the U.S. Chamber of Commerce indicates that half of small businesses worry about having to comply with different state laws on privacy, AI, and technology in states where they are not headquartered, which can lead to increased litigation and compliance costs.[55] Nearly 70% of small businesses state that limiting access to data would harm their bottom line and ability to grow. Most small businesses use some form of technology platform for marketing, and 65% of them expressed the fear that losing their ability to reach consumers with targeted advertising would harm their business.[56]

Data regulations impact small businesses in multiple ways. The passage of data regulations by the various states create multiple, varying and occasionally contradictory requirements for small businesses. This patchwork of regulations may also significantly add to the cost of data, restrict the use of data, and ultimately even force small businesses to limit the extent to which they access data. Intellectual property laws can add to this chilling effect. For example, copyright laws may apply in ways that cause AI product developers and users to unwittingly engage in breaches of copyright.[57]

---

[55] U.S. Chamber of Commerce Technology Engagement Center, *The Impact of Technology on U.S. Small Business 2023*, at 15, available at *https://www.uschamber.com/assets/documents/The-Impact-of-Technology-on-Small-Business-Report-2023-Edition.pdf*.

[56] *Id.*

[57] *See, for example*, Congressional Research Service, *Generative Artificial Intelligence and Copyright Law*, (Sept. 29, 2023), available at *https://crsreports.congress.gov/product/pdf/LSB/LSB10922*.

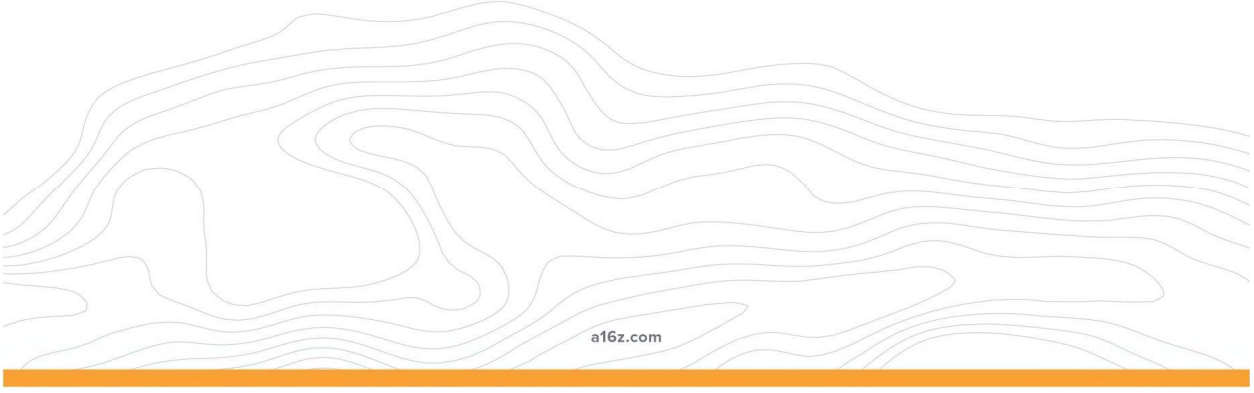

It is not our contention that data protection laws are unnecessary. Nor do we believe that courts should not enforce copyright laws. Our view is only that regulations around data and the use of content should be consistent and clear, and that the costs of compliance should not be excessive. Small businesses are disproportionately affected by compliance costs associated with a confusing and potentially conflicting state-by-state regime of privacy obligations. Small businesses are unlikely to be able to meet the costs required in order to use AI tools that use copyrighted content.

Congress can and should act swiftly to address these concerns by creating a single, easily implementable data standard and by legislating to permit the use of copyrighted material by AI tools and AI developers. The failure to act on these two issues, among others, will not only hinder the deployment of AI tools by small businesses, it will adversely affect the development of AI more broadly.

III. THE NECESSITY OF OPEN SOURCE AI

The RFI asks about the extent to which AI models and tools used by financial institutions are developed in-house, by third-parties, or based on open source code.[58] The RFI goes on to ask about the benefits and risks of using AI models and tools developed in-house, by third-parties, or based on open source code.

Our response, in brief, is that open source code is essential to the development of AI models and tools. We strongly believe that open source AI should be allowed to freely proliferate and to compete with both big AI companies and startups. Regulatory barriers to open source should be minimized so that, in the long run, the widespread availability of open source AI will minimize concentration risk, lower barriers to entry and promote financial stability.

[58] *See* Question 6, 89 Fed. Reg. at 50054.

