DocuSign Envelope ID: 36A3FB47-22DA-4C7F-9F77-903EDD214C2D

andreessen. horowitz

March 27, 2024

BY ELECTRONIC SUBMISSION

National Telecommunications and Information Administration
1401 Constitution Ave.,
NW Washington, D.C. 20230

AH Capital Management, L.L.C.'s Response to the National Telecommunications and Information Administration's Request for Comment on Dual Use Foundation Artificial Intelligence Models with Widely Available Model Weights

Agency Docket Number: 240216-0052
Regulations.gov Docket Number: NTIA-2023-0009

AH Capital Management, L.L.C. ("a16z") welcomes the opportunity to provide comments in response to the National Telecommunications and Information Administration's ("NTIA") Request for Comment ("RFC") on Dual Use Foundation Models with Widely Available Model Weights ("Open Models").[1] As one of Silicon Valley's preeminent venture capital firms with over $35 billion in committed capital, a16z has been investing in artificial intelligence for many years, and today, we manage pooled vehicles which hold investments in nearly 100 AI development firms.[2] As one of the earliest and largest investors in many AI companies and projects, and as

[1] For the purposes of this response, we use the term "Open Models" to mean foundational artificial intelligence models (a) with open weights and training data; and that (b) are accessible by anyone. *See* Section 2, *infra*.
[2] For more detail on our involvement in AI, please see "AI + a16z" on our firm's website, at https://a16z.com/ai/.

DocuSign Envelope ID: 36A3FB47-22DA-4C7F-9F77-903EDD214C2D

one of the largest investment advisers in the private technology space, a16z is well-positioned to evaluate the potential impact and importance of open source models on advancements in AI. We are pleased that NTIA recognizes the importance of these foundation models and appreciate the opportunity to participate in this important discussion.

1. Introduction[3]

Recent advancements in artificial intelligence represent the most significant development in computer technology since the invention of the microchip. From these advancements, we are likely to see developments that measurably improve the human condition. Far from representing an existential threat to humanity, the development of foundation AI models represents a profound opportunity to augment the same human intelligence that has created the world we live in today. We can leverage such development to provide more people with even better outcomes, like the creation of new medicines, and improvements in the way we educate our children, make decisions, and generate creativity, to name a few.

We have a long history in this country of taking a hands-off approach to open source software, and have greatly benefited from the advancements it has enabled. Open source software provides much of the foundation of the internet, is extensively used by government agencies, and is widely considered to be more secure from malicious actors than proprietary software. Consequently, there is no inherent reason to prohibit the creation, dissemination, and use of Open Models. Nor is there reason to create a regulatory preference for models with weights and training data that are closed from public view ("Closed Models"). In addition, Open Models are also a powerful defensive tool to mitigate risk and prevent harm, whether caused by

[3] This comment letter responds to the RFC's questions in narrative format. Consistent with NTIA's instructions, we indicate where appropriate the number of the question to which we are responding.

DocuSign Envelope ID: 36A3FB47-22DA-4C7F-9F77-903EDD214C2D

Open Models, Closed Models, or other non-AI technologies.[4] Finally, Open Models already exist and are actively being distributed and used throughout the stream of commerce.[5] Therefore, as a practical matter, it is likely impossible to shift to a fully closed paradigm.

Like the internet and open source software movement before, Open Models promote innovation,[6] reduce barriers to entry, protect against bias,[7] and allow such models to leverage and benefit from the collective expertise of the broader artificial intelligence ("AI") community. To be sure, bad actors will always look to exploit, misuse, and abuse powerful tools—including open and closed foundational models—to cause harm. All tools have the potential for misuse, and Open Models are not inherently riskier than Closed Models. In fact, Open Models have the distinct advantage of allowing governments, regulators, and the public to evaluate them—

[4] Rowan Zellers et al., *Defending Against Neural Fake News*, NIPS'19: Proceedings of the 33rd International Conference on Neural Information Systems, 9054–9065, (Dec. 2019) https://dl.acm.org/doi/10.5555/3454287.3455099 ("The best existing fake news discriminators are, themselves, deep pretrained language models (73% accuracy) . . . However, we find that Grover, when used in a discriminative setting, performs even better at 92% accuracy. This finding represents an exciting opportunity for defense against neural fake news: the best models for generating neural disinformation are also the best models at detecting it.").

[5] Belle Lin, *Open-Source Companies Are Sharing Their AI Free. Can They Crack OpenAI's Dominance*?, Wall St. J. (Mar. 21, 2024), https://www.wsj.com/articles/open-source-companies-are-sharing-their-ai-free-can-they-crack-openais-dominance-26149e9c; Brian Andrus, *Open-Source AI: 9 Powerful Models You Need to Try*, DreamHost (Jan. 30, 2024), https://www.dreamhost.com/blog/open-source-ai/.

[6] W3techs, https://w3techs.com/technologies/comparison/os-linux,os-windows (indicating that Open source software has been the foundation of technological innovation since the creation of Linux, which, as of March 2024, is used by 41.6% of all websites with known operating systems); Sascha Brodsky, *Mistral AI's New Language Model Aims for Open Source Supremacy*, AI Business (Dec. 19, 2023), https://aibusiness.com/nlp/mistral-ai-s-new-language-model-aims-for-open-source-supremacy#close-modal (quoting computer science professor at Carnegie Mellon stating "Without the foundational building block being open-sourced, a lot of progress in the overall field of computer science would likely have been significantly slower.").

[7] Angela Wang & Olga Russakovsky, *Directional Bias Amplification*, 139 Proceedings of Machine Learning Research 10882–10893 (2021), https://proceedings.mlr.press/v139/wang21t/wang21t.pdf; Inioluwa Deborah Raji & Joy Buolamwini, *Actionable Auditing: Investigating the Impact of Publicly naming Biased Performance Results of Commercial AI Products*, AIES'19 Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society 429–435, https://dl.acm.org/doi/10.1145/3306618.3314244#.

DocuSign Envelope ID: 36A3FB47-22DA-4C7F-9F77-903EDD214C2D

including their capabilities, risks, and underlying training data—in a more transparent manner. It is the very openness of Open Models that enables the more effective identification, understanding, and mitigation of risk.

We are pleased to see that NTIA is proceeding thoughtfully and deliberately with regard to its approach to Open Models. We encourage NTIA to be wary of generalized claims about the risks of Open Models and calls to treat them differently from Closed Models, especially those made by AI companies seeking to insulate themselves from market competition. Open Models should be allowed to freely compete with both big AI companies and startups. We urge NTIA to adopt relevant definitions and rules that allow for the continued development of Open Models rather than a restrictive approach through which the government, and not technological and market developments, picks winners and losers in this important, emerging market.

