Wilton Park

Image: Alexander Raths

Conference report

# Dual-use biology: how to balance open science with security

Sunday 15 – Wednesday 18 September 2013 | WP1260

**In association with:**

MINISTRY OF DEFENCE

Federal Foreign Office

National Coordinator for Security and Counterterrorism
*Ministry of Security and Justice*
The Netherlands

Sandia National Laboratories

Wilton Park

Conference report

# Dual-use biology: how to balance open science with security

## Sunday 15 – Wednesday 18 September 2013 | WP1260

This meeting brought together government and law enforcement officials, practising scientists, NGOs and academics, to conduct an appraisal of security concerns surrounding real world examples of past dual-use research of concern as well as hypothetical experiments. Options for managing concerns at all stages, from conception of an experimental question to publication of results, were assessed.

## Key points

" Dealing with dual-use research cannot be done by any one community, and no single set of actors can or should (or would want to) own this problem. "

- The security risk surrounding dual-use research appears to be generally recognised by those who are aware of the issue, but widening and targeting such awareness is still a challenge, as is reducing scepticism. This issue can be managed and ameliorated, but not resolved. Risk and hazard, and perhaps threat, is something that will have to be lived with not discounted; the price to be paid for the benefits of biological research.
- The growing multi-disciplinary character of biological research means that any regulatory management system that is developed to concentrate on work involving certain dangerous pathogen may fail to capture much future research of concern. This also has implications for outreach and education.
- Dealing with dual-use research cannot be done by any one community, and no single set of actors can or should (or would want to) own this problem. An assessment of what such a process of knitting together disparate communities would look like should thus be a priority.
- Collaborative risk assessment and the acknowledgement of inherent uncertainties and different costs benefit analysis at the start of research is potentially the most fruitful and effective stage at which to intervene, and should involve security input where appropriate.

## Balancing open science and responsible security

" At the heart of the issue is the inescapable fact that any area of biology has the potential for dual use."

1. If the right balance is to be struck between the free pursuit of scientific research and the demands of prudent security, policy makers and the public must have access to sound, fact-based analysis of potential security concerns and a range of appropriate tools and approaches to manage them. What are the security concerns? Where is the evidence that such concerns are valid? How does it relate to past dual-use research of concern? How are likelihood and impact best evaluated, and how are they to be weighed against possible benefits? In practical terms, what options exist for dealing with such concerns?
2. At the heart of the issue is the inescapable fact that any area of biology has the potential for dual use. Henry Dale, for example, discovered acetylcholine in the early nineteenth century, which opened the path to modern neurochemistry and

neurobiology, but also to modern nerve agents like sarin. Much-reported cases, such as the experiments involving the 1918 influenza or mousepox viruses, have tended to stimulate a renewed bout of debate about how to reconcile the drive to exploit benefits with the need to minimize the scope for misuse. If such a reconciliation is to be achieved, it will have to come through an approach that recognizes the multidisciplinary nature of modern life sciences, which now embrace mathematics, engineering and computing as well as more familiar biology and chemistry. Put simply, biology is no longer conducted solely by biologists and has become inherently multi-disciplinary. It therefore follows that the ethical, legislative and operational framework that will evaluate risks and benefits will also have to function across disciplines. The added complexity this generates is only likely to become more prevalent in future, as the convergence of biology with physical sciences and engineering (referred to by MIT as the third revolution in life sciences, after the molecular biology revolution that followed the discovery of DNA, and the genomics revolution of the 1990s) continues.

3. Moreover, debates about how to manage dual-use research have not yet fully integrated a range of new drivers that are shaping biology. Biology has steadily become easier, relatively speaking, to engineer; a rapid decline in government funding has produced a clear trend towards commercialization of scientific research; new organizational business models have emerged that are also being applied in fields of scientific research. These new drivers have produced different patterns of research, as well as different actors and different processes.

4. For a resilient and legitimate framework to operate within a community of researchers, those researchers will need to act in accordance with a clearly-defined set of principles and understandings. A concordat of principles of research, beginning at the national level but aiming to be international, could be an important step in developing this. It would bring together not only the researchers but also their employers, those who fund their research, and the journals that publish its results. Industry will also need to be engaged in order for any concordat to work; attracting smaller businesses may be a stronger challenge than larger multinationals. However, should it be possible to achieve buy-in from a critical mass of industry actors, it is more likely that others will follow.

