Response to the NSCEB's Interim Report and AIxBio Policy Options

Submitted by the Federation of American Scientists, Johns Hopkins Center for Health Security, Nuclear Threat Initiative Global Biological Policy and Programs, and Scowcroft Institute of International Affairs

Executive Summary

Thank you for the opportunity to provide comments in response to the National Security Commission on Emerging Biotechnology's (NSCEB) recent Interim Report[1] and White Paper on Policy Options for Artificial Intelligence and Biotechnology (AIxBio).[2] The comments expressed herein reflect the views of the Federation of American Scientists, Johns Hopkins Center for Health Security, Nuclear Threat Initiative Global Biological Policy and Programs (NTI|bio), and Scowcroft Institute of International Affairs at the Bush School of Government and Public Service and do not necessarily reflect the views of Johns Hopkins University or Texas A&M University.

The Federation of American Scientists (FAS) is an organization dedicated to using science and technology to address global threats and advance policy and innovation for a healthy, safe, and equitable society, founded by atomic researchers after the bombings of Hiroshima and Nagasaki. For 25 years, the Johns Hopkins Center for Health Security (CHS) has protected people's health from major epidemics and disasters and built resilience to those challenges by conducting independent research and analyzing how scientific and technological innovations can strengthen health security. NTI|bio transforms global security by driving system solutions to biological threats imperiling humanity. The Scowcroft Institute of International Affairs is a research institute housed in the Bush School of Government and Public Service with a core mission to foster and disseminate policy-oriented research on international affairs, including pandemic preparedness and biosecurity. Collectively, these organizations represent some of the major civil society organizations with expertise in biosecurity that are developing policies around the governance of AIxBio.

Congress tasked the NSCEB with conducting a thorough review of how advancements in emerging biotechnology and related technologies will shape current and future national defense activities. The NSCEB's Interim Report[3] was submitted in December 2023 and its final report is due to Congress in December 2024, which will include policy recommendations that respond to its mandate. We highly commend the leadership and staff of the NSCEB for its recommendations to date, which aim to promote the responsible innovation and development of biotechnology. In particular, we appreciate that the NSCEB's three pillars of its path forward include: (1) preparing the US government for the age of biology; (2) accelerating innovation and embracing biotechnology; and (3) protecting against misuse and promoting norms for responsible use. These pillars provide a good foundation for the NSCEB to meet its

[1] US National Security Commission on Emerging Biotechnology, *Interim Report* (Dec. 2023), https://www.biotech.senate.gov/press-releases/interim-report/ [hereinafter *Interim Report*].
[2] US National Security Commission on Emerging Biotechnology, *AIxBio White Paper 4: Policy Options for AIxBio* (Jan. 31, 2024), https://www.biotech.senate.gov/press-releases/aixbio-white-paper-4-policy-options-for-aixbio/ [hereinafter *White Paper*].
[3] *Interim Report*.

mission. Additionally, we support many of the NSCEB's recommendations in its White Paper on AIxBio Policy Options.

In conducting its work, we strongly recommend that the NSCEB prioritizes the mitigation of two particular classes of pandemic risk when making recommendations about protecting biotechnology against misuse and promoting norms for responsible use, which we outline below. Additionally, we briefly outline and then provide more detailed responses to each of the NSCEB's recommendations in its White Paper on AIxBio Policy Options and offer additional suggestions for the NSCEB to consider for its final report.

Pursuing the Promise of AI

As has been noted in recent Senate testimony,[4] AI holds great promise for benefits in public health. Potential benefits include earlier disease diagnoses, allowing doctors to intervene sooner in the course of an illness; reduced medical errors; more efficient or less invasive surgeries; reduced administrative burdens on clinicians to allow more time with patients; and faster response times to patient questions. Researchers and companies may be able to create or use AI tools to help them accelerate development of vaccines and medicines and to significantly advance personalized medicine. AI may be able to improve disease surveillance and perhaps even provide earlier indicators of new outbreaks or epidemics. It will place stronger diagnostic and clinical tools in the hands of providers in the field or those in clinics in under-resourced healthcare systems. AI could also assist with more careful monitoring of drug safety and help to improve, and potentially greatly accelerate, clinical trials of new medicines. These are all extraordinary benefits that we should encourage the responsible development and implementation of in the years ahead.