2. Definitions[8]

We propose that NTIA adopt a definition of "open" that includes not just models and weights, but also training data. In order to maximize the benefits inherent in Open Models while leveraging collective experience to minimize harms, openness cannot be restricted to the models or the weights themselves. Some of the most important factors determining how foundation models work are the decisions made regarding the training data, such as what data to use, whether the model is "aligned" during training, what values are used to "align" the model, and whose interests are prioritized to mitigate harms during model deployment.

Open training data provides key insights not only into how the model was developed, but also into the model's capabilities, limitations, and risks,[9] including bias.[10] Access to training data allows users to identify, analyze, and contextualize the information on which the model is acting

[8] This section addresses RFC questions 1, 1.c, and 5.f.
[9] Elizabeth Seger et al. *Open-Sourcing Highly Capable Foundation Models: An evaluation of risks, benefits, and alternative methods for pursuing open-source objectives*, arxiv:2311.09227, Sept. 29, 2023, https://arxiv.org/abs/2311.09227.
[10] Wang, *supra* note 7; Raji, *supra* note 7.

DocuSign Envelope ID: 36A3FB47-22DA-4C7F-9F77-903EDD214C2D

and better understand the model's outputs. In evaluating the training data, users may discover gaps in information, errors, and algorithmic biases; including training data as a central component of the definition of "open" is essential to properly assess the benefits and risks presented by Open Models. In short, data set openness is a key risk mitigation tool.[11] Accordingly, NTIA's definition of "open" should incorporate the idea that open source foundation models are artificial intelligence models, with or without limited user restrictions, utilizing broadly accessible weights *and* training data.

We also propose that NTIA adopt a definition of "widely available" based on whether the model, weights, and training data are released in a manner that would allow access and use by any developer rather than a definition that relies on specific levels of usage. That is, a widely available Open Model is one with open weights and training data that is reasonably accessible to anyone, not one that is distributed to or used by a certain number of developers. Models should still be considered widely available even if they are released with certain access restrictions (*e.g.*, restricting access to individuals in countries subject to governmental sanctions) or terms prohibiting their use for illegal purposes. Finally, availability is not synonymous with use or distribution, and adopting such a definition risks conflating these two principles in a manner that would result in an unduly narrow definition. Rather, the degree to which a model actually is adopted or used is a feature of a free market wherein consumers choose the model that appeals to their preferences, and should not factor into whether the model is considered "widely available."

Our proposed definition differs from those proposed in other contexts by: (1) including training data within its scope; and (2) separating concepts such as staged releases, the availability of companion assets, and license terms from the definitions themselves. For example, a recent paper defines an open foundation model as one that (i) must provide weights-

[11] Aspen Institute, Global Cybersecurity Group, *Generative A.I. Regulation and Cybersecurity: A Global View of Policymaking*, at 6 (Jan. 2024), https://www.aspendigital.org/wp-content/uploads/2024/01/Aspen-Digital_Generative-AI-Regulation-and-Cybersecurity_January-2024.pdf ("Mandating data set openness, human oversight, and transparency in commercial generative models can reduce risks.").

DocuSign Envelope ID: 36A3FB47-22DA-4C7F-9F77-903EDD214C2D

level access; (ii) need not be accompanied by the open release of any other assets; (iii) must be widely available, though some restrictions on users may apply; (iv) need not be released in stages, and (v) may have use restrictions.[12] But the definition does not mention the role of open training data. President Biden's recent Executive Order on artificial intelligence similarly focuses heavily on the model weights of dual-use foundation models.[13] Training data is fundamental to the operation of foundation models. Understanding what data is being used and how it is being used to train models is critical to assessing and mitigating risks. Accordingly, our definition acknowledges the importance not only of access to the weights used within a model but also to the training data.

Influential organizations in the open source community are starting to recognize the importance of developing clear definitions for Open Models that include open training data and wide availability. For example, the Open Source Initiative ("OSI") is engaged in an ongoing project to develop a definition for Open Source AI. OSI defines Open Source AI as an AI System that allows the public to (1) use the system for any purpose and without having to ask for permission; (2) study how the system works and inspect its components; (3) modify the system for any purpose, including to change its output; and (4) share the system for others to use with or without modifications, for any purpose.[14] OSI recognizes that, for the public to be able to modify and use an Open Source AI System, there must be public access to the training data, the code, and the model parameters, including the weights.[15] We encourage NTIA to recognize the

[12] Sayash Kapoor et al., *On the Societal Impact of Open Foundation Models*, Feb. 27, 2024, https://crfm.stanford.edu/open-fms/paper.pdf.
[13] Exec. Order No. 14110, *Safe, Secure, and Trustworthy Development and Used of Artificial Intelligence* (Oct. 30, 2023), 88 FR 75191-75226 (Nov. 1, 2023), https://www.federalregister.gov/documents/2023/11/01/2023-24283/safe-secure-and-trustworthy-development-and-use-of-artificial-intelligence.
[14] Open Source Initiative, *The Open Source AI Definition, version 0.0.6* (Mar. 10, 2024), https://hackmd.io/@opensourceinitiative/osaid-0-0-6.
[15] *Id.* OSI specifies that access to the training data includes access to "the training methodologies and techniques, the training data sets used, information about the provenance of those data sets, their scope and characteristics; how the data was obtained and selected, the labeling procedures and data cleaning methodologies."

DocuSign Envelope ID: 36A3FB47-22DA-4C7F-9F77-903EDD214C2D

importance of including the training data in its definition of Open Models to maximize their inherent benefits while leveraging collective experience to minimize harms.

The remainder of our comment presumes that Open Models are foundation models with open weights and training data that are accessible by anyone, with limited exceptions such as those mentioned above.

3. Advantages and Benefits of Open Models[16]

Open Models provide important benefits, including (1) reducing barriers to entry and promoting innovation, (2) increasing competition, and (3) providing transparency into how models are developed, which promotes democratization and reduces the potential for bias.

a. Open Models Reduce Barriers to Entry and Promote Innovation.[17]

Open Models provide developers with broad access, options for customization, and the ability to run inferences themselves. These features greatly expand who is able to access foundation models, how new models or modifications are created, and how foundation models are used to develop applications.[18]

Developers are not required to invest vast sums of money or time to train Open Models from the ground up. Instead, they can rely on existing models, which allows for iterative

[16] This section addresses RFC question 3.

[17] This question addresses RFC questions 3.a and 6.a.

[18] Brodsky, *supra* note 6 ("With open source, there is a lot more competition, and the speed of innovation is higher. Also, the barrier to entry for someone to get into the field is much lower if there is high-quality open-source software available. That type of democratization in creation and learning for new entrants to the field is also critical.") (quoting Jignesh Patel, computer science professor at Carnegie Mellon University and co-founder of DataChat, a no-code, generative AI platform).

