Policy Working Paper [Revision 0.9]

Beyond risk-benefit analysis: pricing externalities for gain-of-function research of concern

Owen Cotton-Barratt*†, Sebastian Farquhar†, Andrew Snyder-Beattie*

March 2016
*Future of Humanity Institute, University of Oxford
†Global Priorities Project, Centre for Effective Altruism

Executive summary

The recent US moratorium on certain types of Gain-of-Function‡ (GoF) research[1] made it clear that a new approach is needed to balance the costs and benefits of potentially risky research.[2] Current risk management tools work well in the context of most laboratory risk, where risks are local.[3] However, in the case of potential pandemic pathogens, even a very low probability of accident could be unacceptable given the consequences of a global pandemic. Although quantitative assessment is feasible for these low-probability, high-stakes risks, simultaneously comparing these risks with the qualitative benefits of such research is an especially difficult task.[4]

In this policy working paper we outline an approach for handling decisions about GoF research of concern. Our central policy objective is that:

> *Proposals for research projects with the possibility of catastrophic accident should have an independent estimate of the expected damage, and this figure should be explicitly included in the cost of the research project.*

Our policy objective has three key advantages:

1. It keeps decisions about which science is worth funding in the hands of scientists.
2. It incentivizes sponsors to fund research only when the scientific merit outweighs the costs because the negative externalities are considered as part of the cost of the research project - without the need for a direct benefit-cost analysis.
3. It provides a generalizable solution, which can be applied to other emerging risks from science and technology.

We propose and compare two different approaches to achieving the policy objective.

‡ In this paper we use the term 'Gain-of-Function' to refer only to the research covered by the recent White House moratorium.

Notes: We are grateful to Carrick Flynn for research assistance and to Anthony Aguirre, Leah Broad, Marc Lipsitch, Kathryn Mecrow, Piers Millett, and Stefan Schubert for their comments on drafts.

The first is to establish strict liability for any damages that result from GoF research of concern, and to require grant-holders to purchase liability insurance as part of the grant. The strength of this approach is market-based and incentivizes insurers to price externalities correctly.

The second approach is to centrally commission assessments of absolute risk and require a payment to a state or non-state body to cover the expected cost. The strength of this approach is that it works even if there may be no clear liability after the fact, so could address biosecurity as well as biosafety risks.

Framing the issue

Recent controversy has emerged around certain types of GoF research. Scientists remain deeply divided on both the benefits and the risks of such research.[5]

The controversy culminated in a moratorium on GoF research of concern pending an independent assessment of the risks and benefits. The NIH commissioned an assessment from Gryphon Scientific, released in December 2015. The report did not draw firm conclusions on whether the benefits of such research outweighed the risks.[6]

Challenges of risk-benefit analysis for scientific research
In principle, analysing the risks and the benefits of research and weighing these against each other is the correct way to determine whether to pursue risky research. In practice, both sides of this are very difficult to analyse. Gryphon Scientific was able to present a tentative quantitative analysis of absolute biosafety risk, but only a qualitative analysis of the benefits of the research and of the biosecurity risk.[7] That this major review was not able to analyse risks and benefits on a common scale demonstrates the difficulty of this type of analysis.

Existing solution: the scientific grant process
The scientific grant-making process is the primary mechanism for assessing the uncertain benefits of research against their costs, including the opportunity cost of unfunded research. Although it is hard to judge quantitatively, expert reviewers assess the potential for scientific excellence in different proposals. They must regularly make trade-offs between projects with disparate and uncertain benefits.

GoF experiments are the outcomes of successful grants. But these are currently assessed primarily on the basis of scientific merit and potential benefits, with comparatively little emphasis on the scope of possible risks. Risks from research, just like the benefits, impose an externality on the public. Because the risks are not considered as explicitly, a risky project could get funding over a safer one which has equal or only slightly lower expected benefit. This means that the public is implicitly subsidising risky research relative to safe research.

Our policy approach
Since the benefits are difficult to assess, any direct comparison of risks and benefits is extremely difficult, even when the risks are well-quantified. Rather than employing a direct comparison, we

suggest using an absolute risk assessment to price the expected risk, and to explicitly include this cost in grant proposals. This allows the scientists making grant allocations to use their judgement to pick projects with the greatest potential benefits, given their true social costs.