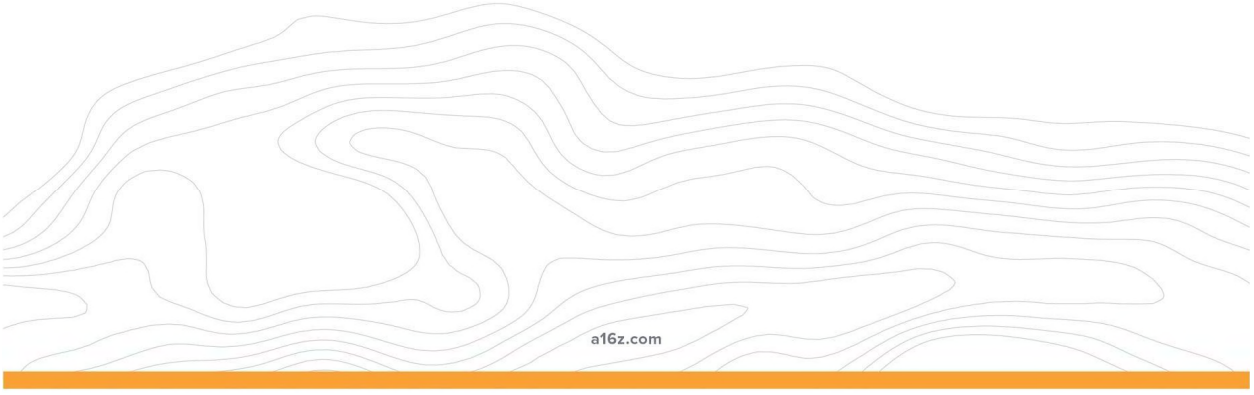

In the immediately foregoing sections of our response, we have outlined certain major economic and structural barriers to the deployment of AI, particularly by smaller firms and start-ups. The increased adoption of open source AI will be an essential ingredient towards removing these barriers — and, as we discuss further below, current trends appear to favor increased adoption of open source AI.

A. Current Trends Favor Open Source AI

We currently appear to be in the middle of a boom in the adoption of open source AI. This is a surprising reversal of the situation in 2023, when, by our estimate, up to 80-90% of respondents we surveyed indicated a preference for closed source AI. The majority of this share went to OpenAI.[59] Heading into 2024, however, 46% of our survey's respondents mentioned that they prefer or strongly prefer open source models.[60] In interviews, nearly 60% of AI leaders noted that they were interested in increasing open source usage or switching when fine-tuned open source models roughly matched performance of closed-source models. In 2024 and onwards, we expect a significant shift in enterprises' usage of AI towards open source. We understand some enterprises are expressly targeting up to 50% usage from open source models — this is up from approximately 20% in 2023.[61]

Understandably, cost is one of the major reasons for the appeal of open source AI. However, from our survey, it appears that control and the ability to customize (i.e., fine tune) open source AI for a given use are more important reasons for why many entrepreneurs chose to adopt open source data.[62] Customization is the key to integrating AI more effectively into

[59] Sarah Wang & Shangda Xu, 16 Changes to the Way Enterprises Are Building and Buying Generative AI, (Mar. 21, 2024) available at *https://a16z.com/generative-ai-enterprise-2024/*.
[60] *Id.*
[61] *Id.*
[62] *Id.*

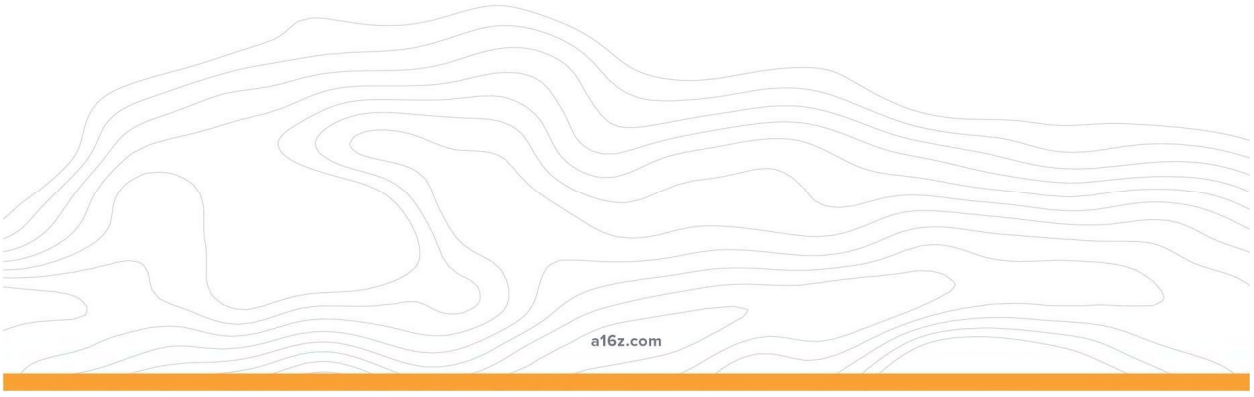

existing information systems. Greater control allows users to have greater security over proprietary data and to understand why models produce particular outputs.

