Response to NIST Request for Information (RFI) on
Safety Considerations for Chemical and/or Biological AI Models
[Docket No. 240920-0247]

December 3, 2024
U.S. AI Safety Institute
National Institute of Standards and Technology
U.S. Department of Commerce
1401 Constitution Avenue NW,
Washington, DC 20230

We are grateful for the opportunity to respond to this Request for Information (RFI).

This comment was prepared by individuals associated with Rosetta Commons and supported by written feedback from researchers in this community and the broader biomolecular structure analysis, prediction, and design communities.

Rosetta Commons is a consortium of developers and scientists who use or contribute to the Rosetta software suite, a set of computational tools used for predicting and designing biomolecules, which has enabled notable scientific advances in computational biology. Recently, the Rosetta Commons has been home to the development of other biomolecular prediction and design tools, including RoseTTAFold, RFdiffusion, and ProteinMPNN. Founded in 2003 by Prof. David Baker, a 2024 Nobel laureate in chemistry, as well as alumni from his lab, the consortium has grown to include hundreds of developers and scientists from over 100 academic and private sector laboratories (see Appendix A).

This comment was supplemented by perspectives shared during a biosecurity-oriented roundtable discussion at Summer RosettaCon in Washington, USA, in August 2024, a workshop related to this RFI conducted at European RosettaCon in Copenhagen, Denmark, in November 2024, as well as a survey administered in person at European RosettaCon and then virtually with the broader Rosetta Commons community (additional details in Appendix B).

This submission has been endorsed by Dr. Jeffrey J. Gray, Director of Rosetta Commons and Principal Investigator of the GrayLab at Johns Hopkins University. However, opinions expressed herein are not official positions of Johns Hopkins University, Rosetta Commons (or its signatory institutions), nor do they necessarily represent those of other members of Rosetta Commons.

We thank you for your consideration of our comment.

Samuel Curtis, MSc
Biosecurity Fellow, Rosetta Commons
Policy Researcher, Open Molecular Software Foundation
Visiting Scholar, Johns Hopkins University

Stephen McCarthy, PhD
Postdoctoral Scholar, University of California, Irvine

Bryce Johnson
PhD Candidate, University of Wisconsin-Madison

Rocco Moretti, PhD
Research Professor, Vanderbilt University

Affiliations are for identification only and do not imply any institutional endorsement.

Context

The ability to make accurate computational predictions of biomolecular structures and to design new functional proteins are goals—once aspirational—that have motivated decades of research.

Computational methods have been central to the success of protein design almost since its inception, exemplified by the development of Rosetta, which employs an understanding of the energetics underlying protein folding, as well as structure and sequence data, to predict protein structures and design new proteins.[1] Significant milestones on the path toward protein design include the first design of a small-molecule binding peptide in 1983,[2] a self-assembling membrane-spanning ion channel in 1988,[3] and the first *de novo* designed protein with a fold unseen in nature in 2003.[4]

Newer methods built on machine learning are rapidly increasing protein design capabilities,[5] and the adoption of platforms that reduce barriers to the sharing of code, model weights, and data—such as GitHub and HuggingFace—has not only accelerated software and methods development but also made these tools accessible to more researchers across a broader range of research disciplines, including those without specialized computational expertise. Hosting protein design tools on web servers, either standalone or through platforms like Google Colab, has further helped the tools reach a larger pool of users by facilitating access to required computational resources.

We anticipate that the scope and capabilities of chem-bio AI models will continue to grow. Data is the lifeblood of AI model training, and the availability of large, high-quality, and curated datasets is currently a limiting factor in the development of more powerful models.[6] While more expressive neural network architectures can help maximize the utility of existing data, automated and high-throughput experimental methods remain an important strategy to overcome this limitation, as are concerted efforts in the research community to collate, curate, and share data suitable for model training (such as the Rosetta Data Bazaar[7]).

---

[1] Kristian W. Kaufmann et al., *Practically useful: What the Rosettaprotein modeling suite can do for you*, 49 Biochemistry 2987–2998 (2010), https://doi.org/10.1021/bi902153g.
[2] Rudolf Moser, Richard M. Thomas & Bernd Gutte, *An artificial crystalline ddt-binding polypeptide*, 157 FEBS Letters 247–251 (1983), https://doi.org/10.1016/0014-5793(83)80555-9.
[3] J. D. Lear, Z. R. Wasserman & W. F. DeGrado, *Synthetic amphiphilic peptide models for protein ion channels*, 240 Science 1177–1181 (1988), https://doi.org/10.1126/science.2453923.
[4] Brian Kuhlman et al., *Design of a novel globular protein fold with atomic-level accuracy*, 302 Science 1364–1368 (2003), https://doi.org/10.1126/science.1089427.
[5] Tanja Kortemme, *De novo protein design—from new structures to programmable functions*, 187 Cell 526–544 (2024), https://doi.org/10.1016/j.cell.2023.12.028.
[6] Francesca-Zhoufan Li et al., *Feature reuse and scaling: Understanding transfer learning with protein language models*, bioRxiv (2024), https://doi.org/10.1101/2024.02.05.578959; Filomeno Sánchez Rodríguez et al., *Using deep-learning predictions reveals a large number of register errors in PDB depositions*, 11 IUCrJ 938–950 (2024), https://doi.org/10.1107/s2052252524009114; Frances Ding & Jacob Steinhardt, *Protein language models are biased by unequal sequence sampling across the tree of life*, bioRxiv (2024), https://doi.org/10.1101/2024.03.07.584001.
[7] Rosetta Data Bazaar, Hugging Face, https://huggingface.co/collections/RosettaCommons/rosetta-data-bazaar-66b6388a83cbf76a213c5f78.

The increasing accuracy of protein design and modeling tools, as well as the ease with which they can be distributed and used by people without specialist training in computational biology, offers the opportunity to realize the potential of the field and make tangible the advances in human health and flourishing—curing diseases, unlocking sustainable energy, tackling pollution, and addressing climate change—envisaged by founders of the field decades ago. At the same time, we recognize that advances in these technologies, as with any biotechnology, are accompanied by the risk that they could be used deliberately or accidentally in a manner that threatens public health or the environment.

Various strategies to mitigate these risks have been proposed, including monitoring the development of chem-bio AI models that surpass some computational power threshold for training,[8] restricting the amount or type of data that can be used to train models,[9] introducing 'refusals' to generate answers to certain requests,[10] limiting access to the trained models,[11] and implementing additional institutional oversight of chem-bio AI model development.[12]

We expect there would not be consensus within research communities on the most effective risk mitigation approach. Responses to the survey we conducted indicated a broad range of opinions in this community on questions related to the presence of risk and risk mitigation (e.g., Figure 3 and Figure 6). Furthermore, when responding to open-ended survey questions, some expressed concern that risk mitigation efforts could unduly constrain computational biological research.

Protein structure prediction and design research also play a critical role in enhancing public health security through, for example, improving our ability to prevent, detect, and respond to pandemics. During the COVID-19 pandemic, the first vaccine using computationally designed proteins was approved for use by major governments—South Korea and the United Kingdom—and even outperformed AstraZeneca's vaccine in clinical tests.[13] The recently-developed deep learning model EVEscape predicts future mutations in viral sequences significantly quicker, cheaper, and more accurately than wet lab methods, and may have allowed prediction of emerging SARS-CoV-2 variants ahead of time.[14] From now into the future, protein structure prediction and design tools will serve as strategic assets in pandemic response by

[8] Executive order on the safe, secure, and trustworthy development and use of artificial intelligence, The White House (2023), https://www.whitehouse.gov/briefing-room/presidential-actions/2023/10/30/executive-order-on-the-safe-secure-and-trustworthy-development-and-use-of-artificial-intelligence/.
[9] Sarah R. Carter et al., *Developing Guardrails for AI Biodesign Tools*, NTI|bio (2024), https://www.nti.org/wp-content/uploads/2024/11/NTIBio_Paper_Developing-Guardrails-for-AI-Biodesign-Tools_FINAL.pdf.
[10] *Id.*
[11] *Id.*
[12] Doni Bloomfield et al., *Ai and biosecurity: The need for governance*, 385 Science 831–833 (2024), https://doi.org/10.1126/science.adq1977.
[13] *Covid-19 vaccine with IPD nanoparticles wins full approval abroad - institute for protein design*, Institute for Protein Design (2024), https://www.ipd.uw.edu/2022/06/covid-19-vaccine-skycovione-wins-full-approval-abroad/.
[14] Nicole N. Thadani et al., *Learning from prepandemic data to forecast viral escape*, 622 Nature 818–825 (2023), https://doi.org/10.1038/s41586-023-06617-0.

enabling rapid pathogen characterization, accelerated diagnostics, and vaccine and therapeutic development.[15]

Many in our research community are familiar with the concept of biosecurity (Figure 2) and recognize its importance; many were involved in the drafting of and were signatories to the "Community Values, Guiding Principles, and Commitments for the Responsible Development of AI for Protein Design."[16] Some have expressed concern that claims of potential biosecurity risks have been used and could continue to be used as a pretext to delay or withhold from sharing a model or method in order to advance a commercial interest. Some expressed skepticism that Google DeepMind's decision to withhold the weights of AlphaMissense to "prevent their use in potentially unsafe applications," (as stated in the accompanying publication) was the sole reason for the lack of disclosure.[17] Some have expressed concern that when journals publish studies without code or model weights, as *Nature* did with AlphaFold3, they effectively serve as a platform for corporate publicity while undermining scientific standards of reproducibility.[18] These concerns were reinforced when, following an open letter about transparency and reproducibility, an employee of Google DeepMind announced that AlphaFold3's weights would be released within six months,[19] raising questions about the original justification for withholding them.