Preventing or Mitigating AIxBio Risks

We commend the NSCEB for including the pillar of "protecting against misuse and promoting norms for responsible use" in its work moving forward. We were pleased to see that the NSCEB included this pillar in its recent AIxBio White Paper on Policy Options by incorporating recommendations that aim to provide oversight of AI models for biotechnology and assess future AIxBio risk.[5] We were also particularly excited to see the following passage under this pillar in the Interim Report:

> "Emerging technologies may themselves provide the technical capabilities to preempt, detect, and mitigate misuse concerns, and we are actively exploring ways that the Commission may further encourage the development and implementation of these technologies. For example, wastewater surveillance could help with early detection of biological threats. We plan to explore best practices for responsible

[4] *Avoiding a Cautionary Tale: Policy Considerations for Artificial Intelligence in Health Care: Hearing Before the S. Subcomm. on Primary Health & Ret. Sec., Comm. On Health, Educ., Labor and Pensions*, 118th Cong. 3 (2023) (Statement of Dr. Tom Inglesby, Director, Johns Hopkins Center for Health Security), https://www.help.senate.gov/imo/media/doc/79536a31-d1cf-25b0-e526-52ba2193d900/Tom%20Inglesby,%20Nov.%208%20HELP%20Subcommittee%20Written%20Testimony.pdf.

[5] *White Paper* at 2.

> innovation that prevents misuse. For example, there are currently no codified best practices for DNA synthesis screening or development of hardware and software safeguards within synthesizers. We plan to assess options for codifying those best practices, including identifying private and government stakeholders responsible for implementing the best practices."[6]

However, we believe that while measures like wastewater surveillance and gene synthesis screening are indeed critical to mitigating risks, we also note that there are important risk prevention and mitigation measures that can and should be considered (and eventually codified) prior to the digital-to-physical transition. At a minimum, mandatory safety evaluations and red teaming that meet established baselines should screen highly capable AI models for pandemic-level risks.

We have learned through much of our work on dual-use research of concern (DURC) and oversight of research with enhanced potential pandemic pathogens (ePPP)[7] that it is important to clearly define the risks that should trigger additional oversight and those that warrant risk assessments prior to proceeding. Some of the extraordinary potential benefits of AI outlined above will necessarily include management of dual-use risks. To efficiently eliminate or reduce those risks, we must articulate which risks need to be addressed as a highest priority and denote if there are unacceptable levels of risks to the public, as compared to potential public benefits.

Accordingly, as the NSCEB works to further operationalize its pillar of "protecting against misuse and promoting norms for responsible use" for its final report, we recommend that the NSCEB, at a minimum, prioritize and mitigate high-consequence biological risks to ensure that America continues to lead in biotechnology innovation safely. We judge two of those high-consequence risks for biotechnology to be AI capabilities that:

(1) Accelerate or simplify the reintroduction of extinct viruses with pandemic potential, or viruses with pandemic potential that only exist now within research labs or virus repositories; or
(2) Enable, accelerate, or simplify the creation of novel or enhanced biological constructs that could start pandemics.

These are not the only risks, but they are potentially particularly severe risks and should be assessed carefully. We support biosecurity-related red-teaming exercises performed on frontier foundation models, and encourage AI labs to ensure a focus of red-teaming efforts is on pandemic-level threats (beyond a focus on non-transmissible agents).The NSCEB could make a significant contribution to risk reduction by recommending that mandatory evaluations should, at a minimum, screen highly capable models for pandemic-level risks prior to model deployment and that any such risks discovered through that process should be mitigated. Mitigation efforts following the identification of pandemic-level risks could include risk-reduction measures such as limiting training data sets, initiating possible un-learning

[6] *Interim Report* at 32.
[7] *See Center for Health Security faculty respond to White House Office of Science and Technology Policy RFI on Dual Use Research of Concern and Potential Pandemic Pathogen Care and Oversight Policy Framework*, CTR. HEALTH SEC. (Oct. 16, 2023), https://centerforhealthsecurity.org/2023/center-for-health-security-faculty-respond-to-whitehouse-office-of-science-and-technology-policy-rfi-on-dual-use-research-of-concern-and-potential.

mechanisms on troublesome training data, developing a process to prevent the public release of pandemic-level risk information, and preventing the digital-to-physical transition to actual biological constructs.