The need for control, one of the major drivers of the increasing preference for open source AI, stems from enterprise data security concerns. Enterprises largely prefer not to share their proprietary data with closed-source model providers for regulatory or data security concerns — companies whose intellectual property is central to their business model are especially conservative in this respect. Some enterprises address this concern by hosting their own open source models, while others prefer models with virtual private cloud integrations. Enterprises are also very interested in customizing models, but with the rise of high-quality open source models, most are opting not to train their own LLMs, preferring instead to fine-tune an open source model for their specific needs.[63]

B. Open Source Models Continue to Face Significant Resource Constraints

The open source ecosystem is starting to develop, and we are beginning to see open source models that rival closed-source alternatives. Hundreds of small teams and individuals are also working to make these models more useful and accessible. These projects push the state of the art in open source AI and instruction-tuning base LLMs; removing censorship from LLM outputs; optimizing models for low-powered machines; building novel tooling for model inference; researching LLM security issues; and many others.

These positive trends notwithstanding, many open source projects lack the resources to sustain themselves or to maintain their efforts. And, as we have previously mentioned, resource constraints are especially acute in AI since fine-tuning or refining models requires significant GPU computing resources, especially as open source models get larger. To help close this resource gap, we run the a16z Open Source AI Grant program, through which we support a small group of open source developers through grant funding (not an investment),

[63] *Id.*

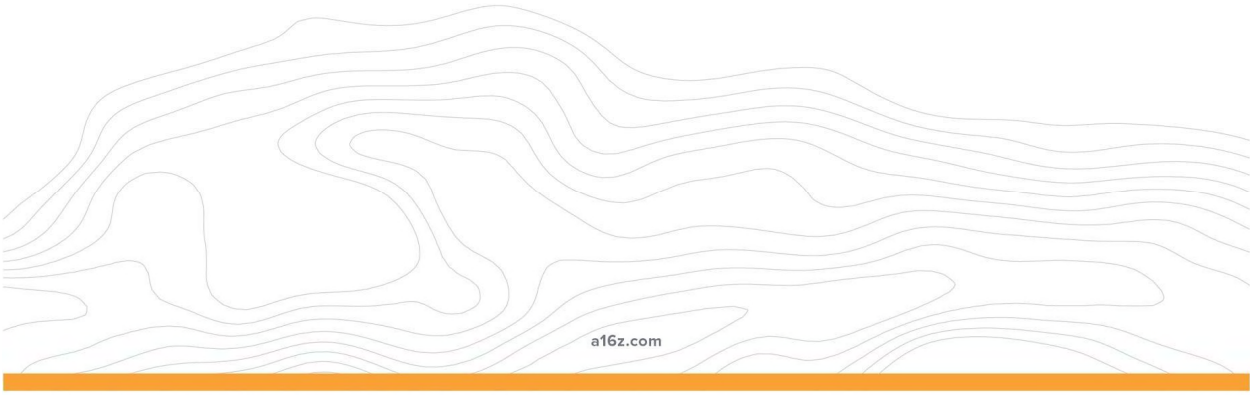

giving them the opportunity to continue their work without the pressure to generate financial returns. Much more, however, needs to be done to encourage the continued growth of open source AI.

C. Open Source AI is Critical to the Growth of AI

Many key developments in AI have been collective efforts, building upon the work of the larger AI community. To take one prominent example, OpenAI's ChatGPT application draws significantly on a 2017 paper published by Google Research which describes the required transformer architecture.[64] Hugging Face, a well known open source AI repository, has assisted AI researchers across areas ranging from medical research to climate change analysis.[65] One Harvard Business School analysis[66] shows that without open source software, businesses would have to spend 3.5 times more on software, leaving fewer resources for other needs.

Open source AI "accelerates innovation, reduces barriers to entry and democratizes access to cutting-edge technologies."[67] As one commentator notes, the openness of open source AI provides at least three distinct advantages:

- **Transparency:** Making the underlying code open to public review makes it easier to identify biases or errors in the training data or the model architecture.