Sentiments of this nature were reflected in our survey: Roughly half of the respondents had encountered at least one published article with safety or security implications that concerned them (see Figure 3), opinions were broadly distributed across response options when asked whether or not restrictions on protein design tools or methods on the basis of safety or security is sometimes justified (Figure 6), and most indicated disbelief in the credibility of the justifications that have been used to limit access to protein design tools or methods on the basis of safety or security concerns (Figure 7).

Moving forward, transparency in decision-making around biosecurity will be essential for building trust within the research community. Approaches to mitigating identified risk should meaningfully consider the practical challenges of weaponizing biological agents and the capabilities and limitations of plausible threat actors, and measures to mitigate risk should be developed in consultation with scientists to ensure they do not unduly hinder research.

The remainder of this comment is divided into three sections, as follows:

1. Assessing dual-use capabilities and mitigating risk of misuse of chem-bio AI models
2. Opportunities to enhance nucleic acid synthesis screening
3. Future safety and security of chem-bio AI models

---

[15] Lynda M. Stuart, Rick A. Bright & Eric Horvitz, *AI-enabled Protein design: A strategic asset for Global Health and Biosecurity*, National Academy of Medicine (2024), https://nam.edu/ai-enabled-protein-design-a-strategic-asset-for-global-health-and-biosecurity/.

[16] Community Values, Guiding Principles, and Commitments for the Responsible Development of AI for Protein Design (2024), https://responsiblebiodesign.ai/.

[17] Jun Cheng et al., *Accurate proteome-wide missense variant effect prediction with Alphamissense*, 381 Science (2023). https://doi.org/10.1126/science.adg7492.; Anthony Gitter et al., *A renewed call for Open Artificial Intelligence in biomedicine*, OSFPreprints (2024), https://doi.org/10.31219/osf.io/2xh3w.

[18] *Id.*

[19] Pushmeet Kohli. X (formerly Twitter) (2024). https://x.com/pushmeet/status/1790086453520691657.

Of the examples of chem-bio AI models listed in this Request for Information, *protein design tools* and *small biomolecule design tools* are most relevant to research efforts in the Rosetta Commons community, and our comments in this response apply principally to these categories, although they may also be relevant to others.

Assessing dual-use capabilities and mitigating risk of misuse of chem-bio AI models

The content of this section concerns Questions 1a-f, 2a-e, and 3a-c of the RFI.

Observations

Evaluating chem-bio AI models for risk is complicated by their inherent dual-use nature. Any capability that a researcher might want to assess for risk most likely also confers some benefit to legitimate research and medical applications. For example, a model's ability to predict ACE2 receptor binding affinity could enable both the prediction of dangerous SARS-CoV-2 variants and the development of life-saving therapeutics. This duality extends across many critical therapeutic areas: cancer treatments rely on precisely targeting and destroying specific human cells, vaccine development often involves stabilizing viral proteins to create better immunogens, and gene therapy requires modifying viruses to target particular cell types—all capabilities that could be concerning in non-therapeutic contexts. This intrinsic overlap between beneficial and potentially harmful applications creates fundamental challenges in risk evaluation. Adding to the complexity, it is often unclear whether any given capability is more likely to be used for harm than for benefit—an ambiguity stemming directly from the dual-use nature of these technologies.

Conventional language model evaluation approaches based on refusing attempts to elicit information are not apt for chem-bio AI models. Large (natural) language models (hereafter, "LLMs") can be evaluated, in part, by testing their ability to withhold specific information—measuring how successfully they resist attempts to elicit harmful content through various prompting strategies.[20] These evaluations work in part because language models can potentially detect harmful intent through linguistic patterns in user prompts. However, this evaluation paradigm is not suitable for chem-bio AI models for two key reasons. First, these models operate on abstract representations—e.g., polymer sequences and atomic coordinates—where user intent cannot be reliably discerned from the inputs alone. A sequence of amino acids or a set of structural coordinates carries no inherent signal about the intended use of the output. Second, these models must accurately generate biological sequences, structures, or other properties to fulfill their scientific purpose—a protein structure prediction model that withholds certain structural predictions would be failing at its core scientific task. These properties, intrinsic to the models and the research for which they are created, indicate the need for evaluation frameworks that go beyond testing discretion in responding to user input.

A grounded understanding of plausible threat actors is essential for designing meaningful evaluations and appropriate mitigations. Preliminary analyses of biotechnology threats suggest that the successful development of novel biological agents requires both significant resources and deliberate harmful intent—factors that substantially narrow the field of plausible threat actors.[21] The level of risk posed by these actors, and therefore the nature and degree of mitigation that is warranted, depends on technical expertise, motivations and incentives, and access to computational, financial, and wet lab

[20] Mantas Mazeika et al., *Harmbench: A standardized evaluation framework for automated red teaming and robust refusal*, arXiv (2024), https://doi.org/10.48550/arXiv.2402.04249.
[21] Michael Montague, *Towards a grand unified threat model of Biotechnology*, PhilSci Archive (2023), https://philsci-archive.pitt.edu/id/eprint/22539.

resources, among other factors. Particularly important is understanding what would motivate actors to attempt to use chem-bio AI models as opposed to conventional approaches. It is beyond our purview to perform such an assessment, but it will be necessary for biosecurity evaluations to be designed with these considerations in mind. Given that different mitigation strategies can have varying impacts on legitimate research—from minimal disruption to significant barriers—security measures should be carefully targeted and proportional to demonstrated risks.

Chem-bio AI models cannot be evaluated with a one-size-fits-all method. Whereas LLMs, e.g., GPT-4, Gemini, and Claude, can be evaluated through standardized methods,[22] this approach is not feasible for all chem-bio AI models. As defined in this RFI, these models encompass a broad range of tools, each designed for specific scientific tasks and requiring distinct evaluation approaches. Consider the following examples:

- RFdiffusion, a diffusion-based generative neural network for the creation of novel 3D protein backbone structures, accepts various forms of conditioning information (e.g., partial sequences, fold specifications, binding hotspots, or symmetry constraints) and outputs 3D backbone structures in PDB format.[23]
- ProteinMPNN, a message-passing neural network for protein sequence design, accepts a 3D protein backbone structure in PDB format and outputs probable amino acid sequences that could fold into that structure along with sequence confidence scores.[24]
- AlphaFold2, a protein structure prediction model that uses a transformer-based neural network architecture, takes as input a FASTA file containing an amino acid sequence and outputs a detailed 3D structural prediction in PDB format, including atomic coordinates and confidence scores.[25]
- ESM3, a generative large language model pretrained on protein sequence, structure, and function datasets, can accept instructions from different input tracks (e.g., sequence, structure coordinates, or functional keywords), and simultaneously reason over the other tokens to produce a desired output (e.g., the sequence and predicted structure of a protein with a particular function). [26]

While performance benchmarks exist for specific classes of models (e.g., protein structure prediction), the diversity across chem-bio AI models—in their functions, architectures, inputs, and outputs—makes a unified evaluation framework impractical.

---

[22] Percy Liang et al., *Holistic evaluation of Language Models*, arXiv (2023), https://doi.org/10.48550/arXiv.2211.09110; Jon M. Laurent et al., *LAB-Bench: Measuring capabilities of language models for Biology Research*, arXiv (2024), https://doi.org/10.48550/arXiv.2407.10362; *Pre-deployment evaluation of Anthropic's upgraded Claude 3.5 Sonnet*, NIST (2024), https://www.nist.gov/news-events/news/2024/11/pre-deployment-evaluation-anthropics-upgraded-claude-35-sonnet.

[23] Joseph L. Watson et al., *De novo design of protein structure and function with rfdiffusion*, 620 Nature 1089–1100 (2023), https://doi.org/10.1038/s41586-023-06415-8.

[24] J. Dauparas et al., *Robust deep learning–based protein sequence design using ProteinMPNN*, 378 Science 49–56 (2022), https://doi.org/10.1126/science.add2187.

[25] John Jumper et al., *Highly accurate protein structure prediction with AlphaFold*, 596 Nature 583–589 (2021), https://doi.org/10.1038/s41586-021-03819-2.

[26] Alexander Rives et al., *Biological structure and function emerge from scaling unsupervised learning to 250 million protein sequences*, 118 Proceedings of the National Academy of Sciences (2021), https://doi.org/10.1073/pnas.2016239118.

Model pretraining computational utilization ("compute") is a flawed proxy for dual-use potential.[27] Policies intending to mitigate risk must balance addressing biological capabilities with the greatest potential for harm while avoiding unnecessary impediments to scientific research. Some existing policies and analyses propose regulatory requirements for chem-bio AI models that surpass a threshold of computational operations (or "compute threshold," measured in floating-point or integer operations) for pretraining.[28] This regulatory framework mirrors similar approaches taken toward LLMs.[29] As they pertain to LLMs, compute thresholds as a regulatory approach are controversial—some researchers argue that they are fundamentally flawed,[30] while others view it as a workable albeit suboptimal proxy for risk.[31] For chemical and biological models, however, compute is an especially flawed metric for assessing risk, as robust performance in relevant capabilities often depends on precise—rather than broad—biochemical understanding. Although biological language models have generally affirmed the scaling hypothesis (or "scaling laws"),[32] recent research has demonstrated exceptions—that medium-sized biological language models can perform comparably to or even outperform much larger models on protein engineering tasks, particularly when working with limited datasets.[33] Consider that one of the most effective current models for predicting viral protein variant effects is a non-generative AI model, GEMME, that can be trained on a typical laptop in minutes.[34] Similarly, ProteinMPNN has demonstrated superior performance in inverse protein design compared to many models with substantially more parameters and greater compute

---

[27] This stance reflects sentiments broadly expressed by participants of a biosecurity-oriented roundtable discussion held at Summer RosettaCon in August 2024 in Washington, USA.