In addition, many organizations, including us, have recommended in a range of forums that the codification of DNA synthesis screening requirements (of both customers and orders) should be established now to both (i) mirror the AI executive order requirement for federally funded life-sciences research[8] but also extend to non-federally funded entities and (ii) require DNA synthesis providers to screen.[9] While the NSCEB considers what it will recommend for codifying best practices,[10] we strongly recommend that the NSCEB assesses DNA synthesis screening practices, software, and hardware options to screen against high-consequence risks, including by prioritizing the prevention of the two classes of pandemic risks outlined above.

Securing biotechnologies against, at a minimum, this narrow subset of high-consequence risks will permit the vast majority of AI biotechnology work to move ahead unobstructed. We hope the NSCEB similarly prioritizes these risks and includes recommendations for guarding against them in its final report. Examination of other potential high-consequence biological risks that could be generated by AI model capabilities should continue, but as an urgent near-term minimum step, a focus on identifying and preventing pandemic-related outcomes should be a high priority. Policy on these issues should aim to delineate unacceptable levels of risk compared to potential benefits. These risks will require federal agency expertise and focus, as well as close AI developer and other stakeholder engagement to develop strong and effective governance.

Select Responses to Strengthen NSCEB's AIxBio Policy Options

Below, we provide our high-level responses to each of a select number of NSCEB's AIxBio policy options to ensure responsible and trustworthy biotechnology innovation. For more information, please see our detailed responses to each policy option noted below, in addition to others, beginning on the next page of this document. Policy options and their responses below are selected and ordered according to either the strength of the NSCEB's recommendation or our response to it.

1) NSCEB: Develop an AIxBio consortium.
    - **Response:** Any such consortium should prioritize information sharing about how to identify

[8] Exec. Order No. 14110, 88 Fed. Reg. 75191 (2023), § 4.4(b)(iii).
[9] *See, eg, Avoiding a Cautionary Tale: Policy Considerations for Artificial Intelligence in Health Care: Hearing Before the S. Subcomm. on Primary Health & Ret. Sec., Comm. On Health, Educ., Labor and Pensions*, 118th Cong. 3 (2023) (Statement of Dr. Tom Inglesby, Director, Johns Hopkins Center for Health Security), https://www.help.senate.gov/imo/media/doc/79536a31-d1cf-25b0-e526-52ba2193d900/Tom%20Inglesby,%20Nov.%208%20HELP%20Subcommittee%20Written%20Testimony.pdf; *Center for Health Security Submits Senate RFC on PAHPA Reauthorization*, CTR. HEALTH SEC. (July 10, 2023), https://centerforhealthsecurity.org/2023/center-for-health-security-submits-senate-rfc-on-pahpa-reauthorization; *Center for Health Security Applauds the Introduction of 2 Critical Bills to Protect Health Security from Potential Threats Arising from Advances in Biotechnology*, CTR. HEALTH SEC. (July 19, 2023), https://centerforhealthsecurity.org/2023/center-for-health-security-applauds-the-introduction-of-2-critical-bills-to-protect-health-security-from-potential-threats-arising-from-advances-in.
[10] *Interim Report* at 32.

and mitigate high-consequence biosecurity risks, prevent the public release of model-generated information that increases such risks, and prevent their digital-to-physical transition.

2) NSCEB: Initiate a global competitive analysis focused on AIxBio.
   - **Response:** Global competitive analyses are a useful input for AI governance policy generation. However, the top goal of AI governance policy should be to develop governance processes to realize benefits while mitigating risks to the public of high-consequence, pandemic-related risks.
3) NSCEB: Establish an independent group to conduct flexible risk assessments.
   - **Responses:**
     - Before establishing an independent group, the NSCEB should make clear that such group will be tasked with, among other things, identifying the highest priority risks, defining the models that fall within scope, and establishing processes for (i) assessing models for such risks, (ii) preventing such risks from becoming fully open-source, and (iii) preventing the digital-to-physical transition of pandemic risks. Knowing these duties before such a group is established ensures that the group's composition/membership is commensurate with the requisite experience.
     - A tiered risk assessment structure should be incorporated into the work of this independent group.
     - Such a group should be well-resourced and have members with appropriate expertise considering the types of risks and models being assessed.
4) NSCEB: Publish standards for potentially harmful algorithms.
   - **Response:** The federal government should monitor and review potentially harmful algorithms prior to the publication stage and develop clear pathways for preventing public communication of results that will lead to new pandemic risks.
5) NSCEB: Establish a national network of cloud labs.
   - **Response:** Strong requirements should be set regarding verification of samples, user access and completed experiments logs, know-your-customer regimes, and other relevant risk assessment/mitigation and security mechanisms prior to, or at least alongside, the establishment of a national network for cloud labs that would serve as a bridge from digital-to-physical transition of AIxBio model outputs with pandemic risks.
6) NSCEB: Collect and standardize common biological data types.
   - **Response:** Careful consideration and attention to governance policy should be given to specific subsets of newly generated biological datasets (eg, generated by automated laboratories or computational methods) that pose pandemic risks as capabilities scale.
7) Additional Recommendation: Require mandatory evaluations and red teaming for, at minimum, high-consequence biological risks to include consideration of how best to implement mitigating measures.
8) Additional Recommendation: Require gene synthesis providers and manufacturers to screen all customers and incoming orders of gene sequences and to require all purchasers of gene sequences to order only from providers and manufacturers who perform such screening.

