

## **Potential Changes to the Policies for Oversight of Dual Use Research of Concern (DURC) and the Potential Pandemic Pathogen Care and Oversight (P3CO) Policy Framework**

An EBRC Response to OSTP RFI 88 FR 60513 Docket ID: EOP-2023-0001  
October 16, 2023

The Engineering Biology Research Consortium (EBRC) appreciates this opportunity to provide comments on potential changes to DURC and P3CO frameworks. EBRC is a non-profit, public-private partnership dedicated to bringing together an inclusive community committed to advancing engineering biology to address national and global needs. We recognize the immense importance of these policies for supporting a robust U.S. life sciences research enterprise that successfully innovates to improve public health, while ensuring that the research does not—through error or intentional misuse—create or exacerbate the problems it seeks to address. We do this through work in four Focus Areas, one of which is Security. Members of EBRC’s Security Working Group, in addition to other interested members of EBRC, contributed to the development of this response. Updates to these policies should focus on aligning review and mitigation practices with risk, and should provide clear direction to the research community.

**Q1a.** As has been noted by the NSABB, these policies have substantive overlap, including prioritizing safety and security in pathogen research while avoiding unnecessary review. Unifying them into a single policy should improve clarity as to how USG addresses the public health benefits and potential risks of pathogen research for all listed entities in addition to the broader public. Importantly, the Revised Policy should not aim to serve as a catch all for all biosafety and biosecurity governance.

Potential **benefits** of streamlining oversight across the entities listed and harmonizing the DURC and P3CO policies into a Revised Policy would include:

- Reducing the number of policies institutions and investigators need to be aware of and consult when developing research directions and assessing relevant research;
- Eliminating redundancy and closing gaps in oversight that currently exist under three separate policies.

There may also be some **challenges** in applying a Revised Policy to these entities:

- Great care should be taken to avoid creating ambiguity or importing parameters from one policy to the unified policy that create undue restrictions or leave significant gaps.
  - Investigators, institutions, and USG departments and agencies tasked with new or revised responsibilities will need proper education, training, and support.
  - Investigators who find the Revised Policy overly restrictive may even leave their publicly funded institutions in favor of institutions beyond the scope of entities listed above, driving pathogen research away from federal oversight.
- P3CO applies to funding decisions whereas DURC policies apply to research. This raises several questions: Will potential DURC research be included in the portion of the policy guiding funding decisions? Will the scope of DURC be altered to include all PPPs? Will there be a unified definition for research that needs review that encompasses all current DURC and ePPP research? Should all research subject to funding review also be subject

to review under current DURC policy? The answers to these questions will impact the entities described above individually and as they interact with one another.

- If the Revised Policy has a broader scope than current DURC or P3CO policies, it will require clear delineation of responsibilities listed for investigators, institutions, funding agencies and federal departments (see Appendix A of the 2023 [NSABB report](#)). There may be challenges or friction in retrofitting a new biosecurity policy onto existing biosafety institutional infrastructure.
  - New or changing responsibilities for funding agencies beyond HHS (e.g., NSF and USDA) should be clearly articulated and codified, perhaps through MOUs.

The US should embrace the opportunity to demonstrate international leadership as it revises DURC & P3CO policies, focusing on providing safeguards around pathogen-based research with biorisk while tirelessly pursuing research that promotes public health.

Given the diversity of entities described and potential magnitude of changes, the Policy should also build-in opportunities to assess its implementation and update and/or reevaluate as life sciences research and best safety and security practices continue to develop across entities.