DocuSign Envelope ID: 36A3FB47-22DA-4C7F-9F77-903EDD214C2D

progress and the ability to leverage previous work.[19] Removing cost as a barrier (and often the most significant barrier) enables developers to engage with, refine, improve, and deploy Open Models to find new and beneficial uses for AI without the need for extensive financial resources. The accessibility of Open Models enables developers to channel time and energy toward improvements and innovative applications to specific and diverse use cases.

Relatedly, Open Models can be customized in a manner that preserves the original model but allows parallel development toward new uses and capabilities that may lack alignment with the original developer's commercial or social goals or technical abilities. In this way, Open Models promote both efficiency and innovation: developers can leverage prior, baseline work in order to focus on making improvements, and specialists with the most knowledge and experience with a particular goal or challenge can more readily focus their efforts on solving these novel problems without needing first to recreate what was done before. While customized models may lead to fragmented model usage,[20] their use of a common Open Model as a starting point makes this less likely and less problematic, since each customization can retrace its steps back to the single origin point if and when it becomes necessary to understand (or mitigate) the effects of each developer's iterative choices. This ability to leverage a single Open Model as a common starting point for specialized model architectures ultimately allows the possibility even for highly customized models to be combined or merged in ways that otherwise may prove technically infeasible.[21] In turn, Open Models enable wide collaboration among software developers. Because open models are widely accessible, software developers with different expertise and perspectives are able to contribute to the existing model and make improvements

[19] Abeba Birhane et al., *Power to the People? Opportunities and Challenges for Participatory AI*, EAAMO '22: Proceedings of the 2nd ACM Conference on Equality and Access in Algorithms, Mechanisms, and Optimization, 1–8 (Oct. 2022) https://dl.acm.org/doi/abs/10.1145/3551624.3555290.

[20] As models become more and more customized, they may become more incompatible and less interoperable with potential impact on further development. Assaf Baciu, *Tackling the Fragmented Nature of Multiple GenAI Tools*, Forbes (Dec. 26, 2023), https://www.forbes.com/sites/forbestechcouncil/2023/12/26/tackling-the-fragmented-nature-of-multiple-genai-tools/?sh=396e4f9a3c10.

[21] Colin Raffel, *Building Machine Learning Models Like Open Source Software*, Communications of the ACM (Feb. 1, 2023), https://cacm.acm.org/opinion/building-machine-learning-models-like-open-source-software/#R5.

DocuSign Envelope ID: 36A3FB47-22DA-4C7F-9F77-903EDD214C2D

beyond what would be possible on Closed Models accessible to only a small pool of individuals relying only on the skillset within that pool.

Accordingly, permitting foundation models to be provided on an open basis facilitates technological research into safety, security, and novel applications of AI in a variety of fields. Further, Open Models enable research into AI interoperability, security, and safety, while Closed Models may inherently be designed not to interoperate and are susceptible to allowing their developers to hide or obfuscate any security or safety concerns from those who wish to study, understand, and mitigate them. For example, Open Models enable developers to leverage their collective expertise to spot and fix bugs, validate model accuracy, and think through strategies for reducing risks, including bias, on a broad and collaborative basis.[22]

Open Models are also essential to scientific research, which requires that results are able to be reproduced before they are broadly accepted as valid: through the accessibility these models provide, researchers will be able to check results using the exact same tools as those originally conducting experiments or analyzing data in a way that has been undermined by—and is fundamentally not possible with—Closed Models. Moreover, Closed Models may have measures in place that may prevent some scientific and technological research entirely and for reasons known only to those who control them.[23] Thus, the existence of Open Models facilitates continued exploration of the ways in which AI can aid the scientific and academic communities in a manner that would all but be foreclosed if regulations prohibited or severely limited such models.

Open source software—which has been at the forefront of technological innovation since Linux was introduced in 1991—provides a useful parallel through which to understand the

[22] Aaron Linskens, *Open source risk management: Safeguarding software integrity*, Sonatype (Oct. 13, 2023), https://blog.sonatype.com/open-source-risk-management; *see also* Ashley Scheutt, Alison Parker & Alex Long, *Open Source Software and Cybersecurity: How unique is this problem?*, Ctrl Forward (Nov. 10, 2022), https://www.wilsoncenter.org/blog-post/open-source-software-and-cybersecurity-how-unique-problem.
[23] Shayne Longpre et al, *A Safe Harbor for AI Evaluation and Red Teaming*, arXiv.2403.04893 (Mar. 7, 2024), https://arxiv.org/abs/2403.04893.

DocuSign Envelope ID: 36A3FB47-22DA-4C7F-9F77-903EDD214C2D

benefits of supporting the development of Open Models. By providing foundational elements on an open basis, open source software allowed computer science to develop more quickly and efficiently than it otherwise would have. As of March 2024, over thirty years after its introduction, Linux is still used by 41.6% of all websites with known operating systems and powers most cloud-based machines.[24] Linux also is the backbone for training large language models. As one computer science professor and AI entrepreneur has stated, "[w]ithout that foundational building block being open-sourced, a lot of progress in the overall field of computer science would likely have been significantly slower."[25] Similarly, Python, a programming language developed under an OSI-approved open source license, is used by 49.28% of software developers worldwide.[26] In the past few decades, users have been able to further contribute to Python's development by tweaking it to support Boolean variables, sets, Unicode, context managers, generators, and keyword arguments as well as thousands of bug fixes.[27] In fact, open source software is so foundational and effective that it is "widely used across the federal government and every critical infrastructure sector."[28] Like open source software, Open Models

[24] W3techs, https://w3techs.com/technologies/comparison/os-linux,os-windows (indicating that open source software has been the foundation of technological innovation since the creation of Linux, which, as of March 2024, is used by 41.6% of all websites with known operating systems); Brodsky, *supra* note 6.
[25] Brodsky, *supra* note 6 (quoting Jignesh Patel, computer science professor at Carnegie Mellon University and co-founder of DataChat, a no-code, generative AI platform).
[26] Lionel Sujay Vailshery, *Most used programming languages among developers worldwide as of 2023*, statista (Jan. 19, 2024), https://www.statista.com/statistics/793628/worldwide-developer-survey-most-used-languages/.
[27] Raffel, *supra* note 21.
[28] *Open Source Software Security*, Cybersecurity & Infrastructure Security Agency, https://www.cisa.gov/opensource#:~:text=Open%20source%20software%20is%20widely,critical%20part%20of%20this%20effort (CISA also notes that it "open sourc[es] much of [its] codes via [its] "open-by-default" software development policy.); *see also, e.g.*, Dep't of Com. Source Code Policy, https://www.commerce.gov/about/policies/source-code; Department of Defense Chief Information Officer. *Clarifying Guidance Regarding Open Source Software (OSS)*. October 16, 2009. http://dodcio.defense.gov/Portals/0/Documents/OSSFAQ/2009OSS.pdf (describing the benefits of OSS that should be considered when conducting market research on software for DOD use); *Frequently Asked Questions regarding Open Source Software (OSS) and the Department of Defense (DoD)*,