A generalizable solution
Biotechnology is not the only research area that could create the potential for small probability, high impact risks. Other fields might need to grapple with similar governance issues.[8] A solution that could be extended to other fields, with moderately straightforward generalizations, might avoid harmful controversy and delays while the issues are resolved. Our approach is likely to be generalizable in this sense.

Policy target: have risks priced into grants

In this section, we outline the intended results of pricing risk externalities into grants. We explain what this would look like, why we think it would be beneficial, and how it could perform better than existing safety approaches. In the next sections, we explore two potential mechanisms for achieving this.

Key policy target
Our central aim is:

> *Proposals for research projects with the possibility of catastrophic accident should have an independent estimate of the expected damage, and this figure should be explicitly included in the cost of the research project.*

If the cost is explicitly included, the project would internalise the negative externality associated with risks to the public. For now, we set aside the issues of where the independent estimate comes from, or where the money to cover this explicit cost goes. We will return to these questions in the next two sections.

Benefits of achieving this target
The principal benefit would be to keep decisions about experiments in the hands of scientists, who are best-placed to evaluate the potential benefits, while removing the implicit subsidy for risky research over safer research.

This would have a number of instrumental benefits. First, it should mean that experiments are funded precisely when the benefits outweigh the costs (including both the risks *and* the opportunity costs of not funding other experiments). Second, it would incentivize scientists and laboratories to look for alternative ways to run experiments that would reduce the risk, as this could reduce their extra costs.

While there are significant existing biosafety measures, in many nations these are driven by regulations focused primarily on occupational health (the safety of lab workers), rather than public health.[9] This focus accurately reflects the median historical risk, since most lab acquired infections

have not been passed on. However, it does not reflect the risks posed by experiments which could cause pandemics, where most of the risk exists in the small chance of catastrophic public damage. Our proposal would incentivize effective ways to minimise these risks.

Effects on grant process and domain of applicability

The independent risk assessment could take place before grant applications are submitted (at the request of the group intending to apply for a grant) or after the grant application is submitted (at the request of the grant body). In either case, some time and work would be needed for a proper risk assessment.

If it affects many areas of research, this requirement would significantly increase bureaucratic overhead. Accordingly, we recommend that if implemented, it initially apply only to the GoF research which is covered by the recent US moratorium. In the future, it could potentially be extended to other areas which pose significant risk to public health.

Effects on funding

Laboratories currently receive an implicit subsidy because they do not fully internalize the probabilistic costs of their dangerous activities. If the larger research community were asked to internalize this cost out of their existing limited budget allocation, it would represent an additional unfunded overhead expense and a functional shrinking of the budget available for actual research. Since research is potentially of great benefit to humanity, this may not be desirable. Instead we recommend that the government proportionally increase funding for life sciences research to compensate for this additional expense. Although this would increase explicit expenditure by the government in the form of larger research budgets, governments are already responsible for public health crisis management. This essentially transfers expenditure, from crisis management to crisis prevention, by making the implicit subsidy explicit. The primary advantage of this budget reallocation is that it allows for the same functional cost to the public, while removing the distortion of incentives created by the hidden externalities.

Reporting and safety culture

Any approach which penalises laboratories for reporting accidents and near misses in a timely way might harm biosafety and biosecurity in the long run. Reduced reporting makes it harder to use lessons from mistakes to improve lab design and impairs accident response. Mechanisms for pricing risk will work best if they avoid creating perverse incentives around reporting, and we believe that the mechanisms we describe below can be constructed in a way that does so.

First potential mechanism: mandatory liability insurance

Our first approach is market-based. Laboratories conducting experiments in the appropriate class could be mandated to purchase insurance against liability claims arising from accidents associated with their research. Ideally, this research should be explicitly classified as an "inherently dangerous activity" by the legislature. This will establish strict liability for any damages caused by accidents,

which means that laboratories would be liable even if there was no negligence. Strict liability is already legally established for other inherently dangerous activities analogous to this research, and might well be the legal standard used in many common law jurisdictions in a GoF case even without legislative intervention. The advantage of making this clearly established is that it would provide laboratories with strong incentives to minimise risk.