### **Q1b.**

#### Challenges:

- Placing primary responsibility for identifying both DURC and ePPP research onto investigators and institutions will decentralize responsibility and decision-making.
  - Different investigators and institutions are almost certain to come to different conclusions when reviewing whether given research falls under DURC and/or P3CO policies, especially if/as the scope expands beyond list-based approaches. Revised policy(s) should make sure to establish a clear framework(s) for decision-making to minimize opportunities for differential interpretation.
  - An institution may interpret policy differently than USG officials. To identify when this might be the case, federal funders and fundees should communicate about biorisks and their management. Funders could still be charged with approving, or minimally having copies of, risk mitigation plans for USG-funded projects. When such research is privately funded at institutions that receive USG funds (as per Institutional DURC policy 6.1.b.ii), different interpretations of policy may not come to light until after research has been conducted and even published, perhaps receiving public scrutiny.
  - There may be confusion as to which individual(s) or committee(s) within an institution have final responsibility for identifying DURC and ePPP research, and what penalties may exist at the individual and institutional level for misclassification. Defining this, or laying out parameters by which an institution can determine these responsibilities, will be important.
- Ambiguity and fear of liability from investigators and/or institutions may result in the avoidance or unnecessary restriction of important pathogen research, leaving the US underprepared for future public health challenges.
- A redistribution of responsibility to institutions will require institutions to draw on resources that, for many, are already very limited. As those responsible for biosafety and biosecurity within an institution take on additional responsibility, institutions may i) shut down areas of inquiry they do not have the capacity to review in accordance with new policy and responsibilities or ii) set inappropriately high thresholds for research review, missing the chance to implement important safeguards and best practices.

### Benefits:

- Institutional review may be more prompt than federal review.
- If institutions and investigators bear more responsibility for the identification of DURC and ePPP research, it could improve all investigators' thoughtfulness for safety and security when designing research projects, not just those that work with the agents and toxins covered by current DURC policy. Given proper tools and resources (see 1c), more investigators could be sufficiently educated and trained on DURC issues such that they recognize risks, weigh risks and benefits, and propose and implement risk mitigation measures, even for research beyond the Policy scope.
  - Increasing awareness for a safe and secure research environment could increase public trust in investigators and institutions if communicated effectively.

### **Q1c.**

Institutional capacity: If the scope of oversight is expanded, institutions and IBCs may be flooded with new research cases for review without the capacity, resources, funding, or training to handle them, delaying and/or compromising the quality of review. Institutions should be given sufficient time and resources to prepare for policy changes. Revised policy or accompanying guidance should include adaptable examples of appropriate institutional structures.

Education & training: Currently, some safety officers and principal investigators at some institutions have little guidance or knowledge about DURC research. For example, smaller institutions lacking research on the 15 agents or toxins in current DURC policy may become newly subject of DURC policy. More generally, education will be needed for all trainees, investigators, and safety officers whose research is relevant to revised policy.

- Such education requires funding, dedicated staff, in person and online education, and incentives for scientists to receive the education. Resources could include informational guides on how to identify research under the scope of the Revised Policy and/or could provide a broader view of safety and/or security issues and how to determine appropriate levels of concern. EBRC's [Malice Analysis Workshops](#) are an example of the latter, designed to train researchers to critically evaluate research for potential security concerns and identify steps to reduce the risks of research.
- The development and availability of these resources could draw trainees into biosafety and (bio)security fields.

Database of examples: A database of example edge cases with detailed explanations of why a particular case does or does not fall under the Revised Policy would also be a useful resource. HHS could create and maintain such a database that could be updated as new cases and best practices emerge within the research community. HHS could alternatively suggest, facilitate, or support the sharing of case studies—real or hypothetical—between institutions. Case study outcomes need not be explicitly endorsed by USG, but could be a useful harmonization and information-sharing tool, perhaps hosted and facilitated by a public-private partnership. A good starting example would be Appendix C from the 2016 [NSABB recommendation report](#), which details the rationale for why or why not experiments were categorized as gain of function research of concern. Input from appropriate research and community stakeholders should be sought to ensure that as many ambiguities are addressed as possible. It should be updated over time as edge cases arise and as new technologies yield new case examples.