DocuSign Envelope ID: 36A3FB47-22DA-4C7F-9F77-903EDD214C2D

will enable the quick and efficient development of foundation models and their application to solve problems in many different contexts. Developers will be able to adapt and tailor Open Models to their specific needs, driving innovation across diverse sectors of the economy with diverse benefits.

b. Open Models Increase Competition.[29]

Open Models increase competition in the development and improvement of foundation models because they do not restrict the use of AI to gatekeeper companies with the most market power or resources. This accessibility increases the prospect of competition and allows for participation by developers who may otherwise have been boxed out of working with AI due to their lacking the requisite access or resources that are necessary components to working within a closed ecosystem. Increasing participation in this manner already has been shown to drive competition and lead to better technological outcomes. For example, in discussing the potential impact of Open Models on competition, the Federal Trade Commission noted that open image-generation models have surpassed the capabilities of the Closed Models on which they were based.[30]

Because developers and startups can use Open Models as the basis for customization and improvement, the continued use of Open Models will contribute to market diversification. Such diversification serves the interests of consumers by allowing them to choose the specific AI

---

https://dodcio.defense.gov/OpenSourceSoftwareFAQ.aspx (stating that "continuous and broad peer-review, enabled by publicly available source code, improves software reliability and security through the identification and elimination of defects that might otherwise go unrecognized"); M-16-21, *Federal Source Code Policy: Achieving Efficiency, Transparency, and Innovation through Reusable and Open Source Software*, Office of Management and Budget, (August 8, 2016), https://www.whitehouse.gov/wp-content/uploads/legacy_drupal_files/omb/memoranda/2016/m_16_21.pdf.

[29] This section addresses RFC questions 3.a, 3.c, and 6.b.

[30] Staff in the Bureau of Competition & Office of Technology, *Generative AI Raises Competition Concerns*, Federal Trade Commission (Jun. 29, 2023), https://www.ftc.gov/policy/advocacy-research/tech-at-ftc/2023/06/generative-ai-raises-competition-concerns.

DocuSign Envelope ID: 36A3FB47-22DA-4C7F-9F77-903EDD214C2D

product that fits their needs and, often, to obtain such a product at a lower cost. By contrast, a regulatory environment permitting only Closed Models contributes to market concentration by restricting consumer choice, which can result in the competitive and systemic risks that arise from homogenization, including higher prices.[31] Furthermore, if a few Closed Models dominate the market and thus are integrated into a wide variety of activities, then consumers have fewer choices about which models to use and may experience a lock-in effect. If consumers have few choices and feel locked-in because the dominant Closed Models permeate their business and the broader economy, the Closed Model developers have less incentive to innovate and improve or address any errors, vulnerabilities, or failures that arise with their models. This lack of incentive can threaten a significant amount of economic activity and harm consumers.

To unleash the benefits of competition that Open Models provide, they should not be treated any differently than Closed Models from a regulatory perspective. Instead, they "should be allowed to freely proliferate and compete with both big AI companies and startups."[32]

c. Open Model Transparency Promotes Democratization and Reduces Bias.[33]

Open Models are transparent. By providing visibility into their weights and underlying training data, Open Models enable developers, regulators, and the public to analyze, critique, fix, and improve the models. They also enable those constituencies to understand and contextualize model outputs in a way that is not possible with Closed Models. It is the very openness of Open Models that facilitates the more effective identification, understanding, and mitigation of the benefits and risks posed by dual use foundation models. Additionally, this transparency allows for responsible innovation and public accountability, facilitating challenges to ingrained

[31] *Id.* ("If a single company or a handful of firms control one or several of these essential inputs, they may be able to leverage their control to dampen or distort competition in generative AI markets. And if generative AI itself becomes an increasingly critical tool, then those who control its essential inputs could wield outsized influence over a significant swath of economic activity.").
[32] Marc Andreesen, *Why AI Will Save the World*, AH Capital Management, L.L.C. ("a16z") (Jun. 6, 2023), https://a16z.com/ai-will-save-the-world/.
[33] This section addresses RFC questions 3.a, 3.b, 3.c, 3.e, 4, and 7.b.

DocuSign Envelope ID: 36A3FB47-22DA-4C7F-9F77-903EDD214C2D

assumptions and established orthodoxies. Broader scrutiny by a variety of stakeholders exposes underlying biases, inaccuracies, and vulnerabilities in these models and facilitates improvements, especially when compared with the lack of transparency inherent in Closed Models. In particular, Open Models facilitate the democratization of AI and the reduction of bias endemic to opacity and ideological uniformity.

We emphasize that Open Model weights alone are insufficient to fully understand the validity, reliability, or bias of a model.[34] Rather, identifying and correcting biases requires true and meaningful openness across the model, enabling analysis of the training data, weights and commands which together generate the outputs. Transparency throughout the model allows for the implementation of appropriate governance and accountability frameworks through which observers and other developers may monitor model performance and more accurately assess risk.[35] Accordingly, as set out in Section 2, *supra*, it is important that NTIA adopt a definition of "open" that includes a requirement that models, weights, and training data be made widely available to public scrutiny.

i. Democratization

By allowing a more diverse group of developers and other stakeholders to access AI, Open Models prevent power and influence over this pivotal technology from being concentrated in the hands of a few large stakeholders and enable broader public participation in the evolution of how dual-use foundation models are built and allowed to operate. Thus, the transparency provided by Open Models, and the increased participation it facilitates, has a democratizing effect on the AI landscape as a whole.

[34] *See* Section 2, *supra*.
[35] Rishi Bommasani et al., *The Foundation Model Transparency Index*, arXiv:2310.12941 (Oct. 19, 2023), https://arxiv.org/abs/2310.12941; Rishi Bommasani et al., *Foundation Model Transparency Reports*, arXiv:2402.16268 (Feb. 26, 2024), https://arxiv.org/abs/2402.16268.

DocuSign Envelope ID: 36A3FB47-22DA-4C7F-9F77-903EDD214C2D

If Open Models ceased to exist, and AI leveraged only Closed Models, it would allow this powerful technology to be developed solely in a series of "black boxes" by companies and developers with unilateral and unaccountable control over model behavior. Closed Model developers entirely control how information is incorporated into and flows through those models and may impose their views and preferences on the public without any meaningful guardrails. By contrast, Open Models allow a more diverse range of developers (and in some ways the public at large) to debate and define acceptable model behavior. Such access and engagement is increasingly important as models come to play a more central role in access to information. By democratizing access, Open Models support fundamental pillars of our constitutional democracy, including the marketplace of ideas, robust debate, and open discourse. These pillars support civic engagement, innovation, creative-thinking, problem-solving, and competition.