[64] Ashish Vaswani et al, *Attention Is All You Need*, available at *https://arxiv.org/abs/1706.03762*.
[65] *See* Alistair King, *The Benefits of Open Source AI*, (May 11, 2024) in The Fieldston News, available at *https://fieldstonnews.com/home/2024/05/op-ed-the-benefits-of-open-source-ai/#:~:text=With%20open%20source%2C%20anyone%20can,tool%20of%20control%20and%20 oppressio*n.
[66] Manuel Hoffman, Frank Nagle, & Yanuo Zhou, *The Value of Open Source Software*, Harvard Business School Strategy Unit, Working Paper No. 24-038, available at *https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4693148*.
[67] King, *The Benefits of Open Source AI*.

- **Auditing:** Third parties can audit open source AI systems without requiring special access, authorization, or nondisclosure agreements. This promotes accountability and increases the chances of spotting errors.

- **Community Involvement:** Diverse communities representing very different views can be involved in open source projects from an early stage.

The development of open source software is instructive in illustrating the importance of open source AI. Open source software has given us some of our most important and widely used technologies, including operating systems (Linux), web browsers (Mozilla), and databases (MySQL).[68] Many of the vital infrastructural elements on which modern computing depends are open source, or arose from the open source software movement. Open source is also expanding to more areas of software. Traditionally, open source software developed around enterprise infrastructure, such as databases and operating systems (e.g. Linux or MySQL). Today, however, open source software is being actively developed for almost every industry – including, very prominently, for fintech.[69]

The seemingly unstoppable rise of open source software suggests that "its rise is due to a virtuous cycle of technology and business innovation." Open source is an optimal method for creating software "because it speeds product feedback and innovation, improves software reliability, scales support, drives adoption, and pools technical talent." When such technological innovation is paired with commercial innovation, the developer community expands, which spurs more technological innovation, further increasing the economic incentives for open source.[70]

[68] Peter Levine & Jennifer Li, *Open Source from Community to Commercialization* (Oct. 4, 2019) available at *https://a16z.com/open-source-from-community-to-commercialization/*.
[69] *Id.*
[70] *Id.*

At least in this context, past is prologue. The same virtuous technological and economic cycles that have propelled the steady rise of open source software will likely also determine the rise of open source AI. The principal task for regulators and governments should be to ensure that barriers to the availability of open source AI are minimized as far as possible. Technological and business communities will take care of the rest.

IV. PROMOTING RESPONSIBLE INNOVATION AND COMPETITION FOR AI IN FINANCIAL SERVICES

The RFI asks what actions are necessary to promote responsible innovation and competition for AI in financial services.[71] It asks what, if any, further actions are needed to protect impacted entities, including consumers, from potential risks and harms. We would respond with three broad principles. First, Treasury should work with other regulators to ensure the existence of vibrant, competitive environment for AI innovation and entrepreneurship. This, in our view, is the single most important policy objective required to ensure continued innovation and competition. Second, Treasury and other regulators should resist premature and inappropriate regulation of AI. Third, Treasury and other regulators should continue to participate in AI innovation, and to deploy and refine AI tools in their ongoing regulatory activities. We elaborate on these principles below.

A. ENSURE A VIBRANT, COMPETITIVE ENVIRONMENT FOR AI INNOVATION

Elsewhere in our response we have pointed to the fact that the market for AI, like all markets, currently has several economic and technological hurdles. Eliminating these barriers is essential for continued AI innovation, but that does not mean that active regulatory intervention is required around cost and availability. Technology is not static, and competitive advantages can erode swiftly — as the arc of the software industry over the last several decades show. Therefore, what is truly required is that regulators, including Treasury, ensure

[71] *See* Question 18, 89 Fed. Reg. at 50055.

that there is no interference with the conditions that encourage entrepreneurship and innovation.

To put this in somewhat more concrete terms, AI start-ups and small AI companies should be allowed to build AI as fast as they can. They should receive no assistance from regulators, but they should also not be required to compete against entities that receive regulatory advantages. In this process, many small AI businesses and AI start-ups will likely fail, but even in failing they will contribute towards a competitive economic environment for AI, that will benefit our economy more generally.[72]

Large AI companies benefit from scale and from their relative insulation from cost. Some also enjoy a dominant position as providers of AI infrastructure. Such large AI firms too should be allowed to build AI as fast as they can, but like small AI companies, they should receive no regulatory advantage when they do so. Perceptions of the nature of the AI markets should not lead regulators to allow large AI companies or to create cartels, or to permit large AI companies to build regulatory moats around their businesses.[73]

The use and availability of open source AI is a critical element of firms' ability to innovate and compete. Regulators should police and resist any attempts to mandate the use of closed source AI, or to impose regulatory barriers around the use of open source AI. Open source is competitive by definition and helps level the playing field at a time when cost is a key metric. The availability of open source AI serves as a resource for students and future innovators — as long as AI continues to be imperfect, as long as AI is expensive and cumbersome to deploy, open source AI must be protected and promoted by regulators and entrepreneurs alike.[74]