Detailed Responses to AIxBio White Paper 4: Policy Options for AIxBio

As noted above, we strongly commend the NSCEB for the inclusion of each of the recommendations

below in its White Paper on Policy Options for AIxBio. Below, we respond to each recommendation and add considerations that we strongly encourage the NSCEB to adopt for its final report to fully harness the benefits of AIxBio technologies as well as to improve their safety and security.

Recommendation to develop an AIxBio consortium

We strongly agree with the NSCEB that an AIxBio consortium of stakeholders from government, industry, and academia should be developed by an agency such as NIST to share best practices, provide a comprehensive understanding of which groups are funding AIxBio research and development, and increase access to critical data resources related to AIxBio.[11] The need for the development of such a consortium was one of the strongest takeaways from a November 2023 meeting CHS convened among leading AI labs, executive branch officials, and biosecurity experts. From that meeting, we saw a compelling need for a public-private forum to facilitate the sharing of important information related to biosecurity risks and red-teaming results.[12] Accordingly, we recommend any such AIxBio consortium should prioritize information sharing about how to identify and mitigate high-consequence biosecurity risks, prevent the public release of model-generated information that increases such risks, and prevent their digital-to-physical transition of pandemic risks. In addition to sharing best practices, the AIxBio consortium should prioritize providing a comprehensive understanding of which groups are funding AIxBio research and development and prioritize increasing access to critical data resources related to AIxBio.

This consortium should also have a broad set of stakeholders from the fields of biosecurity, AI, cybersecurity, public health, national security, law, social sciences, and economics in order to best provide a comprehensive assessment of the risks and benefits of these technologies. The consortium should also strive to ensure diverse representation of stakeholders in the field of AIxBio and balance possible conflicts of interest by participation from companies working in the field of AIxBio.

Recommendation to initiate a global competitive analysis focused on AIxBio

The NSCEB suggested that the executive branch could establish an office to conduct a competitive analysis to assess the state of biotechnology infrastructure and technological advancement in the United States, compared to our strategic adversaries, with a focus on AIxBio.[13] We agree with the NSCEB that the federal government should aim to stay aware of the global competitive landscape of AIxBio. However, AIxBio oversight policies should not be primarily calibrated as a result of competitive analyses.[14] AIxBio oversight and governance policies should also be developed with the principal focus on protecting the public from the potential consequences of AIxBio technologies capable of generating outputs that increase the highest-consequence risks, including, as a highest priority, pandemic risks. Americans rely on the federal government to have such policies in place. Additionally, an important marker of a successful governance policy for AIxBio is that it is also making space for a safe and robust biotech ecosystem.

Recommendation to establish an independent group to conduct flexible risk assessments

[11] *See White Paper* at 1.
[12] *Johns Hopkins Center for Health Security Publishes Key Takeaways from its Meeting on the Convergence of AI and Biotechnology*, CTR. HEALTH SEC. (Dec. 19, 2023), https://centerforhealthsecurity.org/2023/johns-hopkins-center-for-health-security-publishes-key-takeaways-from-its-meeting-on-the-convergence-of-ai-and-biotechnology.
[13] *White Paper* at 2.
[14] *See White Paper* at 2.

We strongly agree with the NSCEB on the need for independent risk assessments that would be both robust and flexible based on the type of AI model and the specific bio-related risks.[15] However, for risk assessments to translate to meaningful risk reduction, further clarity is needed. Before establishing an independent group, the NSCEB should make clear that such a group will be tasked with, among other things, identifying the highest priority risks, defining the models that fall within scope, and establishing processes for (i) assessing models for such risks, (ii) preventing such risks from becoming fully open-source, and (iii) preventing the digital-to-physical transition of pandemic risks. Knowing these duties before such a group is established ensures that the group's composition/membership is commensurate with the requisite experience.