It is beneficial to require insurance, rather than just ensure there is liability, because of the "judgement proof problem."[10] Many universities currently self-insure against the damage of accidents in their research. This makes sense for occupational and small-scale public health issues, but for cases where there is a small chance of catastrophic damage, the institution may not have enough assets to cover the potential damage. Additionally, a blanket policy of self-insurance may mean that financial planners within universities do not even carefully consider liability risks of their specific research activities.

Advantages of the liability approach

There are a number of advantages to taking this market-based approach. First, it is a relatively light intervention, requiring less ongoing work from the state. Second, it incentivizes insurers to accurately estimate risks, reducing possible politicisation of the risk assessment process. Scientists and engineers would also be incentivised to devise effective safety protocols to reduce their institutions' insurance premiums. Imposing liability has been seen to improve outcomes in other domains such as occupational safety, medicine, and general risk management in non-profits and governmental agencies.[11]

Possible issues with the liability approach

A big question about mandatory insurance is whether insurers would in fact be willing to insure against these outcomes. There are two main reasons why they might not.

The first is that the potential risks are simply too large. A bad global pandemic could kill hundreds of millions of people, and even the largest reinsurers would be unable to absorb this cost without bankrupting themselves (costs above this level will be implicitly backed by the state or the public in any case). It is better to be explicit, and cap liability at a specific industry-wide figure. If the cap were sufficiently high, the effect would be improved risk aversion, even if the tail risk for the insurer were not fully internalised.

Secondly, the risks are hard to model and the market would be small. Developing models to estimate the risk could be more costly than the expected profit from participating in the market. Moreover, developing these models is a difficult task requiring rare expertise. However, insurers already have models for other hard-to-anticipate risks, such as terrorism and global pandemics. If necessary, the development of appropriate models to facilitate this insurance could be explicitly subsidised.

Aside from whether or not insurers are willing to take on the risk, there are challenges to international adoption. For example, it may prove difficult to harmonise liability laws across jurisdictions, particularly in international collaborations.

Aside from the question of whether any firms would be willing to insure against this risk, liability insurance can potentially increase moral hazard, by making actors less responsible for the consequences of their actions. If the deductible/excess on the insurance were set correctly, it could reduce the moral hazard while averting the judgement-proof problem. It also seems that in the particular case of GoF experiments undertaken to benefit public health, the nonfinancial consequences for any scientist who was involved in an accident with major harms to the public would be great enough to serve as a deterrent to reckless behaviour, even if financial consequences were mitigated by holding liability insurance.

Finally, depending on which legal framework is used in these types of cases, the market-based approach might not be able to capture various biosecurity risks. This is because it will likely be difficult in these instances to attribute a disaster to a specific project. This is particularly true with information biosecurity risks. It may, however, be possible to employ other methods of attribution, such as 'market share' liability. In contrast, the approach in the next section could potentially treat biosecurity risks in the same way as biosafety risks.

Second potential mechanism: centrally-commissioned risk assessments

The second approach is to centralise risk assessments. When an area of potential concern is identified, a body commissioned by the state would perform an analysis of the risks involved. This might be similar to the recent Gryphon Scientific analysis, except that it would not attempt to analyse the benefits, and it would focus only on producing the best-estimate absolute risk analysis for different kinds of work, rather than leaving it at a qualitative level. This absolute risk analysis would present its outcomes in monetary terms, using Value of Statistical Life figures to convert fatalities into a cost.

In order to do work of the relevant type, laboratories would be required to pay the corresponding cost to a central authority. This would most naturally be the body or bodies likely to absorb the cost in the event of catastrophe; such as the government's public health and disaster management agencies. It could also be used, in part or whole, to support the cost of the risk assessments.

Advantages of centrally-commissioned risk assessments
Compared to the market-based approach, centralising the risk assessments has two main benefits. First, it can be done without needing to persuade insurers to enter the market. Second, it does not suffer from the same legal uncertainties as the market-based approach. This is especially relevant in the biosecurity context, where attribution might be difficult. The Gryphon Scientific report concluded that the biosecurity risks looked at least as large as the biosafety risks, so this may be a significant benefit.