5. If successfully developed, promulgated and adopted, such a concordat could establish collective sensibility all through the research process, and also its preliminary and post-research stages. In so doing, it would provide multiple points at which research can be scrutinised and evaluated. Many of the elements of such a concordat are already in place, and so in many respects the task is only to coordinate and make explicit what is already generally accepted.

6. Nonetheless, it will be important to bear in mind the dangers of looking for ‘magic bullet’ solutions where there are none. Rather than trying to over-quantify risks or find 100% solutions involving scrutiny of every research paper, what is required is a community capable of flagging up issues and willing to do so. Interventions with the aim of restricting availability, preventing the research, or preventing its dissemination all have drawbacks. The alternative is the idea of mitigation, which is to say pursuing science in a way that raises fewer risks.

## Ensuring appropriate risk assessment at the outset of research

“ Appropriate risk assessment should be part of the first phase of the research. Much work needs to be done to identify appropriate risk assessment factors”

7. How can sound risk assessments be made at the outset of the research, and potentially also during it as the research evolves and possible new security considerations emerge? The recreation of the 1918 influenza virus is of strong relevance here. Over several centuries, the influenza virus has episodically taken a catastrophic toll on humans and other species, and it is natural that biologists wanted to understand the nature of the virus and how detection and diagnostics can be enhanced. The 1918 strain was particularly important as it was estimated to be responsible for 50-100 million deaths globally. Whilst subsequent flu pandemics have

not had the same death toll, in part because of the advent of influenza vaccines and improved public health capabilities, flu pandemics remain a cause of concern, particularly as influenza strains have a capacity to introduce mutations from species and spread with relative ease. This is compounded by the limitations in viral drug resistance capacity and the limited number of antiviral drugs available.

8. In the case of the 1918 influenza recreation, CDC researchers were able to sequence the DNA of the virus (H1N1) obtained from a lung tissue sample from a human body frozen in the Alaska permafrost. The sequence was isolated, characterized and then published in the journal Nature, and within a week a different set of scientists had used sequences to reengineer the virus. The research was sufficiently controversial to capture the attention of the media, with the New York Times suggesting this was "extremely foolish. The genome is essentially the design of a weapon of mass destruction".

9. The US government subsequently turned to the NSABB to discuss the benefits and risks of the research, and to make a recommendation. On the one hand, the research generated information important to pandemic preparedness efforts and the development of medical countermeasures; on the other hand, it also generated new safety and security risks, with the manuscripts presenting information which could be exploited to cause harm. Those that argued the sequencing and synthesis of the genome was important to improve medical countermeasure development also argued that it would be difficult to recreate. The group also addressed the consequences of the virus escaping and concluded if it did get out the effects would not be as deadly as 1918. It was pointed out that very few studies looked at immunities across population (something that was not done until 2009), and at the time there was a limited scope for an evidence-based decision. Moreover the research had gone through biosafety reviews at the institutional level.

10. The broader risk analysis came from individuals who remained concerned that it would be easy to recreate the virus through reverse engineering and such a step could lead to the virus escaping from laboratories into the environment. Such sentiments are not unreasonable, particularly in light of the fact that if accidents happen in the laboratory and with a pathogen that is highly transmissible, there are clearly dangers to the work. Indeed some scientists argued that the experiment should never have been carried out at all; others analysed the results and methods and suggested the experiment had not been sufficiently transparent and rigorous.

11. Between 2005 and 2011 the NSABB established a risk/benefit methodology with a series of well-vetted questions people needed to answer in order to publish research in certain areas. However, as a community, life scientists are by no means homogenous and different perspectives tend to produce different answers; although in the case of the 1918 Spanish Flu, risk/benefit analysis clearly reached some agreement that the paper and the details should be published with only minor modifications, recognising some of the biosecurity aspects of the research.

12. Similar arguments were raised during the H5N1 controversy. Cited risks have included the potential insights for approaches (but not blueprints) for producing a pool of recombinant H5N1 viruses with high pathogenicity and sustained, respiratory transmissibility in mammals. On the benefit side the information can be claimed as important in preparedness efforts and enhancing scientific understanding of transmission, as well as through advancing antivirals for public health, but without any proof at the start of a project nor at the end whether this will become reality. Even if some scientists regard the benefits as relatively clear, it may be less clear how the security benefits could be articulated in a manner which generates productive discussion.