AIxBio tools should be able to undergo rapid risk assessments so that these assessments do not constitute undue burden or stifle innovation. However, specific subsets of technologies should undergo more rigorous risk assessment. For example, tools, methods, or approaches that shorten the timeline for, lower the costs of, or decrease the required sophistication for de novo synthesis of pandemic pathogens, or that ease the engineering of pathogens with new or enhanced pandemic properties, should undergo rigorous risk assessments. Accordingly, a tiered risk assessment structure should be incorporated into the work of such an independent group.

Additionally, risk assessment staff and infrastructure must be well resourced so that assessments can occur quickly and so that safe work can be accelerated. These staff should also develop, or work in concert with, others who are developing guidelines for AIxBio developers and users. This will better enable these staff to recognize and mitigate dangerous capabilities and safely limit access to models with dangerous capabilities.

Recommendation to publish standards for potentially harmful algorithms

We strongly agree with the NSCEB that guidance on publishing the computer code from models that have the potential to be used for harm should be developed, as well as guidance on guardrails for publicly available AI models that interact with biotechnology.[16] The federal government, in collaboration with a broad set of stakeholders such as the National Academies of Sciences, Engineering, and Medicine and the National Security Advisory Board for Biosecurity, has previously developed guidance for publishing DURC in the life sciences as well as ePPP research. However, placing a large emphasis on what to publish or not publish will be a failed strategy for numerous reasons—publication occurs as a final step in the technology development process. By the time publication is pursued, dangerous information has been shared across collaborators and devices, often without security protocols in place. Past experience has shown us that publishers do not wish to be put in the position of making decisions regarding protecting potentially hazardous information and have frequently asked for these decisions to be made by the government during earlier stages of the research. As has been recommended for risky research in DURC and ePPP governance, the federal government should monitor and review potentially harmful algorithms prior to the work being funded or initiated. If, despite such early assessments, such work is still not identified until the publication stage, clear pathways should be established regarding

[15] *White Paper* at 2.
[16] *White Paper* at 2.

how to prevent open-source release of AIxBio model outputs that generate new pandemic risks.[17]

Recommendation to establish a national network of cloud labs

We agree with NSCEB's recommendation that the National Science Foundation and Department of Energy (DOE) establish a network of cloud labs across the country where experimental instrumentation for chemistry and biology could be accessed by external researchers.[18] Cloud and automated laboratories provide efficiency gains for researchers and help address capacity constraints with regard to skilled laboratory staff. Cloud labs, whether governmental or private, should be required to provide a safe and secure laboratory venue. Accordingly, safety and security standards and requirements unique to cloud labs should be developed and implemented for this novel class of laboratory infrastructure that poses new security challenges.

Cloud labs reduce the skill required to conduct scientific experiments, and thus could be misused by bad actors or those trying to subvert other processes or controls. Automated laboratories further reduce the cost and skill required to generate large amounts of biological data, which could be used to train AI models for misuse. Companies are developing integrated tools such as "copilots" which would allow for the users to easily program cloud and/or automated laboratories using natural language.[19] Biological samples are shipped directly to such labs, and currently there are limited ways to verify the true contents of samples/reagents.

There are already safety and security oversight concerns regarding cloud and automated laboratories, but a larger national cloud lab network would greatly expand access to wet lab facilities and would come with new urgency to have strong governance systems in place. Wet laboratories are crucial for validated AI-driven in silico predictions against real-world biological phenomena. The ability to build in silico models that reliably model complex biology will be an inflection point in our ability to engineer biology for the betterment of humanity. However, wet lab validation of models that would allow for the accurate prediction of pandemic pathogen characteristics such as transmissibility, virulence, and immune evasion should only proceed when a strong governance policy has been established that would protect against pandemic risks.

Overall, strong requirements and governance should be set regarding verification of samples, logging of user access and experiments completed, know-your-customer regimes, and other relevant risk assessment/mitigation and security mechanisms prior to the establishment of a national network for cloud labs that would serve as a bridge from digital-to-physical transition of AIxBio model outputs with pandemic risks.