Questions and issues for centrally-commissioned risk assessments
Three unresolved questions are:

1) Who would make the risk assessments? This could potentially be an independent agency or

an outside contractor. It might be difficult to build the capacity to do the assessments well; note that the Gryphon Scientific report did not fully support the absolute risk analysis it presented, and did not offer any absolute risk analysis for biosecurity risks.

2) Who should receive the money that is included in the price of the grant? Should it be retained at the national level, or shared internationally (since the risks are global)?
3) How would this system work with international collaboration or non-domestic accidents?

A potentially larger issue is ensuring fair and accurate risk assessments. In the case of liability insurance, market forces help align the incentives to motivate insurers to make accurate risk assessments. If assessments are centrally commissioned, there is no such force keeping them in check, which means they would be at risk of becoming politicised.

Comparisons

We have outlined the advantage of our approach, which works by aligning the incentives for scientists and funding bodies more closely with those of society as a whole. This may be the only way to keep the assessment of the benefits of scientific research purely in the hands of scientists, while also reducing risks when appropriate. We have explored two different ways to achieve this. Each has its own advantages and disadvantages.

Overall, the liability approach is more market based. As a result, the risk-assessors have a financial incentive to accurately estimate risk, and political pressures are diminished. It might also be easier to use as a template internationally. Since the risks are global and the potentially risky research is not being pursued in just one country, being able to build global solutions is valuable.

The main benefit of the centrally-commissioned analysis approach is that it bypasses potential legal difficulties with attribution. It may therefore be a more general tool, able to correct incentives for a larger class of risks (such as biosecurity risks and information hazards).

For both approaches, it is important not to punish reporting of laboratory acquired infections or other accidents and near misses. Accurate information and a culture of open reporting are vital for laboratory safety and disease prevention.

Conclusions

Our approach is not to suggest a specific policy, but rather to outline different options which would facilitate a better evaluation of benefits and risks. We have suggested two quite different methods for achieving this. One relies more on market mechanisms, while the other depends on central oversight. Both would require the strong support of regulatory bodies. Each of them has a number of advantages and disadvantages. We do not feel we are in the right position to conclude decisively in favour of one over the other. We would like to encourage discussion among stakeholders of the relative merits of the two approaches.

References

[1]U.S. Department of Health and Human Services. (17 Oct 2014). U.S. Government Gain-of-Function Deliberative Process and Research Funding Pause on Selected Gain-of-Function Research Involving Influenza, MERS, and SARS Viruses

[3] Wilson, D. E., & Chosewood, L.C. (2009). Biosafety in microbiological and biomedical laboratories.5. *Centers for Disease Control and Prevention*, *12*.

[4] Casagrande, R. (2015). Risk and Benefit Analysis of Gain of Function Research, Draft Final Report. *Gryphon Scientific.*

[5] Duprex, W.P., Fouchier, R.A.M., Imperiale, M.J., Lipsitch, M., Relman, D.A. (2015) Gain-of-function experiments: time for a real debate. *Nature Perspectives 13*

[6] Casagrande, R. (2015). Risk and Benefit Analysis of Gain of Function Research, Draft Final Report. *Gryphon Scientific.*

[7]Ibid.

[8] Beckstead, N., Bostrom, N., Bowerman, N., Cotton-Barratt, O., MacAskill, W., O'Heighearteigh, S., Ord, T. (2014). Unprecedented Technological Risk. *Future of Humanity Institute*.

[9] Kimman, T.G.., Smit, E., Klein, M.R. (2008). Evidence-Based Biosafety: a review of the principles and effectiveness of microbiological containment methods. *Clinical Microbiology Reviews* 21(3): 403-425.

[10] Shavell, Steven. "The judgment proof problem." International review of law and economics 6.1 (1986): 45-58.

[11] Schwartz, Gary T. "Reality in the economic analysis of tort law: does tort law really deter?" *UClA L. Rev.* 42 (1994): 377.