The democratization of AI facilitated by Open Models also helps build trust in AI and reduces fear of the technology. As society's use of AI has expanded, so too has the public's skepticism about what data is being used to train the models, how that data is being protected, and how it is being used to generate outputs. Because Closed Models do not offer users insight into this process, they have spurred a distrust of AI at best and a fear of a Terminator-like future at worst. By contrast, Open Models' characteristic transparency encourages user trust. The public can readily discover the kind of information to which the model has access and what the model is designed to do with that information. The public can also better understand and contextualize the risks and capabilities of AI, which helps mollify unfounded and speculative fears. And, if users do not like how a particular model handles data or weighs certain parameters, or they want to address a particular risk, they can modify the model or choose to use another that meets their needs. Thus, the marked transparency of Open Models encourages users and developers to engage with the technology, allowing it to grow and develop in a way that users understand and trust.

DocuSign Envelope ID: 36A3FB47-22DA-4C7F-9F77-903EDD214C2D

In addition, Open Models provide broader societal benefits, such as greater access to education and training to students and others, regardless of income or geography.[36] Such education is critical to developing the skills necessary to succeed in an age where AI will rapidly reshape the workforce. Even when Open Models do not outperform their closed counterparts, their "widespread availability is a boon to students all over the world who want to learn how to build and use AI to become part of the technological future, and will ensure that AI is available to everyone who can benefit from it no matter whole they are or how much money they have."[37]

ii. Bias

Open Models reduce the potential for bias by allowing broad visibility into model weights and training data through which biases in model parameters, training data, and outputs can be identified. Public scrutiny and the accountability it enables are the most effective means of preventing and combating bias as AI becomes more widely used. This is particularly relevant in regards to bias introduced through the training data itself. Providing this training data on an open basis means that developers with a variety of perspectives and expertise can access and review this data, subjecting it to scrutiny in order to understand how it influences model outputs. Based on this examination, developers can challenge any biases inherent in that data in order to improve model accuracy and neutrality and have visibility into how these biases are likely to influence and bias model outputs, as well.

Closed Models, on the other hand, do not present their weights or training data for public scrutiny and therefore are more vulnerable to undisclosed biases. Because the models are closed, these effects might not be observed until the models are in use and the harm has already materialized, with the potential for significant impact on downstream activity. To be sure, Closed Models have relied on other methods to try to achieve neutral outputs. But these efforts have had varying degrees of success. For example, evidence suggests that some Closed Models

[36] Shrestha, Y.R. et al., *Building Open-Source AI.*, Nature Computational Science 3, 908-911 (Oct. 26, 2023), https://www.nature.com/articles/s43588-023-00540-0#citeas.
[37] *Id.*

DocuSign Envelope ID: 36A3FB47-22DA-4C7F-9F77-903EDD214C2D

exhibiting so-called unbiased outputs are, in fact, merely manipulating prompts or weights to hide bias rather than actually reducing the bias.[38] Developers may program models to achieve such an artificial outcome through technical manipulation, such as adding terms to users' prompts post-hoc to generate particular favored outputs or tweaking the model such that certain favored outputs are displayed first.[39] These findings underscore the importance of open training data as well as model weights, since poorly selected or otherwise biased or unrepresentative training data, in addition to manipulated prompts or weights, may infect a model in a way that otherwise is difficult to detect. In contrast, Open Models not only allow developers to dissect a model's weights and training data to find its flaws and biases, but also enable developers to improve the models, or develop new ones, to fix those flaws, limit bias, and reduce errors.

Open Models also prevent a limited group of gatekeeping stakeholders from determining what outputs are and are not considered socially acceptable. In a world where AI is likely to be the control layer for how we engage with the world, how it is allowed to operate will influence what we consume, say, and think. Closed Models put control over those decisions in the hands of the concentrated few, powerful companies and individuals who get to pick and choose what we consume and say based on their view of what is good for society.[40] As a result, outputs may reflect only the limited perspectives and biases contained within largely ideologically uniform Closed Model development teams and, through offering only these limited perspectives, unnaturally shape public opinion to match their own. Additionally, those same teams are often solely responsible for evaluating, identifying and correcting biased outputs—a practice antithetical to peer-review and the widely established benefits thereof. By democratizing access, Open Models support principles of free expression, diversity of thought, and public debate of difficult and contentious issues. Developers and users can identify and engage with

[38] Gerrit De Vynck and Nitasha Tiku, *Google takes down Gemini AI image generator. Here's what you need to know.*, Wash. Post (Feb. 23, 2024), https://www.washingtonpost.com/technology/2024/02/22/google-gemini-ai-image-generation-pause/; Sigal Samuel, *Black Nazis? A woman pope? That's just the start of Google's AI Problem*, Vox (Feb. 28, 2014), https://www.vox.com/future-perfect/2024/2/28/24083814/google-gemini-ai-bias-ethics.
[39] De Vynck and Tiku, *supra* note 38.
[40] Andreesen, *supra* note 32 ("AI is highly likely to be the control layer for everything in the world.").

DocuSign Envelope ID: 36A3FB47-22DA-4C7F-9F77-903EDD214C2D

each other's ideas, perspectives, and biases and openly, freely, and publicly debate how to resolve challenging issues. The future of AI should be determined by everyone, not by the agendas of a select few, and Open Models are essential to that future.

4. Rules Should not Prohibit Open Models or Regulate Open Models Differently than Closed Models Based on a Misconception of Comparative Risk[41]

While Open Models offer numerous significant benefits, they, just like Closed Models and any other tool, pose a risk that they will be misused to cause harm. But the types of output that often raise concerns, such as how to make a bomb, lethal pathogen, or cyberweapon, can already be found on the public internet or dark web.[42] Therefore, it is important to think about Open Models in terms of marginal risk and opportunity cost.

Take privacy, for example. A common criticism of Open Models is that they pose a significant privacy risk if personal data is included in the training data and made public. Certainly, this is a risk; but risk should not be and cannot be addressed in a vacuum. The relevant point of comparison, Closed Models, also presents risk. Closed Models are not impenetrable. Personal data used to train Closed Models is vulnerable to exploitation by threat actors. For example, threat actors can cause Closed Models to disclose non-public training data, including sensitive personal data, by crafting prompts that circumvent the models' safeguards.[43] Personal

---

[41] This section addresses RFC questions 2.a, 2.d, 2.d.i, 2.d.ii, 3.b, 3.d, 5.b, 5.d, 6.a, and 7.d.

[42] Adrian Bridgwater, *Avoiding The Swinging Pendulum In The Great AI Debate* (Mar. 5, 2024), https://www.forbes.com/sites/adrianbridgwater/2024/03/05/avoiding-the-swinging-pendulum-in-the-great-ai-debate/?sh=31ef32fb185d.