USG consultation: Finally, there should be a mechanism for seeking and receiving USG determination in situations where an institution is uncertain as to whether or not research is covered under Revised Policy. USG employees filling this role should be empowered to make

regulatory classifications akin to the regulatory classification of goods for export licensing (<https://snapr.bis.doc.gov/snapr/>) available through the Department of Commerce. Additionally, a technical advisory committee for the Revised Policy, similar to those advising the Department of Commerce (<https://tac.bis.doc.gov/index.php>), should be established to provide on-going stakeholder feedback on Policy implementation and impact. These committees could meet quarterly in public and closed sessions to provide regular, highly technical guidance on rapidly changing technology and security landscapes.

**Q2a.** Adoption of Recommendation 10.1 would place a significantly higher burden on institutions. The current policy for institutional oversight of DURC requires a PI to identify any research in their lab “that involves one or more of the 15 agents or toxins” listed in the policy. That identification is followed by an institutional review process to assess whether the research could produce one or more of the experimental effects of concern. This change would result in a **tremendous increase in the amount of research identified by PIs that requires institutional review for the 7 experimental effects**. There could be significant initial confusion as to which strains of which pathogens (e.g. E. coli) are subject to the policy. The strain placed on institutional staff would be substantial, which could result in other important work receiving less attention. Such increased responsibility may also result in low risk research involving pathogens being needlessly slowed down. The final policy details of some of the following points would impact the institutional burden:

- **Trigger for institutional review:** One possibility to ease the institutional burden is to ask PIs to identify research in their labs that uses any human, animal, or plant pathogen, toxin, or agent *and* is reasonably anticipated to result in one or more of the 7 experiment effects. Then, only research meeting both criteria would be elevated to institutional review for determination of whether the research meets the definition of DURC. This would create a slightly higher burden on PIs, but that burden would be distributed as opposed to being concentrated on institutional safety staff.
  - However, while there is little subjectivity as to whether or not a lab is working with one of the 15 agents and toxins in current DURC policy, there is subjectivity as to whether research can be “reasonably anticipated” to result in one of the 7 experimental effects of concern; tasking PIs with predicting such outcomes could have heterogeneous results. PI determinations could be influenced by:
    - PI perception of their own personal liability.
      - One of the 7 experimental effects unanticipated in advance by a PI could be identified by the PI or lab members as experimental results are analyzed, by a Program Officer reviewing funding reports, by editors or reviewers during the publication process, or at other points. In such cases, would there be a review process to judge whether the PI should have reasonably anticipated the outcome? Would there be a consequence for failing to anticipate the effect?
      - This actual or perceived increase of liability for under-resourced PIs without sufficient institutional support could have a chilling impact on research and result in subpar outcomes. It may dissuade researchers from approaching any work that may come under the

umbrella of this policy out of concern for making a review or reporting error and being held liable.

- PIs have a self-interest in avoiding delay and review, and may be subject to (conscious or unconscious) motivated reasoning when considering whether or not research in their lab could be “reasonably anticipated” to produce one of the 7 experimental effects of concern.
- Attenuated strains: Investigators and institutions must have clear understanding of if or which types of attenuated strains do or do not fall within the policy(s) scope. (Attenuation is addressed in the current [DURC Companion Guide](#); this could be referenced in policy.)
- White list: Perhaps a “white list” of pathogenic organisms and/or strains deemed sufficiently un concerning could be developed to avoid unnecessary bureaucratic burden (e.g., plant pathogens that do not cause serious disease in agronomic crops). This could be done at the federal level, or institutions could be empowered to know and understand the research occurring on their premises and develop contextualized white lists on an institutional basis. Any pathogen—or any microorganism, really—could be engineered enough to raise concern, but the actual risk for many pathogens is quite low.
- “Reasonably anticipated”: The precision with which this is defined will impact institutional burden (See Q4). Examples of what can be reasonably anticipated would be useful in final policy and/or companion document(s). E.g., it can be “reasonably anticipated” that inserting DNA sequences with known “functions of concern” into other genomes may confer some risk that should require appropriate precautionary measures. Or, when adding cell penetrating motifs to toxic compounds, one may “reasonably anticipate” that this could cause a security risk if aerosolized and take appropriate precautions before testing the toxicity of those hybrid compounds.