[28] Executive order on the safe, secure, and trustworthy development and use of artificial intelligence, The White House (2023), https://www.whitehouse.gov/briefing-room/presidential-actions/2023/10/30/executive-order-on-the-safe-secure-and-trustworthy-development-and-use-of-artificial-intelligence/; Provisions Pertaining to U.S. Investments in Certain National Security Technologies and Products in Countries of Concern, 89 Fed. Reg. 55846 (July 5, 2024) (to be codified at 31 C.F.R. pt. 850); Doni Bloomfield et al., *AI and biosecurity: The need for governance*, 385 Science 831–833 (2024), https://doi.org/10.1126/science.adq1977.

[29] Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence and amending Regulations (EC) No 300/2008, (EU) No 167/2013, (EU) No 168/2013, (EU) 2018/858, (EU) 2018/1139 and (EU) 2019/2144 and Directives 2014/90/EU, (EU) 2016/797 and (EU) 2020/1828 (Artificial Intelligence Act), http://data.europa.eu/eli/reg/2024/1689/oj; Executive order on the safe, secure, and trustworthy development and use of artificial intelligence, The White House (2023), https://www.whitehouse.gov/briefing-room/presidential-actions/2023/10/30/executive-order-on-the-safe-secure-and-trustworthy-development-and-use-of-artificial-intelligence/

[30] Sara Hooker, *On the limitations of compute thresholds as a governance strategy*, arXiv (2024), https://doi.org/10.48550/arXiv.2407.05694.

[31] Girish Sastry et al., *Computing power and the governance of Artificial Intelligence*, arXiv (2024), https://doi.org/10.48550/arXiv.2402.08797.

[32] Jared Kaplan et al., *Scaling laws for neural language models*, arXiv (2020), https://doi.org/10.48550/arXiv.2001.08361; Eric Nguyen et al., *Sequence modeling and design from molecular to genome scale with Evo*, 386 Science (2024), https://doi.org/10.1126/science.ado9336; Thomas Hayes et al., Simulating 500 million years of evolution with a language model, bioarXiv (2024), https://doi.org/10.1101/2024.07.01.600583; Yeqing Lin et al., *Out of many, one: Designing and scaffolding proteins at the scale of the structural universe with Genie 2*, arXiv (2024), https://doi.org/10.48550/arXiv.2405.15489.

[33] Luiz C. Vieira, Morgan L. Handojo & Claus O. Wilke, *Scaling down for efficiency: Medium-sized transformer models for protein sequence transfer learning*, bioRxiv (2024), https://doi.org/10.1101/2024.11.22.624936.

[34] Elodie Laine, Yasaman Karami & Alessandra Carbone, *Gemme: A simple and fast global epistatic model predicting mutational effects*, 36 Molecular Biology and Evolution 2604–2619 (2019), https://doi.org/10.1093/molbev/msz179; Pascal Notin et al., *ProteinGym: Large-scale benchmarks for protein design and fitness prediction*, bioarXiv (2023), https://doi.org/10.1101/2023.12.07.570727.

demands.[35] These examples indicate the need to identify more appropriate risk factors (which might include examining data characteristics, model architecture, intended applications, and other relevant considerations), devise corresponding regulatory frameworks, and develop mechanisms by which these frameworks can adapt to advancements in technology and our understanding of risk.

The impact of LLMs on the risk landscape is uncertain. While LLMs have the possibility to be employed directly for dual-use purposes, they also have the potential to alter the risk profiles posed by other chem-bio AI tools. For example, they may be able to interpret natural language inputs and produce the scripts needed to run other protein design tools, or even create and run entire pipelines. This may increase the number of potential threat actors by lowering the technical skill requirements to perform modeling. LLMs can make novel inferences across their training data, but they are ultimately bounded by the information contained within that training data (or which exists on the internet, for systems with retrieval-augmented generation). As such, the ability of LLMs to assist potential threat actors in using chem-bio AI models nefariously is limited by the ease of use and the quality of documentation available. For example, researchers assessed the ability of GPT-4 to build a Rosetta protocol to create a protein-based drug targeting the SARS-CoV-2 spike protein.[36] While GPT-4 was able to provide a reasonable-looking protocol, the material provided for the most part regurgitated already existing tutorials (though with a number of mistakes), examples, and guidelines that are accessible online from the Rosetta documentation resources. Thus, while LLMs like GPT-4 may make accessing model documentation easier, it seems unlikely that they would significantly affect the threat profile of a knowledgeable or determined actor.

Examples

Our survey was accompanied by an open-ended question asking respondents how they had conducted evaluations relevant to biosecurity in their work. Some reported using methods to assess designed proteins individually, such as immunogenicity prediction, homolog comparison, allergen, or toxicity evaluations. No examples were given of researchers assessing chem-bio AI models themselves for risk.[37]

We are, however, aware of instances in which other researchers have made efforts to evaluate dual-use implications and/or mitigate potential risks of their models (in chronological order):

- Before pretraining Evo, a genomic foundation model, researchers at the Arc Institute removed sequences from viruses that infect eukaryotic hosts from the pretraining dataset.[38] After pretraining Evo, researchers engaged with experts in other research domains, including biomedical informatics and epidemiology, to identify the safety and ethical implications of the tool, including the potential for misuse, contribution to social and health inequity, and environmental disruption. Based upon these findings, researchers proposed several next steps in

[35] J. Dauparas et al., *Robust deep learning–based protein sequence design using ProteinMPNN*, 378 Science 49–56 (2022), https://doi.org/10.1126/science.add2187; Fei Ye et al., *Proteinbench: A holistic evaluation of Protein Foundation models*, arXiv (2024), https://doi.org/10.48550/arXiv.2409.06744.
[36] Microsoft Research AI4Science & Microsoft Azure Quantum, *The impact of large language models on scientific discovery: A preliminary study using GPT-4*, arXiv (2023), https://doi.org/10.48550/arXiv.2311.07361.
[37] The survey was not specifically targeted at *developers* of chem-bio AI tools; consequently, the participants may principally be *users* of these tools rather than involved in their creation.
[38] Eric Nguyen et al., *Sequence modeling and design from molecular to genome scale with Evo*, 386 Science (2024), https://doi.org/10.1126/science.ado9336.

their publication accompanying Evo, including the establishment of ethical guidelines, oversight mechanisms, the promotion of transparent technological use, forging community partnerships and international collaborations, investing in education and capacity building, and mechanisms to collect and integrate feedback from those involved in or impacted by Evo's applications.

- Prior to releasing AlphaFold3, a protein structure and interaction prediction tool, researchers at Google DeepMind conducted analyses to understand the tool's benefits and risks.[39] Described in a supplementary document to the tool's announcement, these analyses included: an ethics and safety assessment to identify and analyze potential risks and benefits, including their potential likelihood and impact; interviews with external experts in domains ranging from DNA synthesis to virology and national security to understand their view on the tool's benefits and risks; and comparisons of the output of AlphaFold3 to other existing resources (such as the Protein Data Bank). Google DeepMind furthermore reports to be using its AlphaFold Server as a testbed for the efficacy of screening and filtering processes, including by blocking a small number of viral protein sequences.[40]
- Researchers at EvolutionaryScale recently developed two variations of their latest, largest protein language model: ESM3, which has not been publicly released, and ESM3-open, an open-source version that incorporates several precautionary risk-mitigation measures. For the open-source version, researchers removed sequences unique to viruses from the pretraining dataset, as well as viral and non-viral sequences from the CDC and USDA's Select Agents and Toxins List, and sequences matching a subset of keywords associated with viruses and toxins. Researchers then performed evaluations to measure different aspects of the performance across versions of their ESM models: ESM3-open, ESM3 (which lacked any such sequence/keyword filtering), and ESM2 3B (a previous-generation model).[41] These evaluations included structure prediction, representation learning, function keyword prediction, and zero-shot viral fitness prediction, allowing researchers to see how dimensions of performance were affected by filtering sequences of potential concern.
- When publishing a preprint for Protein Set Transformer, a protein-based genome language model for viromics, researchers at the University of Wisconsin-Madison incorporated a risk-benefit analysis.[42] This analysis included a calculation of the percentage of viruses within the training set that infect humans and mammals, as well as the number of sequences on various watchlists—the CDC's list of bioterrorism agents and the National Respiratory and Enteric Virus Surveillance System. They also report having independent experts consider the impacts of this tool before releasing its code and model weights.

Thus, to our knowledge, there do not seem to be common methods by which chem-bio AI models are evaluated for risk. Approaches to evaluations have involved both qualitative (e.g., interviews with domain experts) as well as quantitative (e.g., comparing performances of viral fitness predictions) methods. None

[39] C. Griffin et al., *Our approach to biosecurity for AlphaFold 3,* Google DeepMind (2024), https://storage.googleapis.com/deepmind-media/DeepMind.com/Blog/alphafold-3-predicts-the-structure-and-interactions-of-all-lifes-molecules/Our-approach-to-biosecurity-for-AlphaFold-3-08052024.
[40] *Id.;* AlphaFold Server, https://alphafoldserver.com/welcome.
[41] Thomas Hayes et al., Simulating 500 million years of evolution with a language model, bioarXiv (2024), https://doi.org/10.1101/2024.07.01.600583.
[42] Cody Martin, Anthony Gitter & Karthik Anantharaman, *Protein set transformer: A protein-based genome language model to power high diversity viromics*, bioRxiv (2024), https://doi.org/10.1101/2024.07.26.605391.

have involved *in vitro* validation (which would be expensive and potentially dangerous). All of the evaluations appear to have been commenced after pretraining; in other words, we are not aware of efforts to evaluate the impact of the desired chem-bio model *ex-ante*. Furthermore, we are not aware of any approach consisting of predefined risk-tolerance thresholds, i.e., results that would indicate or demand some form of risk mitigation.