[72] Marc Andreessen, *Why AI Will Save the World*, (Jun. 6, 2023), available at *https://a16z.com/ai-will-save-the-world/*.
[73] *Id.*
[74] *Id.*

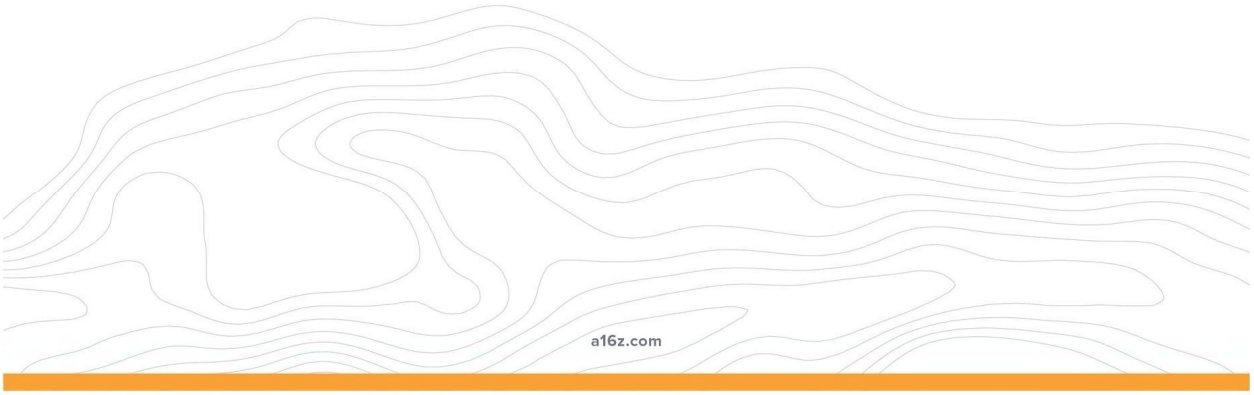

B. AVOIDING INAPPROPRIATE OR PREMATURE REGULATION

Perhaps even more important than the actions that regulators can take to encourage AI development are the actions that regulators should not take. Regulators should avoid succumbing to perceptions of "AI risk" that suggest that AI requires bespoke and immediate regulation. Regulators should avoid pressures to regulate AI so as to achieve social goals or to undertake social engineering. Regulators should not undertake AI regulation at a stage when its development is still nascent, and where regulation could have the consequence of entrenching dominant players, stifling innovation or strengthening barriers to entry. Unfortunately, there is already evidence that regulators, both in the United States and elsewhere, are failing to avoid these pitfalls.

AI is a set of technologies with transformative potential, but it is just that — one set of technologies among many. U.S. regulation has rarely been technology-specific. It is functional, rather than formal.[75] Among other things, the U.S. legal system has robust antitrust laws, civil and constitutional rights and protections against anti-discrimination, consumer protection laws, and laws that protect the integrity of our financial system. Each of these laws has been successfully used across a range of technologies and in a range of different situations. These laws are available to regulators, and in many cases to private litigants. There is no reason to believe that existing laws do not apply to AI, or are insufficient to regulate AI so that new AI-specific regulations are immediately required. The concerns that have been raised around the AI industry — allegations relating to bias, systemic risks, and abuses of dominance — are addressable by existing law. They do not require AI-specific regulation.

Relatedly, regulators should avoid seeking to regulate AI at a time when its development is barely beginning to take shape, and much data remains to be collected and studied. For example, a notable failure in this regard is the Securities and Exchange Commission's recently

[75] *See, for example*, Division of Corporation Finance, Division of Investment Management, and Division of Trading and Markets, *Statement on Digital Asset Securities Issuance and Trading*, (Nov. 16, 2018), available at *https://www.sec.gov/newsroom/speeches-statements/digital-asset-securities-issuance-trading*.