**Q2b and c.** Recognition that research beyond the 15 agents and toxins and 7 experimental effects currently described in DURC policies can be of concern is important. However, defining the scope of agents, toxins, and experimental effects that need this oversight is an immense challenge, and it may be wise to leverage previous USG work and standards to do this. As a broader scope is considered, it is important to recognize that casting too broad a net would carry its own public health risks associated with a slower-moving research enterprise.

Members of EBRC expressed different opinions on a reasonable scope for DURC policy, including that the scope should:

- not change; the most concerning research is identified and addressed under current policy;
- include all Select Agents & Toxins (not just the subset in question 2.b.i. of this RFI);
  - Synthesis of viruses with high homology or percent sequence identity to SATs could also be included.
- include all human, plant, and animal viruses with high transmissibility, thereby focusing on pathogens that could cause the greatest community or global harm. High transmissibility could potentially be defined via R0 or other quantitative measures;
- include any research producing an agent or toxin akin to those on the Select Agents and Toxins list, regardless of the actual species worked on or experiments performed. The capabilities to apply this standard may be developed in the future, but present capabilities cannot support it (e.g., what are the defining characteristics of a SAT and how can those be quantified such that they can be applied to other agents and toxins?).

Amongst these ideas, the overarching needs to tie policy scope to actual risk, to conscientiously consider institutional review capacity, and to not needlessly slow research were consistently reiterated. Policy revision approached with an experimental mindset, including plans and timelines for revision and assessment processes, should enable responsiveness to needed changes, evolving technologies, changing biorisk landscapes, and changes to societal and scientific risk/benefit calculations.

**Q2d.** A broad expansion of DURC agents and toxins could be limited by changes to the 7 experimental effects:

1. Enhance the harmful consequences of the agent or toxin:  
This is a broad classification and could be eliminated. “Harmful consequences” are mostly captured by experiments 2-6. Alternatively, this could be made more specific as to which harmful consequences, in addition to those in 2-6, are of particular concern.
6. Enhance the susceptibility of a host population to the agent or toxin:  
“Host population” could be narrowed. Under an expanded scope of DURC that includes all plant pathogens, for example, a plant host population may become susceptible to a bacterial pathogen if it that pathogen is given just one virulence (Avr) gene for which the plant does not have a cognate Resistance (R) protein. Alternatively, the *removal* of an Avr gene for which a plant *does* have a cognate R gene could also enhance susceptibility of the host plant. However, knocking out and knocking in these types of virulence and resistance genes is crucial to our understanding of plant disease and the future of agriculture and does not necessarily cause significant risk. Often, such research is done in laboratories far away from susceptible plant populations. However, these would still be captured by the expanded definition of DURC and this experimental effect:
7. Generate or reconstitute an eradicated or extinct agent or toxin or will synthetic biology techniques be used to construct a pathogen, toxin, or potentially harmful product  
Applying this experimental effect beyond the 15 agents and toxins currently in DURC policy would represent a significant expansion not necessarily reflective of risk. “Potentially harmful product” is a low bar and reflects process-based thinking (made with synthetic biology) as opposed to product-based thinking; non-synthetic biology research where a “potentially harmful product” may result is not otherwise captured by DURC policy. Additionally, the use of synthetic biology to construct a pathogen does not inherently pose more risk than obtaining the same pathogen from nature. The experimental effect could be narrowed to just extinct or eradicated pathogens.

**Q2e.** USG should provide as many resources and tools that help PIs and institutions understand their obligations and responsibilities as possible. Such resources should include as many examples as possible demonstrating what falls within the scope of the revised policy, particularly if the current list of 15 agents and toxins were expanded. USG could consult with academia and industry in a structured way to develop such examples to best capture nuance and areas of gray.