Recommendations

Recognizing the need to identify and mitigate biosecurity risks, as well as the fundamental importance of advancing research to enhance biosecurity, we believe the following should be considered when developing biosecurity evaluations and mitigations for chem-bio AI models:

Biosecurity evaluations and mitigations should measure and attend to biological capabilities that have the greatest capacity for harm. Prior research has sought to identify dual-use biological capabilities of concern in the life sciences broadly,[43] and correspondingly, brought about by AI models.[44] Steps should be taken to reach a scientific consensus on which capabilities of chem-bio AI models would be of greatest consequence, such as those that would enable the creation of novel or enhanced pandemic-capable pathogens. The National Science Advisory Board for Biosecurity (NSABB) has acknowledged the decreasing utility of "list-based" approaches to oversight,[45] and the National Science and Technology Council has adopted a progenitor-agnostic definition of "pathogen with enhanced pandemic potential."[46] Following these precedents, evaluations should not be grounded in a model's ability to make any *specific* pathogen variant, but rather on its capability to reliably confer *characteristics* that pose a significant threat to public health or the environment. A recent expert assessment of AI capabilities in biological research echoed this characteristics-based approach by focusing on concerning capabilities rather than specific pathogens, identifying key risks including those enabling immune evasion, enhanced transmission, automated synthesis, host adaptation, and predictive modeling of disease spread.[47]

Establishing robust biosecurity evaluations for *in silico* research should be a top priority. We acknowledge the need for a clear understanding of the risks associated with chem-bio AI models, and that mitigations should be grounded in empirical evidence. As we discuss later in this comment, we also recognize the lack of a shared understanding of what "biosecurity evaluations" should constitute, and of the data necessary to make appropriate, informed decisions with respect to biosecurity mitigations.

---

[43] National Academies of Sciences, Engineering, and Medicine, *Biodefense in the age of synthetic biology*, The National Academies Press (2018), https://doi.org/10.17226/24890.

[44] Jaspreet Pannu et al., *Prioritizing high-consequence biological capabilities in evaluations of Artificial Intelligence Models*, SSRN (2024), https://dx.doi.org/10.2139/ssrn.4873106.

[45] National Science Advisory Board for Biosecurity. *Proposed Biosecurity Oversight Framework for the Future of Science* (2023), https://osp.od.nih.gov/wp-content/uploads/2023/03/NSABB-Final-Report-Proposed-Biosecurity-Oversight-Framework-for-the-Future-of-Science.pdf.

[46] United States government policy for oversight of dual use research of concern and pathogens with enhanced pandemic potential, The White House (2024), https://www.whitehouse.gov/ostp/news-updates/2024/05/06/united-states-government-policy-for-oversight-of-dual-use-research-of-concern-and-pathogens-with-enhanced-pandemic-potential/.

[47] Jaspreet Pannu et al., *AI could pose pandemic-scale biosecurity risks. Here's how to make it safer*, 635 Nature 808–811 (2024), https://doi.org/10.1038/d41586-024-03815-2.

Therefore, establishing mechanisms for conducting evaluations should be an overarching priority, and we believe they should adhere to the following design principles:

- Evaluations should be grounded in well-defined threat models, assessing risks in the context of capabilities, resources, and motivations of plausible threat actors.
- Evaluations should aim to narrowly target only chem-bio AI models most likely capable of conferring characteristics of concern; attempting to evaluate all models would be both impractical and unnecessary.
- Evaluations should involve both quantitative and qualitative metrics, developed with interdisciplinary expertise.[48]
- Evaluations should be conducted on an ongoing basis.
- Evaluations should take into consideration the *marginal risk* that a new tool or method confers (i.e., the extent to which capabilities could be achieved with similar effort through existing or traditional tools or methods).[49]
- Evaluations should be iteratively assessed for construct validity, as technological advancements may render existing evaluation frameworks inadequate or obsolete.

Approaches to biosecurity mitigations should be based on empirical evidence, account for the practical research requirements of scientists, and consider how mitigations might inhibit technical advancements that could enhance biosecurity. Different approaches to risk mitigation could have drastically different consequences for the research community. For example, approaches that limit access to code or model weights prevent researchers from being able to reproduce and verify research findings, a prerequisite for advancing scientific knowledge.[50] Furthermore, much of academic research relies on the ability to access, customize, and openly share code; there are a number of scientific advancements by third-party researchers that rely on the ability to access the internals of previously published chem-bio AI models in ways that would be impossible with a "black box" or server-based approach.[51] Additionally, chem-bio AI models may be developed by small academic research groups that may not have the facilities to implement ongoing mandatory access control mechanisms, which may result in the model becoming unavailable to anyone. It is important that biosecurity mitigations are developed carefully, with the following considerations in mind:

- Biosecurity mitigations should be based on empirical evidence—reaffirming the need to establish rigorous evaluations.

---

[48] The Small Molecule Steering Committee launched by Polaris is an example of a concerted effort to develop "standardized, domain-appropriate datasets, guidelines and tools for the evaluation and comparison of methods": Cas Wognum, *Introducing our Small Molecule Steering Committee*, Polaris (2024), https://polarishub.io/blog/introducing-our-small-molecule-steering-committee; WelQrate is another example for developing standards in small moledule drug discovery: Yunchao Liu et al., *Welqrate: Defining the gold standard in Small Molecule Drug Discovery Benchmarking*, arXiv (2024), https://doi.org/10.48550/arXiv.2411.09820.

[49] Sayash Kapoor et al., *On the societal impact of Open Foundation models*, arXiv (2024), https://doi.org/10.48550/arXiv.2403.07918.

[50] Stephanie Wankowicz et al., *AlphaFold3 transparency and reproducibility,* Zenodo (2024), https://doi.org/10.5281/zenodo.11391920.

[51] Casper A. Goverde et al., *De novo protein design by inversion of the AlphaFold structure prediction network*, 32 Protein Science (2023), https://doi.org/10.1002/pro.4653. Henry Dieckhaus et al., *Transfer learning to leverage larger datasets for improved prediction of protein stability changes*, 121 Proceedings of the National Academy of Sciences (2024), https://doi.org/10.1073/pnas.2314853121.

- Biosecurity mitigations should meaningfully take into account the difficulty of physically expressing and weaponizing agents of concern.
- Biosecurity mitigations should be designed with care not to inadvertently undermine research that could itself enhance biosecurity.[52]
- Mitigation approaches should be tested and validated by researchers before they are broadly recommended or applied.
- If access control or "Know Your Customer"-style approaches are warranted for academic models, external support for such access controls should be made available, and such controls should not limit further development by third-party researchers.

[52] For example, while limiting documentation for chem-bio AI models might seem like a potential security measure, it seems unlikely to be effective and would create disproportionate challenges for research. The legitimate users of chem-bio AI models, while knowledgeable and accomplished scientific researchers, are not necessarily experienced with using computational tools. As computational biology has flourished, researchers who have previously used exclusively non-computational methods are increasingly accelerating their research through the incorporation of computational techniques. This usage requires robust, easy to use tools and clear documentation to allow researchers to leverage these techniques without having to shift their research focus specifically to computation. See also: Lynda M. Stuart, Rick A. Bright & Eric Horvitz, *AI-enabled Protein design: A strategic asset for Global Health and Biosecurity*, National Academy of Medicine (2024), https://nam.edu/ai-enabled-protein-design-a-strategic-asset-for-global-health-and-biosecurity/.

Opportunities to enhance nucleic acid synthesis screening

The content of this section concerns Questions 4a-c of the RFI.

When considering the biosecurity implications of chem-bio AI models, it is crucial to bear in mind that computational designs can only cause harm if physically synthesized.[53] Therefore, while evaluating and mitigating the risks of model misuse is important, strengthening biosecurity measures at the stage of nucleic acid synthesis should be a priority. This is also an area where developers of chem-bio AI models could make significant technical contributions.

As computational biomolecular design capabilities—and in particular *de novo* design[54]—improve, existing nucleic acid synthesis screening tools will be less capable of identifying sequences with the potential to cause harm.

Developing screening tools that are more robust to *de novo* design capabilities may require technical approaches that are different from those employed in existing screening tools—which rely on comparing sequences against curated databases of known threats through direct similarity searches,[55] diagnostic signatures extracted from comparing threat and non-threat sequences,[56] or pre-generated functional variants filtered against known benign matches.[57] New approaches could include using protein multimer structure prediction tools to screen *de novo* designed protein sequences for binding affinity to known targets of select agents and toxins or employing deep learning or large language models to screen for changes in protein function that could be malicious.[58] A benefit of open-source chem-bio AI tool development, in this context, is enabling defensive screening capabilities to advance in parallel with advances in protein design capabilities.

There are many different conceptual approaches worth investigating, and, for each conceptual approach, many different ways in which they could be implemented technically. Technical implementations, in turn, are likely to vary in efficacy, interpretability, computational demands, time complexity, vulnerabilities, and dual-use considerations.

It is difficult to say from this vantage point which approach(es) might be most practical to implement in the real world. However, this challenge is particularly suited to scientists in our field.