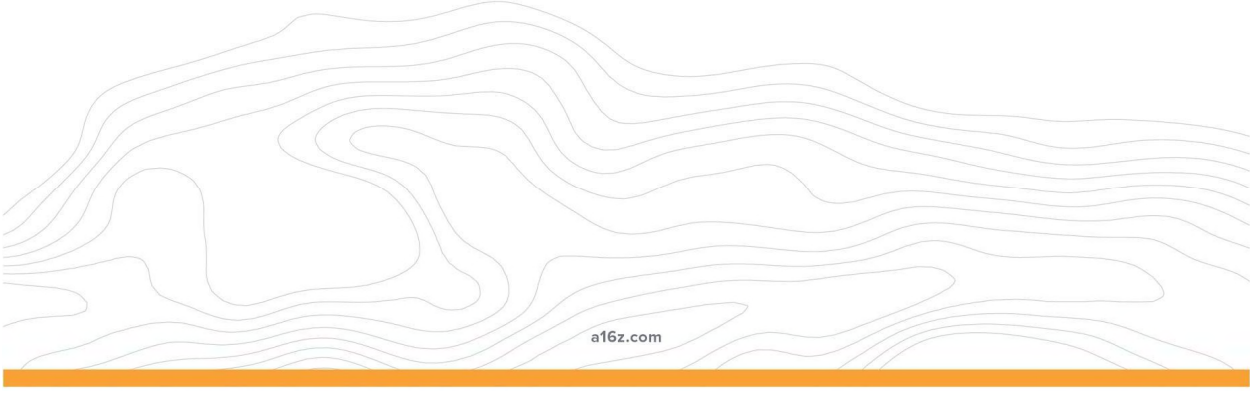

proposed Predictive Data Analytics rule.[76] The proposed rule would, among other things, have required broker-dealers and investment advisers who use AI or other predictive data tools to eliminate all conflicts of interest arising from the use of such tools, irrespective of the costs or feasibility of such elimination. Following significant criticism from a range of industry participants, including around the proposal's failures to collect data or evaluate costs,[77] SEC Chair Gensler has recently indicated that the rule may now be re-proposed.[78]

Instead of the SEC's attempt to regulate AI prematurely and in an overly broad manner, we would instead recommend an approach that, for the moment, primarily focuses on observing the developing AI industry, collecting data on participants and outcomes[79] and evaluating such data towards the development of a national policy. For the moment, as we discussed above, regulators should take the steps necessary to permit innovation to flourish, and to allow both large and small AI businesses to build as much AI as possible.

We would also suggest that regulators avoid the temptation to mandate AI to perform social engineering or to achieve social goals. In this regard, the fortunes of the AI industry in China should serve as a cautionary tale, particularly when it comes to the regulatory pressures placed by the Chinese state on generative AI.

---

[76] *Conflicts of Interest Associated with the Use of Predictive Data Analytics by Broker-Dealers and Investment Advisers*, 88 Fed. Reg. 53,960 (Aug. 9, 2023).

[77] *See, for example*, a16z's own comment on the proposed rule, available at *https://www.sec.gov/comments/s7-12-23/s71223-271620-654582.pdf*.

[78] Declan Harty*, SEC may open up AI, custody plans for further comment, Gensler says*, Politico, (May 16, 2024), available at *https://subscriber.politicopro.com/article/2024/05/sec-may-open-up-ai-custody-plans-for-further-comment-gensler-says-00158329*.

[79] For an example of this, consider the recent Executive Order on the Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence (Oct. 20, 2023), available at *https://www.whitehouse.gov/briefing-room/presidential-actions/2023/10/30/executive-order-on-the-safe-secure-and-trustworthy-development-and-use-of-artificial-intelligence/*.