[43] *See* Jeremy White, *How Strangers Got My Email Address From ChatGPT's Model*, N.Y. Times (Dec. 22, 2023), https://www.nytimes.com/interactive/2023/12/22/technology/openai-chatgpt-privacy-exploit.html; *see also* Nicholas Carlini et al., *Extracting Training Data from Large Language Models*, arXiv:2012.07805 (Jun. 15, 2021), https://arxiv.org/pdf/2012.07805.pdf; Lance Eliot, *Prompt Engineering Boasts New Practice Of Telling Generative AI To Be On Your Toes And Stridently Question Wishful Miracles*, Forbes (Mar. 21, 2024), https://www.forbes.com/sites/lanceeliot/2024/03/21/prompt-engineering-boasts-new-practice-of-telling-

DocuSign Envelope ID: 36A3FB47-22DA-4C7F-9F77-903EDD214C2D

data is also subject to misuse and abuse by the companies providing the Closed Models, with limited visibility into their data handling practices to be able to hold them accountable. Moreover, Open Models possess certain privacy-enhancing features, such as greater visibility into their data handling practices than Closed Models, which increases user trust in how models are protecting their personal data. And, in cases where the models fail to protect such data, this visibility enables users and the public to monitor for misuse of personal data, hold companies accountable for such misuse, and choose to use models that do adequately protect that information. Additionally, unlike Closed Models, which require data to be sent to the Closed Model developers for training and inferences, Open Models allow users to directly use the models without the need to share confidential information or sensitive personal data with third parties.[44] Instead, developers can adapt Open Models to their purposes and perform inferences locally, reducing privacy concerns.

Viewed through the lens of marginal risk and considering the privacy-enhancing features of Open Models, it is much less clear what risk Open Models pose that could justify treating them differently than Closed Models. Accordingly, the relevant question when assessing Open Models is what additional risks do Open Models create that do not already exist on the internet or that Closed Models do not present, and what benefits and advances we potentially sacrifice by restricting Open Models.[45]

---

generative-ai-to-be-on-your-toes-and-stridently-question-wishful-miracles/?sh=4d22d73a1275; Hayden Field, *Inside the largest-ever A.I. chatbot hack fest, where hackers tried to outsmart OpenAI, Microsoft, Google, CNBC* (Aug. 15, 2023), https://www.cnbc.com/2023/08/15/def-con-hackers-try-to-crack-chatbots-from-openai-google-microsoft.html; Jordan Pearson, *ChaptGPT Can Reveal Personal Information From Real People, Google Researchers Show*, Vice (Nov. 29, 2023), https://www.vice.com/en/article/88xe75/chatgpt-can-reveal-personal-information-from-real-people-google-researchers-show.

[44] *Open-Source vs OpenAI: Is it Time to Move On?*, HoneyHive, https://www.honeyhive.ai/post/openai-vs-open-source-models ("Open-source models, on the other hand, offer the advantage of retaining sensitive data within a company's own cloud environment, thereby safeguarding data privacy and avoiding the risks associated with transmitting sensitive information to external entities."). For example, Baseten offers users the option to deploy its open source models within the users' own cloud environments so that users' customized models and data never leave their *virtual* private clouds. *See* Baseten, https://www.baseten.co/.

[45] *Id.*

DocuSign Envelope ID: 36A3FB47-22DA-4C7F-9F77-903EDD214C2D

Foundation models, both Open and Closed, are nothing more than mathematical representations of a defined world—tools that are used by humans in order to accomplish tasks—and are not by their nature a threat to national security or economic interests. Technology is a tool. And any technology can be used for good or bad.[46] In any event, the cat of Open Models is already out of the bag and has been released into the wilds of commerce such that it likely is impossible, as a practical matter, to claw back their development and distribution without doing fundamental damage to the existing AI infrastructure.[47] Once an Open Model has become widely available—and there are already many out there—you cannot lock it back up, especially when it comes to bad actors.

Governments rarely have found success in banning tools simply because they have the potential for misuse, nor has this generally been the approach the United States has taken regarding developing technology. While it is possible for bad actors to use Open Models to cause harm, regulations and prohibitions should focus on the conduct the government wishes to prohibit rather than the tools with which such conduct is undertaken. The harms for which Open Models may be leveraged already are prohibited by existing federal laws such that it is unnecessary to create rules targeted at the Open Models, themselves, rather than the prohibited conduct.[48]

---

[46] Andreesen, *supra* note 32 ("Technology is a tool. Tools, starting with fire and rocks, can be used to do good things—cook food and build houses—and bad things—burn people and bludgeon people. Any technology can be used for good or bad.").

[47] Andreesen, *supra* note 32 ("The AI cat is obviously already out of the bag. You can learn how to build AI from thousands of free online courses, books, papers, and videos, and there are outstanding open source implementations proliferating by the day. AI is like air – it will be everywhere.").

[48] *See, e.g.*, Computer Fraud and Abuse Act ("CFAA"), 18 U.S.C. § 1030 (establishing the primary federal statutory mechanism for prosecuting cybercrime, including hacking and some related extortion crimes in the context of ransomware); 18 U.S.C. §§ 1341-1351 (federal statutes prohibiting fraud including mail fraud, wire fraud, bank fraud, healthcare fraud, and securities and commodities fraud); 15 U.S.C. § 45 (Section 5 of the Federal Trade

DocuSign Envelope ID: 36A3FB47-22DA-4C7F-9F77-903EDD214C2D

Furthermore, Closed Models present their own opportunities for misuse such that it would be misguided to structure rules in a manner that distinguishes between whether a model is open or closed. The restricted nature of Closed Models is likely to prevent risks from being identified in the first place, creating a false sense of security in models that are not subject to outside scrutiny. Perhaps this false sense of security is why proponents of Closed Models argue that they are inherently more secure than Open Models. However, there are confirmed reports of foreign adversaries that have infiltrated most of the companies that develop and advocate for Closed Models,[49] and there are also reports that foreign adversaries have used such models to cause harm.[50] Assuming those reports are accurate, either Closed Models are not inherently more secure than Open Models or the companies developing Closed Models have not taken appropriate steps to secure them. Either way, it is difficult to justify treating the two types of models differently based on Closed Model proponents' claim of superior security.

If Closed Models that dominate the market are vulnerable to malicious attacks and a threat actor exploited a vulnerability, such a breach could result in a successful attack against significant parts of the economy reliant on the particular breached model. Algorithmic monocultures resulting from reliance on a few Closed Models can create resilience problems and generate systemic risk. If those models are compromised, the impacts could be widespread and pervasive. As in the proprietary software context, Closed Models present a security risk

---

Commission Act prohibiting "unfair or deceptive acts or practices in or affecting commerce"); 18 U.S.C. 113B §§ 2332. 2332a, 2332b, 2339, 2339A, 2339B, 2339C (federal statute criminalizing terrorism and material support to terrorists); 18 U.S.C. § 175 (federal statute establishing prohibitions with respect to biological weapons).