---

[53] Community Values, Guiding Principles, and Commitments for the Responsible Development of AI for Protein Design (2024), https://responsiblebiodesign.ai/.

[54] Joseph L. Watson et al., *De novo design of protein structure and function with rfdiffusion*, 620 Nature 1089–1100 (2023), https://doi.org/10.1038/s41586-023-06415-8; Sarah Alamdari et al., *Protein generation with evolutionary diffusion: Sequence is all you need*, bioRxiv (2024), https://doi.org/10.1101/2023.09.11.556673.

[55] International Biosecurity and Biosafety Initiative for Science, *Frequently Asked Questions: Common Mechanism for DNA Synthesis Screening* (2024). https://ibbis.bio/wp-content/uploads/2024/02/IBBIS-Common-Mechanism-FAQ.pdf

[56] Jacob Beal et al., *Development and transition of fast-na screening technology*, Zenodo (2024), https://doi.org/10.5281/zenodo.10214870.

[57] SecureDNA, *Features,* https://securedna.org/features/.

[58] Vladimir Gligorijević et al., *Structure-based protein function prediction using graph convolutional networks*, 12 Nature Communications (2021), https://doi.org/10.1038/s41467-021-23303-9.

Unfortunately, there presently seems to be a lack of infrastructure, funding, opportunities, or other incentives to accelerate research in this domain as screening requirements for nucleic acid synthesis become more sophisticated.[59] Complementary to efforts to evaluate and mitigate risks of misuse of chem-bio AI models, the federal government should allocate resources towards establishing research infrastructure—funding grants, hosting competitions, and establishing partnerships—to accelerate research in next-generation nucleic acid synthesis screening tools that apply techniques beyond sequence comparison with listed agents.

Competitions, which have long been popular in this research community, may be a promising model for advancing these efforts.[60] Critical Assessment of Structure Prediction (CASP) experiments date back to 1994 and are still held every two years.[61] More recently, there have been several large-scale protein design competitions hosted by the non-profit Align to Innovate[62] and private company Adaptyv Bio,[63] among others.[64] In running competitions, precautions may need to be taken in handling certain information, such as test datasets, which themselves present a risk of misuse.

Finally, we acknowledge that securing nucleic acid synthesis involves challenges beyond technical solutions. The most urgent priorities may lie in achieving universal adoption of screening tools internationally and establishing systems to track synthesis orders[65]—even more so than developing next-generation screening techniques.

[59] National Science and Technology Council, Framework for Nucleic Acid Synthesis Screening, The White House (2023), https://www.whitehouse.gov/wp-content/uploads/2024/04/Nucleic-Acid_Synthesis_Screening_Framework.pdf; UK screening guidance on Synthetic Nucleic Acids, Department for Science Innovation and Technology (2024), https://www.gov.uk/government/publications/uk-screening-guidance-on-synthetic-nucleic-acids.
[60] Similar competitive funding models are being explored internationally. For example, in November 2024, the UK's AI Safety Institute (AISI) launched a bounty program for developing novel AI system evaluation methods and agent scaffolding tools, demonstrating growing international recognition of the need to assess and govern advanced AI capabilities through competitive research initiatives. See: UK AI Safety Institute, *Bounty programme for novel evaluations and agent scaffolding,* Gov.UK (2024), https://www.aisi.gov.uk/work/evals-bounty.
[61] Protein Structure Prediction Center, https://predictioncenter.org/index.cgi.
[62] The Tournament, Align to Innovate, https://alignbio.org/the-tournament-program.
[63] Protein Design Competition, Adaptyv Bio, https://design.adaptyvbio.com/.
[64] Jonathan Hsu, *Driving Innovation in Drug Discovery: The Role of ML Competitions,* Polaris (2024), https://polarishub.io/blog/driving-innovation-in-drug-discovery-the-role-of-ml-competitions.
[65] David Baker & George Church, *Protein design meets biosecurity*, 383 Science 349–349 (2024), https://doi.org/10.1126/science.ado1671.

Future safety and security of chem-bio AI models

The content of this section concerns Questions 5a-f of the RFI.

It is critical that policies concerning chem-bio AI models are developed with meaningful involvement of the scientific community, as even non-binding government guidance can have far-reaching effects on research through the incentives and precedents it creates. For example, guidance stemming from this RFI could shape federal funding priorities, influence organizations' oversight practices or information-sharing norms, or even inform future regulations that affect this field of research.

Working in collaboration with scientists who understand both the technical capabilities and limitations of these tools will help ensure that government actors are better equipped to identify and address legitimate security concerns while mitigating the risk of policies that place unduly burdensome barriers on beneficial research. A majority of the researchers who responded to our survey were at least moderately familiar with the concept of biosecurity (Figure 2) and felt confident that they could evaluate the biosecurity risks of their own work (Figure 4); together, these findings suggest that the research community could substantively inform policy research.

We recognize, however, that several structural realities stand in the way of meaningful engagement with scientific communities. Academic and industry researchers often lack dedicated time, resources, incentives, or an interest in engaging in policy work. While nonprofit organizations working at the intersection of science and security policy can help bridge this gap, their perspectives may not adequately represent those of active researchers in the field.

Another fundamental challenge in developing sound science policy in this domain, in particular, is the current lack of empirical evidence regarding the actual risks posed by *in silico* research. While there is broad agreement that policies should be evidence-based, we currently lack robust mechanisms to gather and evaluate data that would help us understand where risks may lie.

Given these challenges, we recommend that the National Institute of Standards and Technology consider approaches that seek to address some of the structural challenges in developing biosecurity evaluations and mitigations (and other policies concerning chem-bio AI models).

A particularly promising approach would be establishing frameworks, advisory bodies, or other resources to support researchers in evaluating potential dual-use implications of their research methods in the early stages of their research planning, i.e., before a chem-bio AI model has been developed. DARPA's adoption of Ethical, Legal, and Societal Implications assessments in early program development exemplifies this approach.[66] These early evaluation frameworks could benefit from incorporating preregistration principles—where researchers commit to specific evaluation plans before conducting their analyses—which has been shown to enhance research reproducibility and credibility by clearly

[66] "Voices from DARPA" Podcast, Episode 79: Integrating ELSI, DARPATV (2024), https://www.youtube.com/watch?v=PYSbEnamSDA.

distinguishing between predicted and post-hoc findings.[67] NIST or other federal agencies could develop standardized evaluation toolkits and provide expert consultation services to help researchers assess dual-use concerns before initiating AI development projects. To drive the adoption of such resources, funding bodies could prioritize grants for entities that incorporate these evaluations, as demonstrated by the Bio Funders Compact's use of institutional commitments to promote biosecurity reviews.[68] These mechanisms should include structured feedback channels—such as standardized assessment forms and regular multi-institution workshops—to build a shared knowledge base of common challenges and best practices.

The development and oversight of chemical and biological databases could also enhance biosecurity while advancing legitimate research. Creating next-generation AI models will require extensive, high-quality datasets—a challenge exceeding the means of individual research groups. By supporting community-led data standardization efforts, governments can both distribute the costs and implement biosecurity practices. For example, the Protein Data Bank,[69] an open-access database of experimental protein structures collected over decades, has demonstrated how centralized, well-governed databases can accelerate scientific progress while maintaining oversight, and has formed the basis for modern AI protein folding prediction algorithms. This data-sharing model enables sensitive data to be screened, the tracking of data usage patterns, and provides foresight into potential model capabilities based on training data characteristics. Moreover, standardized data protocols could improve our ability to assess potential hazards by making model training data more transparent and traceable.

We also recognize the challenge of connecting technical expertise to where it is needed in government teams developing relevant guidance, such as those at NIST. Surmounting this hurdle could involve establishing policy fellowships for scientists that are well-suited for typical career development trajectories (perhaps modeled on the scheme organized through UK Research and Innovation[70]), creating technical advisory committee roles that are appropriately compensated, commissioning studies by the National Academies of Sciences, Engineering, and Medicine (NASEM), or mixed methods research designs that couples low-commitment, high-volume surveys with in-depth interviews.

Finally, we acknowledge that the protein design community has members worldwide. A core principle of our consortium is to build a welcoming environment where researchers from all backgrounds and nationalities can contribute meaningfully to scientific advancement.[71] Thus, we believe that international engagement will be essential for creating guidelines that are supported by the entire research community and are compatible across legal jurisdictions.

---

[67] Brian A. Nosek et al., *The Preregistration Revolution*, 115 Proceedings of the National Academy of Sciences 2600–2606 (2018), https://doi.org/10.1073/pnas.1708274114.
[68] International Bio Funders Compact NTI|bio (2024), https://www.nti.org/about/programs-projects/project/bio-funders-compact/.
[69] Helen Berman, Philip Bourne & John Westbrook, *The Protein Data Bank: A case study in management of Community Data*, 1 Current Proteomics 49–57 (2004), http://dx.doi.org/10.2174/1570164043488252.
[70] UKRI policy fellowships 2023, UK Research and Innovation (2023), https://www.ukri.org/opportunity/ukri-policy-fellowships-2023/.
[71] Diversity, Equity and Inclusion in the Rosetta Commons, Rosetta Commons (2018), https://rosettacommons.org/about/diversity/.