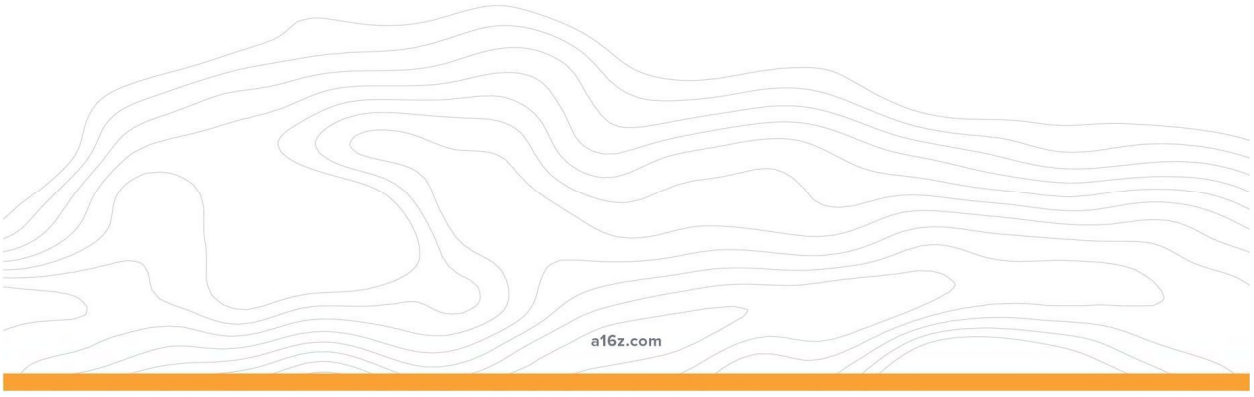

Today, most generative AI models in China need to obtain the approval of the Cyberspace Administration of China before being publicly available.[80] The Cyberspace Administration requires companies to prepare between 20,000 and 70,000 questions designed to test whether the models produce safe answers, according to one newspaper report.[81] Companies must also submit a data set of 5,000 to 10,000 questions that the model will decline to answer, many of which concern political topics or the Chinese Communist Party. The Chinese state's requirements for control also threaten to limit Chinese firms' access to training data. Start-ups struggle to find Chinese-language data for training AI systems. Less than 5% of the data in Common Crawl, a widely used open source database, is Chinese-language data.[82]

The United States is not China and our political and constitutional systems serve to protect individuals and firms from such governmental pressures. Nevertheless, the broader point is instructive — governments should not use AI to achieve social or political goals, irrespective of how well-intentioned those mandates may be.

C. Government and Regulators Should Deploy AI and Benefit From Its Development

Governments and governmental agencies, like individuals and companies, will make efficiency gains from the strategic deployment of AI. The rise of SupTech provides powerful examples of this. To take just one such example, the Federal Reserve System recently explored the implementation of Natural Language Processing with Lexical Expansion ("NLP-LEX") to improve operational efficiency and effectiveness.[83] NLP-LEX uses machine learning to analyze and expand financial text data. The Federal Reserve had found it labor-intensive and time-consuming to extract valuable insights from the large mass of financial data to which it had

[80] Liza Lin, *China Puts Power of State Behind AI—and Risks Strangling It*, The Wall Street Journal, (Jul. 16, 2024), available at *https://www.wsj.com/tech/china-puts-power-of-state-behind-aiand-risks-strangling-it-f045e11d*.
[81] *Id.*
[82] *Id.*
[83] Cambridge SupTech Lab, *State of SupTech Report 2023*, available at *https://lab.ccaf.io/wp-content/uploads/2024/03/Cambridge-State-of-SupTech-Report-2023.pdf*.

access. Deploying NLP-LEX allowed it to streamline its processes, and enhance its decision-making capabilities.[84]

Deploying AI and training staff to use AI will not only allow regulators to more efficiently perform their functions, they will likely allow regulators to regulate markets and activities more effectively. We have earlier noted that we find the concept of "AI risk" to be overbroad and misleading. AI developers are not reckless, in our view, and AI does not pose the kinds of vast systemic and social risks that some critics have falsely claimed.[85] However, if regulators do have concerns around AI use, or around the potential use of AI by bad actors, the best remedy would be for regulators to vigorously engage with the AI industry in each area of risk to use AI to maximize society's defensive capabilities.[86]

D. THE INTERNET'S LEGACY: LESSONS FOR AI REGULATION

In the early days of the internet, it was nearly impossible to foresee the profound impact it would have on every aspect of society, from financial services and healthcare to innovation and daily life. The internet has become so integral to our existence that imagining a world without it is now near-impossible. We view AI through a similar lens, recognizing its potential to revolutionize the world on a comparable scale. Just as with the internet, the transformative capabilities of AI are vast and largely unpredictable.

We must acknowledge that "we don't know what we don't know" regarding AI's future innovations and breakthroughs. Therefore, it is crucial to allow AI innovation to flourish without imposing overly burdensome regulations that could stifle progress or lead to unforeseen consequences. Historical examples such as Section 230 of the Communications Decency Act (1996), which allowed platforms like social media to flourish without being liable for user-

[84] *Id.*
[85] Marc Andreessen, *Why AI Will Save the World*.
[86] *Id.*

generated content[87], and the Digital Millennium Copyright Act of 1998 with its "safe harbor" provisions that helped balance the interests of copyright holders and online platforms[88], demonstrate how thoughtful regulation can enable innovation while protecting essential interests. Over-regulation, premature regulation, and inappropriate regulation each run the risk of not only hindering domestic advancements but also ceding global leadership to other nations with less stringent controls. It is vital for the United States to foster an environment where AI can develop robustly, ensuring that we remain at the forefront of this pivotal technological development.

V. OBSERVATIONS AND RECOMMENDATIONS

Last year, President Biden noted that "Responsible AI use has the potential to help solve urgent challenges while making our world more prosperous, productive, innovative, and secure."[89] We agree, but we are concerned by the fact that the growth and development of AI has been matched only by governmental zeal to regulate AI firms and AI tools, not just in the United States, but elsewhere in the world, including the European Union. The proliferation of new technologies invariably sparks panic, and AI has been no different in that regard.