Because rare and ambiguous instances will inevitably arise after the policy’s development, having an office—ideally within a new government agency dedicated to biorisk management—that is able to support its implementation, would be of tremendous benefit. Such an agency would have mandate and funding to help PIs and institutions implement this new policy, provide education and training, answer questions, etc.

**Q3a.** The change from “or” to “and/or” between parts 1 and 2 of the definition of a PPP would certainly expand the scope of the policy, but expansion may be appropriate. It recognizes that

highly transmissible pathogens can still be devastating, even if their morbidity and mortality rates are relatively low. The change from “highly transmissible” and “highly virulent” to “moderate or highly” would also represent an expansion, but again not an unreasonable one when taken within the context that the policy directs the oversight of research to *enhance* PPPs. There would, of course, be challenges to this expansion:

- Under the proposed language change, pathogens with moderate transmissibility but low virulence, such as the common cold, could become more difficult to study.
- Just as “highly” virulent or transmissible is a subjective term, “moderately” is subject to interpretation. Where possible, parameters that define these terms should be developed along with guidance for handling uncertainty when data is insufficient or of low quality.
- Departmental review called for in the P3CO policy requires the time of Federal employees. If more research proposals require this Department review, more employee time will need to be dedicated to it. That could result in slower progress in other HHS priority areas. Worse, the accumulation of proposals requiring review could result in long delays between the time a proposal is submitted and the time a decision is reached.
  - As a result, PIs may work hard to avoid research that may be subject to P3CO policy. This is positive when the same advances to public health can be made with less risk, but becomes a risk in and of itself when researchers avoid priority research areas because of slow federal review.
- The policy must be articulated with sufficient clarity such that researchers who consult the policy are able to design experiments and build research programs understanding what research is likely or unlikely to be subject to the policy. For example, are experiments using commercial viral replicon systems subject to the policy? What about custom replicon systems that have not been subject to peer review?

**Q3b.** While respiratory transmission might be the most concerning route for an uncontrolled global pandemic, vector-borne and blood-borne pandemic pathogens still raise considerable concerns. For example, ebola is not primarily spread via a respiratory route, but given its propensity for outbreaks causing significant morbidity and mortality, any research to enhance its transmissibility and/or virulence could raise significant biosecurity and biosafety concerns and thus benefit from inclusion within the scope of this policy.

**Q3d.** Tying “moderately virulent” and “moderately transmissible” to features or characteristics would be very useful for helping to triage research potentially subject to this policy. A characteristic of transmissibility worth exploring for use in developing a standard is R0. By developing high, moderate, and low transmission categories based on R0, it might be possible to quickly identify which category a pathogen falls into and thus whether or not it meets the “moderately or highly transmissible” definition.

However, R0s are estimates based on available data, which is often very limited and specific to given strains or genotypes. R0 may change season by season as a result of pathogen evolution, seasonal vaccine efficacy, and current social practices and norms (e.g., social distancing measures during the COVID-19 pandemic decreased the transmission of other viruses, too). R0 values could also change as a researcher modifies the pathogen genome, which would be impossible to measure in humans and expensive and time-consuming (and not without risk) to measure in other species. Thus, HHS should work to provide up-to-date categorical descriptions and designations with the best available information. It would also be useful to quantify the degree of uncertainty around an R0 estimate for a given virus, e.g. if a 95% CI around an R0

estimate contains the threshold for moderate or high transmissibility, is oversight is triggered (even if the R0 point estimate is below the threshold)?

Perhaps moderate and high virulence could be defined by morbidity and/or mortality rates. These face many of the same challenges as described above for transmissibility.

**Q4a.** “Reasonably anticipated” does not have clear bounds, but “non-trivial likelihood” and “highly unlikely” are similarly subjective. Ultimately, since efforts to reasonably anticipate outcomes require making predictions about the future, there is no way to ensure uniformity or total agreement. One change might recognize that individuals with expertise often disagree, thus “an individual with scientific expertise” could be expanded, e.g., to “a majority of individuals with scientific expertise.”