Appendix A: Rosetta Commons member laboratories (as of November 2023)

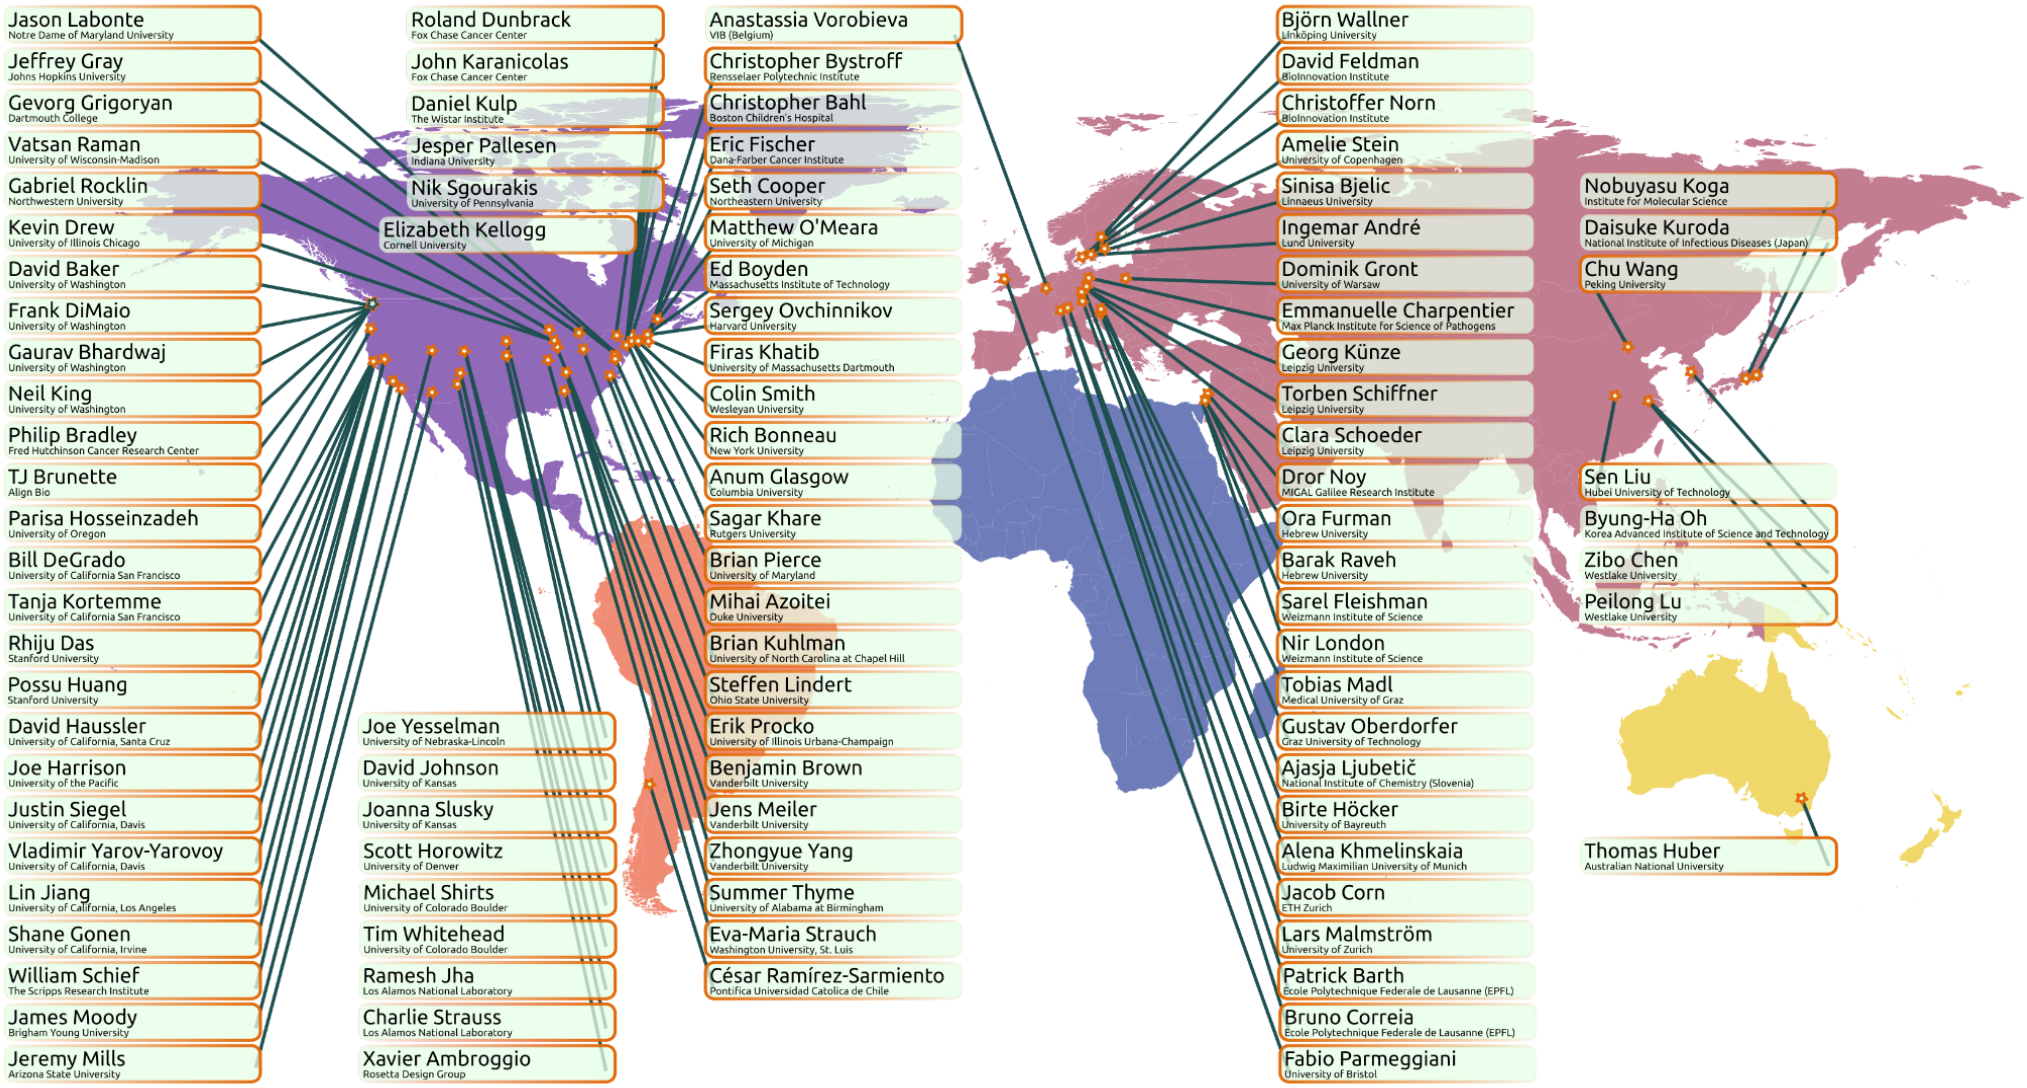

Appendix B: Survey

This section contains select findings from a survey developed in consideration of this Request for Information, for attendees of European RosettaCon, which took place in Denmark in mid-November 2024. Attendees were invited to access the survey via a QR code displayed on an auditorium screen. In the interest of collecting more input from Rosetta Commons community members, the survey was then shared on a Slack space open to all Rosetta Commons members. A majority of responses (~75%) were received during European RosettaCon—hence the majority of respondents being from Europe.

The survey was preceded by the following text:

> *Responses to this ~2-minute survey may supplement a comment on the U.S. National Institute of Standards and Technology's Request for Information (RFI) on Safety Considerations for Chemical and/or Biological AI Models.*
>
> *A term frequently used in the RFI is "biosecurity." The UN's Food and Agriculture Organization (FAO) and World Health Organization (WHO) define "biosecurity" as "a strategic and integrated concept that encompasses the policy and regulatory frameworks that analyse and manage risk in food safety, public health, animal life, and health, and plant life and health, including associated environmental risk."*[72]
>
> *In the context of computational biology, "biosecurity" typically refers to the dual-use implications of research artifacts—how emerging technologies can improve public health systems, and how, simultaneously, they modulate the risk of biological data and tools causing harm (deliberate or accidental). Discourse sometimes looks at how emerging technologies could be used to, for example, create toxic proteins or enhanced potential pandemic pathogens (ePPPs).*
>
> *If you do not know the answer to a question or are confused by or disagree with its formulation, please leave it blank (and leave a comment at the end, if you'd like).*
>
> *Your responses are anonymous unless you opt in to reviewing a draft of the comment in the final question.*

After the first question, which asked respondents about their familiarity with the term "biosecurity," the remaining questions substituted the terms "safety and security" rather than biosecurity, given that "biosecurity" might have been unfamiliar. Respondents were informed of this substitution of terms.

The survey also provided respondents the opportunity to share open-ended comments. These responses are not shared verbatim here, but we strived to reflect their sentiments in this document.

---

[72] Véronique Renault, Marie-France Humblet & Claude Saegerman, *Biosecurity concept: Origins, evolution and perspectives*, 12 ANIMALS 63 (2021), https://doi.org/10.3390/ani12010063.