We would also note that governmental actions that relax unnecessary and restrictive regulation often spark technological innovation, competition and progress. One example is telecommunications, where the overhaul of the industry in the 1990s reduced the regulatory burden on internet service providers, allowed different companies to enter the space and ultimately led to innovations like broadband and wide-spread use of mobile phones.[90] Competition in the markets for electric power was substantially increased by the Energy Policy Act of 1992, which also helped increase clean energy use and improve overall energy efficiency

[87] Kosseff, J. (2019). The Twenty-Six Words That Created the Internet. Cornell University Press
[88] Litman, J. (2001). Digital Copyright: Protecting Intellectual Property on the Internet. Prometheus Books
[89] Executive Order on the Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence.
[90] Stuart N. Brotman, "Was the 1996 Telecommunications Act Successful in Promoting Competition?", Brookings Institute, (Feb. 8, 2016), available at *https://www.brookings.edu/articles/was-the-1996-telecommunications-act-successful-in-promoting-competition/*.

in the United States. A similar push towards industry-appropriate regulation could generate a wave of entrepreneurial and technological activity around AI, yielding benefits that are at once wide-ranging and long-lasting.

We view the rise of AI with hope, not concern. And it is in that hope that we suggest the following observations:

- AI as a technology is still developing, but the use of AI across financial services is widespread, well-entrenched, and showing continued potential.

- The rise of AI in financial services has produced efficient decision making, greater accessibility across asset classes for investors, and a reduction in concentration risk.

- There continue to be significant barriers to the use of AI by small firms and start-ups, principally around cost and increasing data-related regulation. However, technological advances have the potential to chip away at cost barriers if AI entrepreneurs and innovators are allowed to flourish.

- The principal aim for Treasury and other regulators should therefore be to preserve and promote the conditions that would lead to innovation and competition in AI.

- The existence of open source AI and open source tools is perhaps the most critical element for continued AI innovation and broad-based AI entrepreneurship. Regulators should strongly resist moves to mandate closed source AI, and all attempts to narrow access to open source AI.

- Regulators should also resist pressures to engage in premature regulation of AI or to commandeer AI to achieve social or political goals. They should instead take

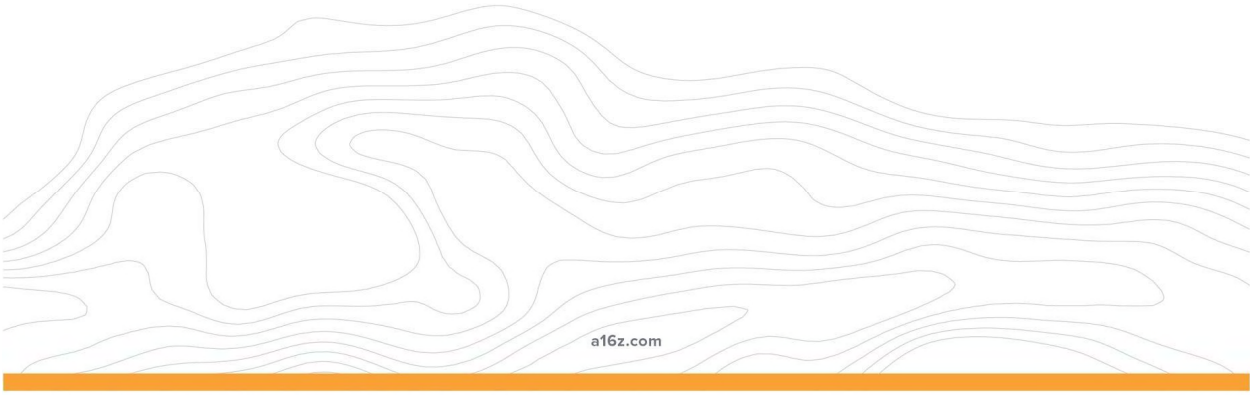

this opportunity to observe, collect data on, and evaluate the workings of the AI industry both in the U.S. and abroad.

- Regulators should also be participants in the AI revolution, engaging with the industry, training staff and deploying AI tools to make their own processes more effective and more efficient.

In the words of President Biden, "The rapid speed at which AI capabilities are advancing compels the United States to lead in this moment for the sake of our security, economy, and society."[91] We stand ready to work with Treasury towards that leadership.

[91] *Id.*

Respectfully submitted,
A.H. Capital Management, L.L.C.

By:

/s/ Jai Ramaswamy
Jai Ramaswamy
Chief Legal Officer

/s/ Scott Walker
Scott Walker
Chief Compliance Officer