[49] Ian Thomas, *China and cybercriminals are targeting American AI companies, FBI Director Wray says*, CNBC (Jan. 9, 2024), https://www.cnbc.com/2024/01/09/china-and-cybercriminals-are-targeting-american-ai-companies.html.

[50] Karen Weise, *Hackers for China, Russia and Others Used OpenAI Systems, Report Says*, N.Y. Times (Feb. 14, 2024), https://www.nytimes.com/2024/02/14/technology/openai-microsoft-hackers.html; Christopher Hutton, *Foreign adversaries using AI to improve cyberattacks, Microsoft and OpenAI warn*, Wash. Examiner (Feb. 14, 2024), https://www.washingtonexaminer.com/news/2855805/foreign-adversaries-ai-improve-cyberattacks-microsoft-openai-warn/; Microsoft Threat Intelligence, *Staying ahead of threat actors in the age of AI*, Microsoft (Feb. 14, 2024), https://www.microsoft.com/en-us/security/blog/2024/02/14/staying-ahead-of-threat-actors-in-the-age-of-ai.

DocuSign Envelope ID: 36A3FB47-22DA-4C7F-9F77-903EDD214C2D

because users cannot customize the model to fit their needs and requirements, and the models may contain security vulnerabilities that users cannot access, evaluate, or change.

By contrast, the transparency afforded by Open Models allows potential harms to be identified more quickly, thereby allowing the collective expertise of the broader AI community to aid in developing strategies and technical mitigations either to make those harms more difficult to exploit or—where such prophylaxis is unfeasible—to develop technical countermeasures designed to restore the pre-harm status quo. Parallels can be seen in the fraud and cybersecurity sectors, which have moved from closed threat modeling to more collaborative and open systems, thereby reducing the likelihood that threat actors can move from one institution to another causing similar harms while circumventing common controls and evading detection. Moreover, Open Models achieve security through decentralization. Even if a threat actor does find a vulnerability in a particular model, because Open Models promote proliferation and diversification of models, it will be more difficult for a threat actor to exploit that vulnerability across many different models.

As the Aspen Institute's Global Cybersecurity Group has observed, "[t]raditional wisdom in computer security is that security by obscurity does not work, and that systems built on open source are *more secure* than closed source systems."[51] That's because open source systems allow anyone to examine the software, identify vulnerabilities, report them, and apply fixes before attackers find them and use them. This type of distributed security does not depend on a single vendor to protect the software, often resulting in vulnerabilities being identified and fixed much more quickly than in closed source systems.[52] Thus, it does not make sense to treat Closed Models more favorably or leniently from a regulatory perspective when Closed Models do not actually fulfill their purported security advantage over Open Models and they lack all the benefits of Open Models discussed above.

[51] Aspen Institute, *supra* note 11 at 32.
[52] Seger, *supra* note 9 at 18.

DocuSign Envelope ID: 36A3FB47-22DA-4C7F-9F77-903EDD214C2D

5. Government Should Not Pick Winners and Losers Among AI Approaches in this Nascent and Pivotal Stage of Development[53]

NTIA should proceed cautiously and in a technology-agnostic manner when developing rules. While it is important to recognize that, as foundation models develop, there will undoubtedly be unforeseen risks and benefits, given the pace of technological change, it is very difficult—and quite possibly counterproductive—to try to make decisions about or set thresholds to mitigate risk for foundation models in the future. As discussed above, one way to address the potential risks posed by foundation models, both open and closed, is to enforce existing federal law prohibiting harmful conduct such that new rules and regulations targeted at foundation models themselves are unnecessary.

In considering how to maximize the benefits and minimize the harms from Open Models, NTIA should look to existing U.S. and international frameworks that already provide a basis for AI governance without requiring the government to develop a new regulatory paradigm and pick winners and losers. While we do not endorse any of these frameworks, we nonetheless consider them illustrative as to the kinds of alternate, consensus-driven, and technology-agnostic approaches that are preferable to broad regulations and those that prohibit Open Models outright. For example:

- The NIST AI Risk Management Framework provides a consensus-driven, collaborative approach to incorporating and improving trustworthiness considerations into the design, development, and use of AI products, services, and systems. The Framework is intended for voluntary use and seeks to improve developers' ability to incorporate

[53] This section addresses RFC questions 7, 7.d, 7.d.i, 8, and 8.b.

DocuSign Envelope ID: 36A3FB47-22DA-4C7F-9F77-903EDD214C2D

trustworthiness considerations into the design, development, and use of AI products, services, and systems.[54]

- The U.S. Department of Homeland Security's ("DHS") Cybersecurity and Infrastructure Security Agency ("CISA") and the United Kingdom's National Cyber Security Centre ("NCSC") released nonbinding guidelines for providers to deploy 'secure by design' AI systems. Sixteen countries signed on to the guidelines along with the United States and the United Kingdom. The guidelines are intended to help developers of any systems that use AI make informed cybersecurity decisions at every stage of the development process. The guidelines provide essential recommendations for AI system development and emphasize the importance of adhering to Secure by Design principles.[55]

- While we disagree with the European Union's overall regulatory approach in its Artificial Intelligence Act ("AI Act"), we note that, despite the Act's extensive regulation of AI systems, it does not ban the use of Open Models and, in fact, acknowledges that there is a role for Open Models in the AI ecosystem and provides certain exceptions for them from the regulation's obligations.[56]

---

[54] U.S. Dep't of Com., Nat'l Inst. of Standards and Tech., NIST AI 100-1, Artificial Intelligence Risk Management Framework (AI RMF 1.0) (Jan. 2023), https://doi.org/10.6028/NIST.AI.100-1. NIST has also published an AI Risk Management Framework Playbook that suggests recommended actions for organizations to achieve the objectives set out in the AI RMF 1.0. U.S. Dep't of Com., Nat'l Inst. of Standards and Tech., *NIST AI RMF Playbook*, https://airc.nist.gov/AI_RMF_Knowledge_Base/Playbook.

[55] Press Release, Cybersecurity & Infrastructure Security Agency, *DHS CISA and UK NCSC Release Joint Guidelines for Secure AI System Development* (Nov. 26, 2023), https://www.cisa.gov/news-events/news/dhs-cisa-and-uk-ncsc-release-joint-guidelines-secure-ai-system-development; Cybersecurity & Infrastructure Security Agency & National Cyber Security Centre, Guidelines for secure AI system development (Nov. 26, 2023), https://www.ncsc.gov.uk/files/Guidelines-for-secure-AI-system-development.pdf.