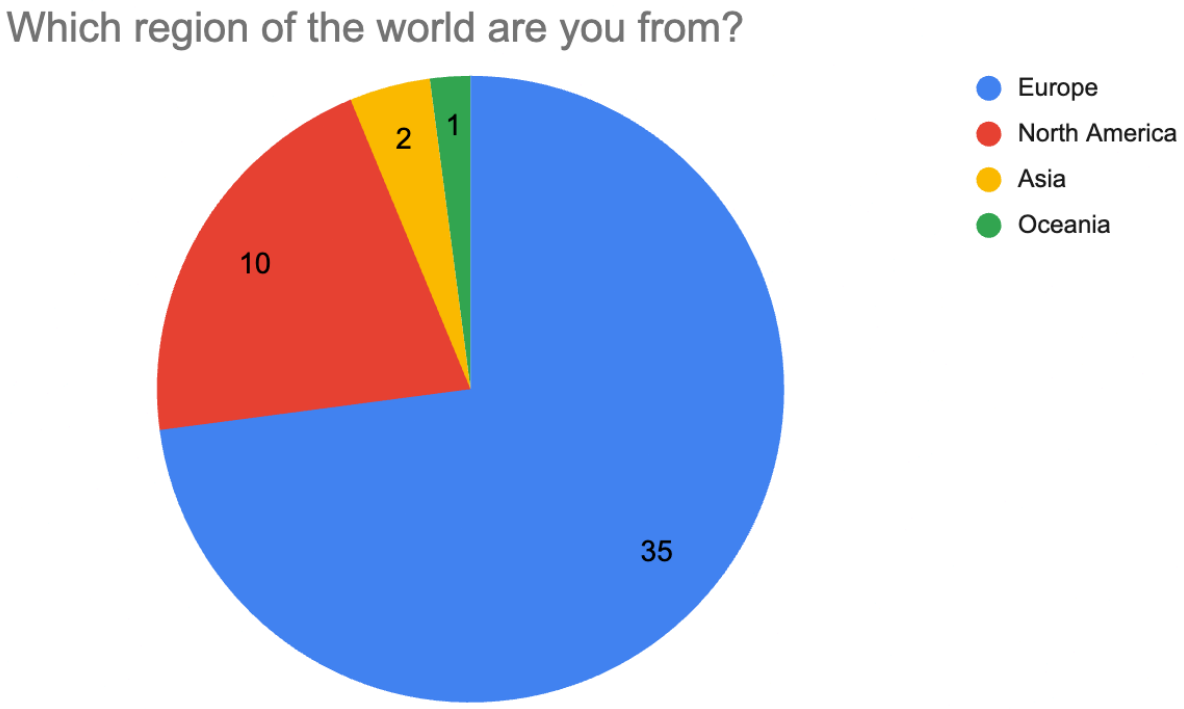

Figure 1. Respondents' geographic representation. (n=48)

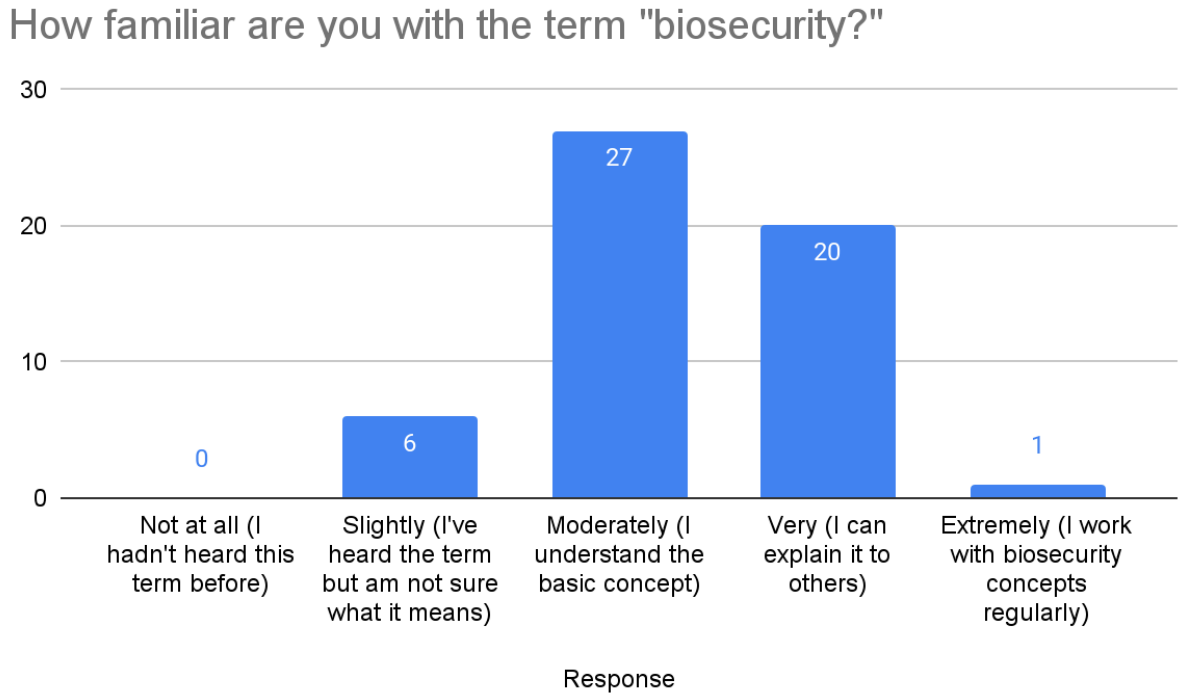

Figure 2. Familiarity with biosecurity. (n=54)

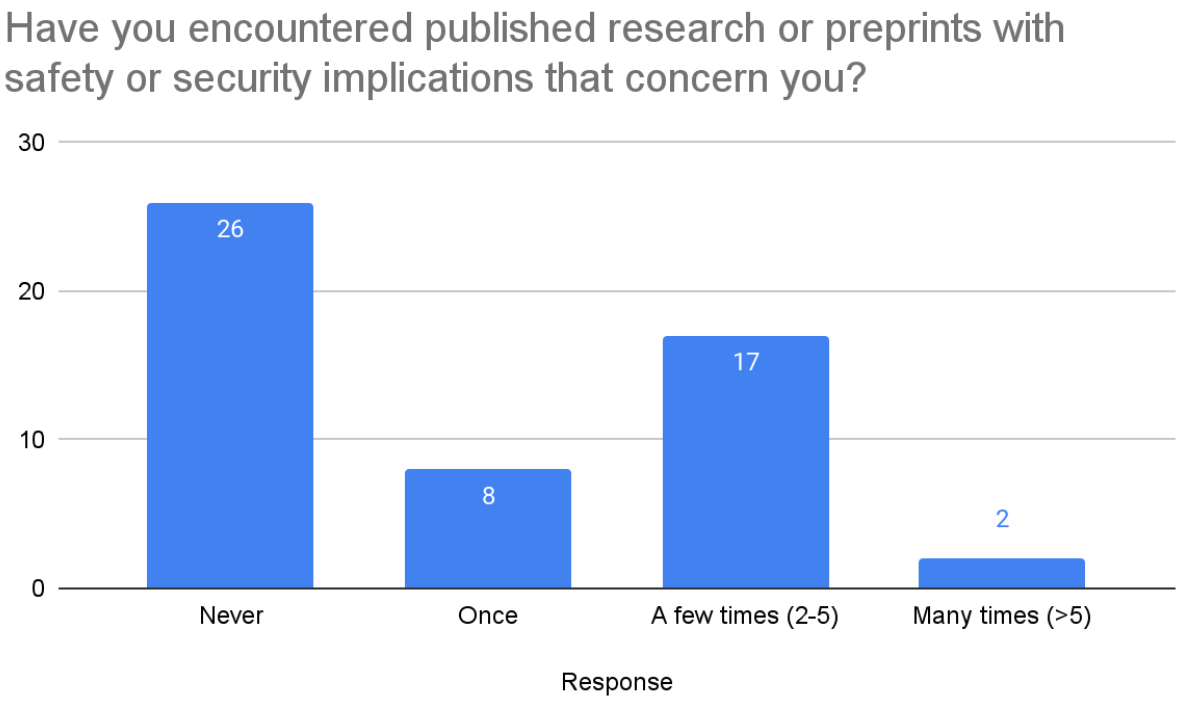

Figure 3. Concern with safety or security implications of published research or preprints. (n=53)

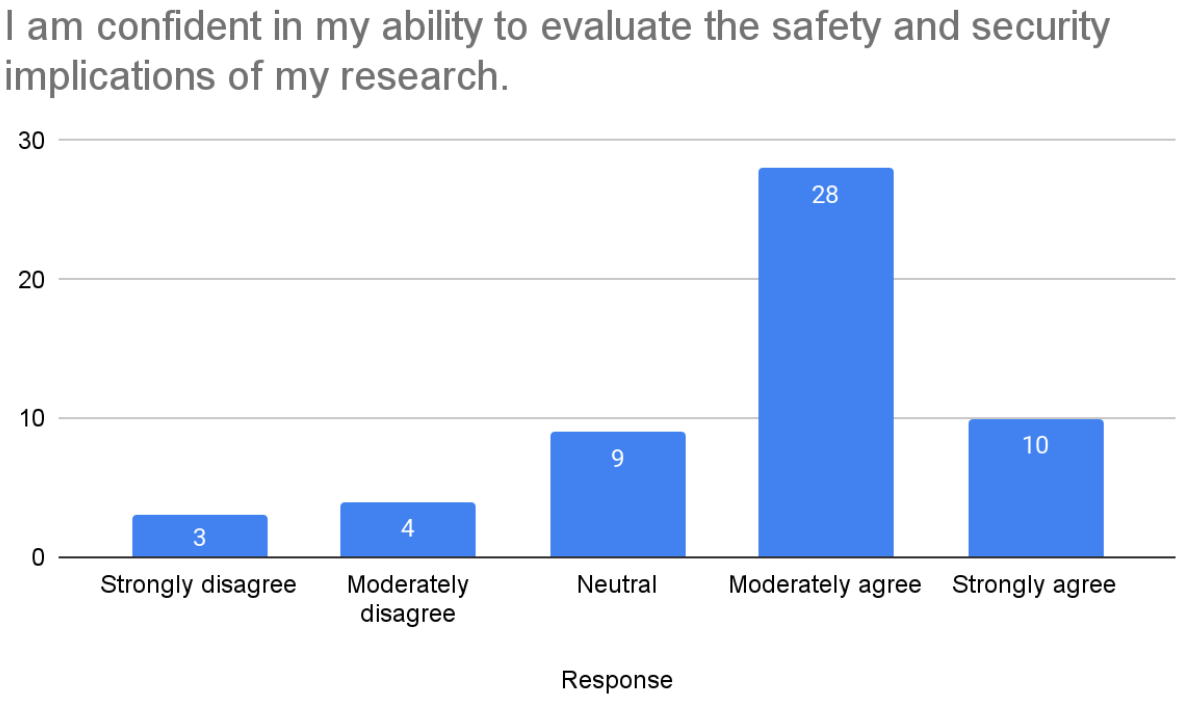

Figure 4. Confidence in evaluating the safety and security implications of one's own research. (n=54)