Recommendation to collect and standardize common biological data types

The NSCEB suggested that Congress or the President could establish a central office that would require

---

[17] *Center for Health Security Faculty Respond to White House Office of Science and Technology Policy RFI on Dual Use Research of Concern and Potential Pandemic Pathogen Care and Oversight Policy Framework*, CTR. HEALTH SEC. 5 (Oct. 16, 2023), https://centerforhealthsecurity.org/sites/default/files/2023-02/220629-recstostrengthenusgepppanddurcpolicies.pdf.

[18] *White Paper* at 1.

[19] *See, eg,* Ryan Heath, *AI Copilots and Robo-labs Turbocharge Research*, AXIOS, Jan. 9, 2024, https://www.axios.com/2024/01/09/ai-copilots-cloud-labs-science-research.

agencies to coordinate, collect, manage, and store high-quality biological data and encourage wider data availability while protecting privacy and ensuring data security.[20] We agree with the NSCEB that the success of AI methods depends on both the development of strong learning algorithms as well as the presence of useful datasets on which to train them, and computational resources to perform the training. Ensuring data are available and standardized for research and industry applications is desirable, as insufficient training data (whether in availability, size, quality, relevance, or other forms of bias) will prevent AIxBio technologies from being accurate or effective. Collecting and standardizing data will require additional funding support but is likely to unlock subsequent economic gains.

While recognizing that automated lab capabilities are not yet operating at this level, we recommend that careful attention to governance policy be given toward specific subsets of newly generated biological datasets (eg, generated by automated laboratories or computational methods) that pose pandemic risks as capabilities scale. New technologies, including automated labs and computational methods, may allow for the rapid and scalable generation of new biological datasets. For example, in the future, one could potentially program an automated lab to run thousands of experiments that generate vast amounts of structured and useable biological data. We must ensure that bad actors are not able to use automated labs and computational methods to generate new, substantial datasets that increase pandemic risk, such as data on transmissibility and immune evasion. Controlled access should be established for these model outputs that increase pandemic risks, as also recommended as a consideration for Congress by a Congressional Research Service report.[21]

Additionally, recent AIxBio research[22] indicates that, at least for a subset of biological design tools (BDTs), the relationship between dataset size and model performance is not only a matter of scale but also may be linked to model size and training strategies. Optimal dataset and strategic training approaches, including model sizes, may play a critical role in advancing superior performance across tasks. Further empirical research is required to identify the specific conditions and tasks for which data scale, paired with training strategies, confer a computational advantage for targeted biological functions.

Recommendation to establish an international working group focused on AIxBio

We understand the sentiment behind the NSCEB's recommendation that Congress could pass legislation to create an AIxBio working group within a multilateral body such as Five Eyes (Australia, Canada, New Zealand, the United Kingdom, and the United States).[23] Given that AI and biotechnology are being pursued globally, international cooperation will be crucial to ensure these technologies are developed safely and securely. While intelligence sharing with trusted allies should be encouraged, intelligence gathering and sharing alone should not be considered a major mitigation strategy for reducing AIxBio risks. The ultimate goals regarding AIxBio safety should include the establishment of international norms, standards, and requirements that identify and prevent high-consequence risks from emerging in countries around the world. Additionally, because the biosecurity implications of AIxBio research are

[20] *White Paper* at 1.
[21] TODD KUIKEN, CONG. RESEARCH SERV., R47849, ARTIFICIAL INTELLIGENCE IN THE BIOLOGICAL SCIENCES: USES, SAFETY, SECURITY, AND OVERSIGHT 17 (2023), https://crsreports.congress.gov/product/pdf/R/R47849.
[22] Francesca-Zhoufan Li et al., *Feature Reuse and Scaling: Understanding Transfer Learning with Protein Language Models,* BIORXIV (2024), https://www.biorxiv.org/content/10.1101/2024.02.05.578959v1.full.pdf.
[23] *Working Paper* at 1.

not confined to the US or their allies, the US should consider meaningful participation in an international AIxBio Forum beyond the Five Eyes to share best practices for implementing effective AIxBio guardrails, identify emerging biological risks associated with ongoing AI advances, and prompt the development of new tools and sharing of existing practices to manage these risks.[24]

Recommendation to dedicate high-performance computing capabilities to AIxBio

We agree with the NSCEB that agencies such as the DOE have existing high-performance computing available to them and that broadening access to such compute to additional researchers in academia and industry would be beneficial.[25] The scale of computing power needed for state-of-the-art algorithms is becoming prohibitive for smaller organizations, such as most academic labs or early-stage startups. We also recognize that limited access to advanced computing for academic and research groups could concentrate progress in the hands of industrial entities capable of heavy investment in technology. However, we strongly recommend that access to high-performance computing should be tied to compliance with safety and security measures, as has been the case with the Five Safes framework recommended by the National AI Research Resource (NAIRR) Task Force for its Implementation Plan[26]: safe projects, safe people, safe settings, safe data, and safe outputs.