[56] Press Release, Council of the EU, *Artificial intelligence act: Council and Parliament strike a deal on the first rules for AI in the world* (Dec. 9, 2023), https://www.consilium.europa.eu/en/press/press-releases/2023/12/09/artificial-intelligence-act-council-and-parliament-strike-a-deal-on-the-first-worldwide-

DocuSign Envelope ID: 36A3FB47-22DA-4C7F-9F77-903EDD214C2D

Several other recent agreements relating to AI confirm that the prevailing approach is to proceed in a technology-agnostic manner, focusing on addressing harmful conduct perpetrated using AI rather than the models themselves. For example, the Leaders of the Group of Seven (G7) recently reached agreement through the Hiroshima AI Process on the Hiroshima Process International Code of Conduct for Organizations Developing Advanced AI Systems.[57] The Code of Conduct provides voluntary guidance for actions by organizations developing the most advanced AI systems. Notably, the Code of Conduct encourages a risk-based approach to its recommendations and does not discuss treating Open Models differently than Closed Models.[58] Similarly, the Council of Europe's Committee on Artificial Intelligence reached agreement on the text of a treaty that would require AI developers and deployers to respect human dignity, privacy, the rule of law, and democracy.[59] The treaty also takes a risk-based approach and applies its framework to all foundation models, both open and closed.

In addition to being technology-agnostic, NTIA should proceed cautiously and maximize flexibility when developing rules for foundation models. The technology is changing so rapidly that regulations may quickly become obsolete or even counterproductive or harmful. Accordingly, NTIA should take an approach that is consistent with a nascent and evolving technology. Self- and cooperative-regulatory approaches are the most effective and practical way to prevent and address AI-related problems within the boundaries set by sector-specific regulation. Such approaches allow and encourage participation by industry and civil society,

---

rules-for-ai/c; Regulation 2021/0106 (COD) of the European Parliament and of the Council of Jan. 26, 2024, Laying Down Harmonised Rules on Artificial Intelligence (Artificial Intelligence Act) and Amending Certain Union Legislative Acts, at ¶¶ (57e), (60i), (60i+1), (60f), (60o); art.2, 5g.; art. 28, 2b.; art. 52c, -2.; art. 52ca, https://data.consilium.europa.eu/doc/document/ST-5662-2024-INIT/en/pdf.

[57] Group of 7 (G7), *Hiroshima Process International Code of Conduct for Organizations Developing Advanced AI Systems* (Oct. 30, 2023), https://www.mofa.go.jp/file.

[58] *Id.*

[59] Council of Europe, Committee on Artificial Intelligence, Draft Framework Convention on Artificial Intelligence, Human Rights, Democracy and the Rule of Law (Dec. 19, 2023) (agreed Mar. 14, 2024), https://rm.coe.int/cai-2023-28-draft-framework-convention/1680ade043.

DocuSign Envelope ID: 36A3FB47-22DA-4C7F-9F77-903EDD214C2D

which have a broader level of expertise in burgeoning technologies. These approaches also allow governance frameworks to remain flexible in a way that static regulation cannot, evolving and adapting over time as the technology rapidly changes.

Open Models are particularly well-suited to and support flexible self- and cooperative-regulatory approaches. The need to set predetermined computational resource thresholds, or, more generally, to evaluate risk and take action based on necessarily incomplete information about the future derives from the opacity of Closed Models. Because of the lack of transparency, regulators cannot properly assess the risks and benefits posed by such models and can only guess at how to regulate an obscured, moving target. Consequently, any regulatory action is likely to be over- or under-inclusive, inefficient, counterproductive, or harmful. By contrast, Open Models do not require regulators to make future predictions or adopt prescriptive regulations based on an incomplete view of how AI will develop or concepts of risk that may prove outdated as the technology changes. Open Models allow for adaptation and adjustment of regulations collaboratively as the technology changes because of the transparency they provide. "[T]ransparently disclosing information makes that information available, shareable, legible, and verifiable."[60] That information can be used to inform the assessment of risks and benefits and improve decision-making in real time. Such an approach can promote accountability and tailored risk-mitigation without the efficiency and innovation loss that may result from rigid one-size-fits-all standards imposed in advance with imperfect information. Therefore, NTIA should leverage the built-in accountability that comes with the transparency of Open Models to manage the risks, and allow Open Models to develop freely—with all the attendant benefits such development will entail for the world.

The markets, rather than government, should determine which technologies ultimately are successful based on their performance, utility, and ability to earn consumer confidence and trust. Given our discussion of the benefits and risks of Open Models and Closed Models, there is no reason from a regulatory perspective to treat these types of foundation models differently. Indeed, no other regulatory regime in the United States or around the world has adopted such

[60] Bommasani, et al., *The Foundation Model Transparency Index*, *supra* note 35 at 10.

DocuSign Envelope ID: 36A3FB47-22DA-4C7F-9F77-903EDD214C2D

an approach. It is entirely possible that Open Models and Closed Models each find their own place in the widely evolving technological landscape. NTIA should not short-circuit this process by picking winners and losers among foundation models in this nascent and pivotal stage of development.

6. Conclusion

a16z appreciates the opportunity to share its perspective on NTIA's RFC. AI in the United States is at an inflection point, and we urge NTIA to consider the likely impact of any forthcoming regulations or guidance on this emerging and essential industry. While doomsayers fear AI as an existential threat to humanity, this pessimism is unfounded. The continued development of foundational AI models represents a profound opportunity to augment human intelligence and bring untold benefits to people around the world.

The transformative nature of this technology also offers tremendous financial gain to AI companies, some of which may seek to hamstring their competition by calling for regulation of Open Models under misguided and alarmist claims about risk. Thus, we encourage NTIA to be wary of generalized claims about the risks of Open Models and calls to treat them differently from Closed Models, especially when such claims are made by AI companies seeking to insulate themselves from market competition. In fact, as we have explained, Open Models have the distinct advantage of allowing governments, regulators, and the public to understand them—including their capabilities, risks, and underlying training data—in a more transparent manner. It is the very openness of Open Models that enables the more effective identification, understanding, and mitigation of risk.

Accordingly, NTIA should not treat Open Models any differently than Closed Models. As one of the founders of our firm has stated:

"Startup AI companies should be allowed to build AI as fast and aggressively as they can. They should neither confront government-granted protection of big companies, nor should they receive government assistance. They should simply be allowed to compete. If and as startups

DocuSign Envelope ID: 36A3FB47-22DA-4C7F-9F77-903EDD214C2D

don't succeed, their presence in the market will also continuously motivate big companies to be their best – our economies and societies win either way."[61]

We urge NTIA to adopt relevant definitions and rules that allow for the continued development of Open Models rather than a restrictive approach through which the government, and not technological and market developments, picks winners and losers in this important, emerging market.

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

By:

DocuSigned by:
Jai Ramaswamy
48A50A932BBE431...
Jai Ramaswamy
Chief Legal Officer

DocuSigned by:
3E3E8C67CB98483...
Collin McCune
Head of Government Affairs

[61] Andreesen, *supra* note 32.