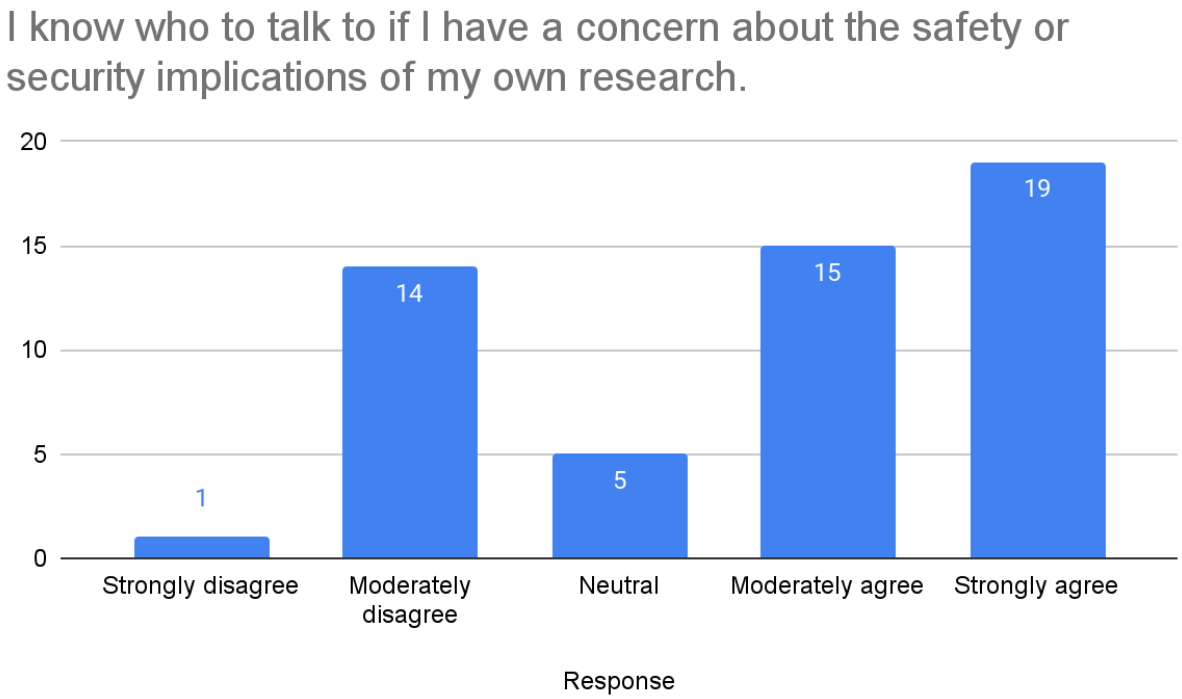

Figure 5. Knowledge of whom to talk to with concerns about the safety or security implications of one's own research. (n=54)

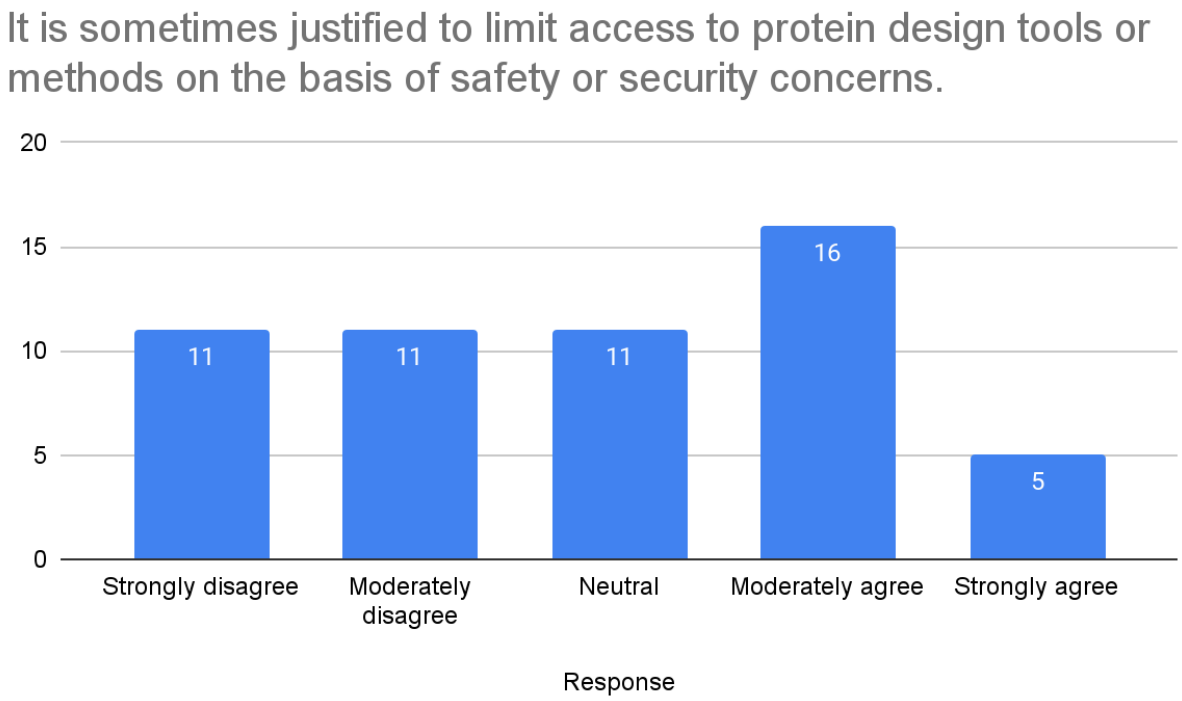

Figure 6. Whether it is sometimes justified to limit access to protein design tools or methods on the basis of safety or security concerns. (n=54)

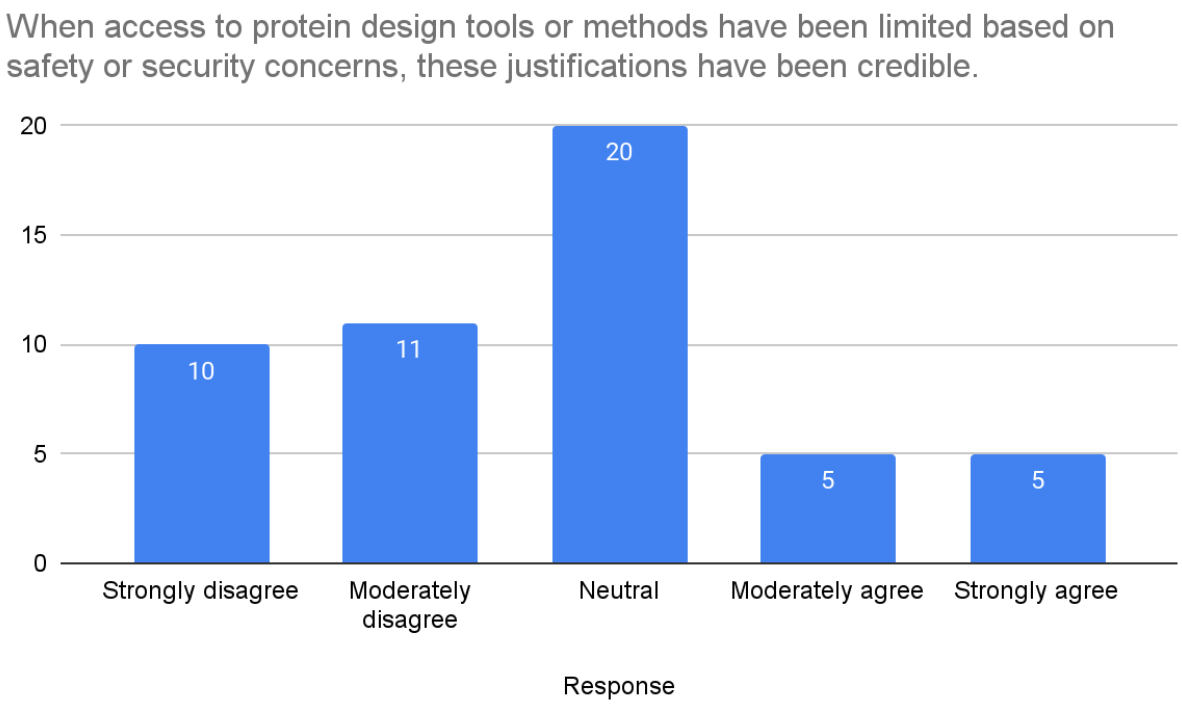

Figure 7. Whether justifications for when access to protein design tools or methods has been limited on the basis of safety or security concerns have been credible. (n=51)