**Q5a.** It may be reasonable to remove blanket exclusions for research activities associated with surveillance and vaccine development, as recommended by the NSABB. If exemptions are in fact removed from Revised Policy, it is crucial that an accelerated framework simultaneously be developed to enable research to be undertaken rapidly should a biothreat develop or intensify. Instead of a blanket exemption covering all experiments within certain application areas of research (i.e. vaccine development and surveillance), perhaps it would be appropriate to approve a subset of specific experiments within those application areas which could be exempt. For example, exemptions for sampling and sequencing could continue, but not genetic manipulation of surveilled potential pathogens. Or vaccine research could continue to test vaccine efficacy against existent pathogens, but generation of high-growth strains may need to seek additional oversight. Determining which experiments are appropriate for such a policy should be determined in consultation with investigators, (bio)security experts, and public health officials and updated regularly, perhaps as an additional charge to NSABB activities.

**Q5b.**

Vaccine development:

Benefits to removing exemption:

- Research for vaccine development is typically undertaken because a pathogen has significant morbidity and/or mortality. Thus, making enhanced versions of those pathogens poses a public health and safety risk. A deliberative process to weigh risks and benefits could help protect international and domestic public health, as opposed to the current de facto assumption that the benefits outweigh the risks.

Benefits to retaining exemption:

- Prevent delays in vaccine development, which can result in lives saved.
- Anticipation of delays or regulatory uncertainty could move vaccine research, supply chains, and manufacturing offshore, harming national health security in times of need.

Surveillance:

Current P3CO policy is unclear as to how or which kinds of surveillance activities could enhance a pathogen’s pandemic potential.

- Surveillance for known pathogens, for example through sampling and sequencing wastewater, without further experimentation to enhance or alter pathogen virulence or transmissibility, is common. Such work poses little risk and could be considered outside the scope of P3CO policy if not explicitly exempted.
- When sampling and sequencing surveillance activities identify a novel (potential) pathogen of interest, or a known pathogen with a novel genotype, microbial



characterization may be warranted; it is nearly impossible to predict the overall virulence and transmissibility of a (potential) pathogen from its genome. Characterization can be pursued via many experimental methods and ultimately the degree of risk is determined by which experimental route is chosen.

- Surveillance activities may also include the enhancement of a known pathogen's virulence and/or transmissibility characteristics to model and study what might occur in nature. Doing so enables the preparation of appropriate countermeasures. If new genotypes of concern are identified in these experiments, then surveillance efforts can watch closely for corresponding mutations and initiate public health responses immediately if found. In the Revised Policy or accompanying guidance documents, further clarification of types of surveillance experiments is needed because “surveillance activities, including sampling and sequencing” is insufficient.

Benefits to removing exemption:

- Engineering a known pathogen to aid future surveillance efforts and/or engineering a novel pathogen identified by sampling and sequencing surveillance activities involves risk. Creating the very pathogen characteristics that we hope to be able to quickly respond to and mitigate brings pathogens into existence that did not previously exist (to our knowledge) and that might never occur naturally. Even if lab accidents and leaks are rare, they do happen. It therefore makes sense for USG to have additional oversight of these activities through the P3CO framework.
- Bioprospecting can be considered a type of sampling and sequencing surveillance. These activities can be hazardous because there is little known about the properties of what is being collected. By subjecting bioprospecting activities to P3CO, a thorough assessment of risks, benefits, and assurance of biosafety precautions and oversight can support public health and instill public confidence in the safety and security of surveillance research.