Recommendation to launch a National AI Research Resource for Biotechnology

As with the previous recommendation, we agree with the NSCEB that broadening access to compute for academics and smaller industry players strengthens our nation's ability to innovate in biotechnology.[27] We are excited to see that the NSCEB describes this program as providing a "safe computational environment." In order to ensure this program provides value, safety and security standards should be established in advance with input from cybersecurity, biosecurity, and AI professionals, and we recommend here again the Five Safes framework recommended by the NAIRR Task Force for its Implementation Plan[28]: safe projects, safe people, safe settings, safe data, and safe outputs. We further recommend that any such Research Resource be established with one goal being the advancement of trustworthy biotechnology in the same way that one of the NAIRR goals is to advance trustworthy AI.[29]

Recommendation to establish an AI and biotechnology sandbox

We support NSCEB's recommendation to explore the creation of an AIxBio sandbox focused on the development of near-term use cases and pilot demonstrations of AI toward biotechnology for national security applications.[30] It is a strategic initiative that could hold several benefits, including speeding up innovation efforts while mitigating the risks linked with implementing unproven technologies and

---

[24] *See* Carter et al., *The Convergence of Artificial Intelligence and the Life Sciences: Safeguarding Technology, Rethinking Governance, and Preventing Catastrophe*, NTI 47 (Oct. 2023), https://www.nti.org/wp-content/uploads/2023/10/NTIBIO_AI_FINAL.pdf.

[25] *Working Paper* at 1.

[26] National Artificial Intelligence Research Resource Task Force, *Strengthening and Democratizing the US Artificial Intelligence Innovation Ecosystem: An Implementation Plan for a National Artificial Intelligence Research Resource 26–7* (Jan. 2023), https://www.ai.gov/wp-content/uploads/2023/01/NAIRR-TF-Final-Report-2023.pdf.

[27] *See White Paper* at 1.

[28] National Artificial Intelligence Research Resource Task Force, *Strengthening and Democratizing the US Artificial Intelligence Innovation Ecosystem: An Implementation Plan for a National Artificial Intelligence Research Resource 26–7* (Jan. 2023), https://www.ai.gov/wp-content/uploads/2023/01/NAIRR-TF-Final-Report-2023.pdf.

[29] *Id.* at v.

[30] *White Paper* at 2.

allowing setbacks to become opportunities for learning instead of expensive errors. A sandbox for national security applications may also enable a more secure environment to investigate the ethical and regulatory dimensions of novel technologies. Policymakers, by monitoring the practical use of these new tools, can develop well-informed regulations that promote safety and privacy while still promoting innovation.

Recommendation to initiate a regular intelligence assessment of future AIxBio risks and possible countermeasures

We support the NSCEB's recommendation that the Intelligence Community conduct annual assessments of emerging possible threats and countermeasures and were pleased to see the NSCEB's intention that such assessments could provide justification for the establishment of different oversight considerations for AIxBio.[31] As mentioned above, while intelligence assessments are useful, they alone do not reduce risk. Risk mitigation efforts must be prioritized alongside intelligence assessments.

Additional recommendations

Beyond the above recommendations that NSCEB included in its White Paper on AIxBio Policy Options, we suggest that the NSCEB evaluate these additional recommendations for inclusion in its final report:

- **Require mandatory evaluations and red teaming for, *at minimum*, high-consequence biological risks**: Require entities developing models with significant dual-use risks to red team and evaluate their models, and task an agency with: (1) auditing those models; and (2) submitting a report to Congress with recommendations for new authorities that will be needed by the agency to take any appropriate remedial action should red teaming, evaluations, or audits uncover vulnerabilities.[32] In setting requirements for the evaluation and red teaming of such models, we strongly recommend that, *at a minimum*, the federal government prioritize preventing the two classes of pandemic risks we have highlighted in this response. As NIST is currently developing guidelines for AI developers, deployers, and other actors to recognize when and how to mitigate dangerous capabilities identified through evaluations, it's unclear whether any US government agency would have the authority to audit those models, by which we mean the assessment of developers' red-teaming efforts as well as an evaluation of frontier models by the government itself. Nor is it clear by what authority the US government could take remedial action should its evaluation, or that of the developers, find a model dangerous.