Benefits to retaining some surveillance exemptions:

- Without data, it is very difficult to draw accurate conclusions about a novel pathogen's threat characteristics. Some of the needed data can be collected through sampling, sequencing, and microbe characterization. If those relatively lower risk efforts remain exempt (or are deemed not to involve the development of an ePPP and thus not subject to P3CO), we can continue important surveillance work to inform public health without overly burdening researchers or posing unacceptable risk. Additionally, if they are not subject to P3CO policy, investigators and biosafety officers / reviewers can focus their reviewing efforts and development of mitigation plans for research that is more hazardous.
- There may be some benefits to retaining exemptions for higher risk surveillance activities / experiments that characterize pathogenic traits such as virulence or transmissibility through organism engineering (potentially enhancing those pathogenic traits). Certainly it would enable investigators to move their work forward more quickly, lessen institutional review load, and may make some researchers more likely to pursue these important, but challenging lines of research. While this may involve riskier experiments to understand what factors enhance (or diminish) a pathogen's pandemic potential, associated risks *could* be worth the benefits of more informed public health decisions.

**Q5c.** The sample language narrows the scope of the exemption by adding “directly” before “associated.” Of course, “directly” can still be subjective, so if such language is used, examples in addition to “generation of high growth strains” should be included to maximize clarity.

Exemptions should be as clear and well-defined as possible. Finally, it is important that NSABB continue to review existing or needed exemptions as life sciences rapidly evolve and regularly propose updated exemption language that maximizes public safety and security. Ultimately, researchers must feel confident in what does and does not fall under exemption regardless of political or threat situation in order to maximize readiness.

**Q6a.** Advances in computational capabilities, including modeling and *in silico* experimentation, do cause legitimate safety and security concerns. However, DURC and ePPP policies may not be the most appropriate place for their oversight. DURC and ePPP are pathogen-focused policies. *In silico* life sciences research can have safety and security concerns that extend beyond pathogen research (e.g., engineering proteins to degrade materiel, engineering metabolic pathways to synthesize narcotics). Thus, while USG should establish processes to continually recognize emerging *in silico* capabilities and assess resultant security and safety concerns, it may be best under the purview of a different policy or entity such as a nationally coordinated Biosafety and Biosecurity Innovation Initiative. Then, future review processes for DURC and ePPP oversight policies could pay special attention to any relevant emerging capabilities *enabled by in silico* research and consider if/how relevant aspects of those capabilities might be captured by policy.

Rather than mitigate all risk at the design / *in silico* stage, other appropriate safeguards such as those that control access to pathogen DNA, must be robust. The design of the deadliest enhanced pathogen still cannot cause disease unless it is made in the physical world, which is not a trivial undertaking. Importantly, some of the types of modeling and *in silico* research that conceivably could fall within modified DURC or ePPP policies can also enable those who control access to pathogen DNA (e.g., DNA synthesis companies, plasmid repositories) to better screen the orders they receive. Because computational tools can be developed, shared, and accessed by anyone, anywhere, with any source of funding, making it more challenging for US-funded researchers to use dual use computational capabilities could ultimately enhance biorisk.

**Q6b.** The parameters described in (i) could capture a very wide swath of computational work with vastly greater benefits than risks. For example, perhaps AlphaFold2 could be used to design a more heat-stable version of a botulinum neurotoxin. But subjecting all federally-funded research to develop protein engineering tools to ePPP and/or DURC policies would be a misallocation of time and attention, would be challenging to implement, and could cause US researchers to fall behind their international peers.

Of course, more targeted and specific models could also be developed, for example with the narrow purpose of designing enhanced ePPPs. While such models could deliberately be misused, they might ultimately enhance laboratory safety by reducing the amount of experimental trial and error in the physical world. Their safety could be enhanced by the development of probabilistic risk estimates assigned to each design. Then, the safety / oversight / review mechanisms for the physical experiments resulting from those designs could be informed by those risk estimates. Access tracking and controls would be important, although not necessarily appropriate for governance under DURC / P3CO policy(s).

The suggestion in (ii) would likely have sweeping impacts on common databases and datasets and is too far removed from the creation and release/escape of a pathogen to be appropriate under these policies. ePPP and DURC review should focus on the potential harm that could be caused as a result of the experiment itself. Other policies and governance frameworks may be more useful for considering how tools or outputs of research could be misused by others.