- **Require mandatory gene synthesis screening:** Require gene synthesis providers and manufacturers to screen all customers and incoming orders of gene sequences and to require all purchasers of gene sequences to order only from providers and manufacturers who perform such screening. We were pleased to see the NSCEB Interim Report's plan to assess options for codifying best practices for DNA synthesis screening.[33] We encourage NSCEB to continue to

[31] *See White Paper* at 2.
[32] *Avoiding a Cautionary Tale: Policy Considerations for Artificial Intelligence in Health Care: Hearing Before the S. Subcomm. on Primary Health & Ret. Sec., Comm. On Health, Educ., Labor and Pensions*, 118th Cong. 8–9 (2023) (Statement of Dr. Tom Inglesby, Director, Johns Hopkins Center for Health Security), https://www.help.senate.gov/imo/media/doc/79536a31-d1cf-25b0-e526-52ba2193d900/Tom%20Inglesby,%20Nov.%208%20HELP%20Subcommittee%20Written%20Testimony.pdf.
[33] *Interim Report* at 32.

engage with the biosecurity expertsduring such assessment, in addition to private and government stakeholders responsible for implementing the best practices.

- **Assess strategies to manage access to biological design tools (BDTs) with high-consequence biosecurity risks.** Many strategies for safeguarding AI models depend on managing, overseeing, and limiting access to highly capable models, and such strategies should be considered for those BDTs with specific model capabilities that could increase high-consequence biosecurity risks. Many LLM developers already employ application programming interfaces (APIs) to maintain control of their models; however, some LLMs and many current BDTs make their model weights widely available. BDTs are often developed by academic scientists in collaborative groups, who consider open sharing of resources an important norm for scientific advancement. Funders, developers, and distributors of AI models that could pose high-consequence biological risks should work with biosecurity experts, the research community, and appropriate government programs to critically evaluate open-source norms for BDTs with specific model capabilities that could increase high-consequence biosecurity risk.[34]

- **Ensure the US has a strong and robust AIxBio workforce by investing in education and training, especially in red teaming:** Establish specialized education programs at universities and vocational schools that combine AI and biotechnology, offering degrees and certifications. Such programs are critical for cultivating a robust AIxBio workforce capable of driving innovation and addressing future challenges.[35] This strategy must be complemented by policies that encourage high-skilled immigration, as there are significant challenges in talent attraction and retention within the national security innovation base.[36] To assist with this, Congress should liberalize immigration policies for STEM professionals. Specifically, increasing the number of green cards for applicants with STEM PhDs in critical and emerging technologies is crucial.[37] This will not only help alleviate workforce shortages but also ensure the US remains competitive and a global leader in these pivotal fields by attracting and retaining the world's most talented individuals. Importantly, it seems there is a workforce shortage of individuals capable of red teaming leading AI models for biosecurity threats. The NSCEB should consider how additional funding for NIST could help alleviate some of these shortages, as well as support specialized education programs at universities and vocational schools.

[34] Carter et al., *supra* note 24 at 49.
[35] *See, eg,* US National Security Commission on Artificial Intelligence, *Final Report* (Oct. 2021), https://reports.nscai.gov/final-report/ (detailing potentially applicable talent and workforce-related recommendations from the National Security Commission on Artificial Intelligence's final report).
[36] *National Security Innovation Base Report Card,* RONALD REAGAN PRESIDENTIAL FOUNDATION AND INSTITUTE (Mar. 2024), https://www.reaganfoundation.org/media/362366/2024-nsib-report-card.pdf.
[37] Dr. Mark J. Lewis and Divyansh Kaushik*, High Skills Immigration Is a National Security Issue*, NATIONAL DEFENSE (Aug.19, 2022), https://www.nationaldefensemagazine.org/articles/2022/8/19/high-skills-immigration-a-national-security-issue; *See, eg,* Elliott Abrams et al., National Security Leaders' Letter to House Select Committee on CCP (May 2023), https://www.documentcloud.org/documents/23813309-national-security-leaders-letter-to-house-select-committee-on-ccp (showing over five dozen national security leaders urging the House Select Committee on the Strategic Competition between the United States and the Chinese Communist Party to address immigration challenges).