13. It is clear that there are unintentional risks through, for example, accidental exposure, as well as the possibility of deliberate attacks including those by 'lone wolf' operators (something notoriously difficult to detect or deal with). However such risks are believed

to be difficult to account for in an empirical or scientific manner, and the discussion is further hampered by difficulties in communicating risks to non-scientists. Moreover there are complicating factors such as concerns caused by withholding research information and or impeding preparedness.

14. Developments in life sciences show that existing instruments are insufficient. Moreover the increasing convergence of biology, nanotechnology and chemistry will lead to new challenges. The Dutch regime has measures such as export licensing, a legal regime for biosafety (but not for biosecurity) and a biosecurity bureau that provides an overview on biosecurity and advice to researchers and research institutions. Turning the legal biosafety officer into a bio-risk professional responsible for safety and security is also being considered. In addition there is also a code of conduct for scientists, which is currently being reconsidered. However the components of the regime are still fragmented.

15. The Dutch drew from the H5N1 case that at least four pillars of policy need to be developed: a comprehensive assessment of risks; intelligent distribution of responsibilities within the scientific community; public trust in science and in government; and in particular, international collaboration.

16. The first of these needs to be multidisciplinary and independent, particularly as there are several different interests (economic, scientific and security) at stake. Advisory committees or review committees that conduct initial assessments are meant to be the first line of assessment and conducted at the point of the research being considered, perhaps even before the funding stage. In their consultations about H5N1, the Dutch government encountered questions of strong relevance to possible future cases: were security concerns (sufficiently) addressed in the initial assessment and during the research itself? To whom was (and should) advice about the (proposed) research be addressed: to the scientists or host institution; to the government; to funding agencies; or to a combination thereof? Should the advice be binding or non-binding and should it be made public? To what extent are publications about dual use research subject to export controls?

17. Current policies rely on the scientists, which is inevitable given that a great deal of research takes place outside the view of governments. However, that reliance depends upon the right sensitization of the community, and here there is room for improvement. Guidelines and codes of conduct are required, but clarity on which research is of concern and better tools for the risks assessment process are also needed. At the moment things are still developing, but the process of public debate on science keeps people aware of the issues and on their toes. The Dutch aim is the promotion of a high-trust high-penalty system, in which scientists self-regulate. National differences remain, however, and many governments have yet to look at this in any detail.

18. Other options that might be considered include the possibility of building in a sort of insurance system. In medical experimentation committees, for example, it is common practice that the interests of the patient are guaranteed before funding is requested. In this context medical experimentation committees are not concerned about burdens on scientists but more general burdens on people. Could parallels be drawn in scientific research? This also points to how fundamental the social contract between science and society is to scientific progress. This can sometimes be overlooked or undervalued, but the idea of building legitimacy into the research process is important.

19. Appropriate risk assessment should be part of the first phase of the research. Much work needs to be done to identify appropriate risk assessment factors relevant to DURC, taking into account the wide range of possible security concerns. In the future, a broader approach to risk could assess physical safety; economic security costs; diplomatic security; social and political stability; fear and anger and risk of research leading to the diminishing trust in government. It should also look at probability and take into account possible actors motives as well as intelligence on terrorist actors. Current DURC risk assessments have been largely “risk-benefit” analyses, and there is

a need for much more comprehensive and quantitative risk assessments that specifically evaluate what could go wrong with certain research. The assessment should not be left solely to researchers and we need to incorporate all bodies and have a debate including governments which are responsible for crisis management and therefore need to consider responses.

" At the heart of the dual-use dilemma is the transfer of knowledge, and thus far there are no internationally valid standards on this. "

## Managing security risks during the research process

20. If the risk assessment raises concerns, how can they be managed during the research process? What happens if a risk is identified at a later stage, during the conduct of an experiment, perhaps as a result of the experimental design, rather than the actual research goal? What are the options for interventions, and are these the same or different in regions with lower capacity?

21. The much cited case of mousepox virus research is a good example of how security issues can arise unexpectedly during research. An Australian research group, working on making mice infertile as a pest-control measure, unexpectedly produced results that in effect demonstrated how to modify poxviruses to break through smallpox immunization. The case did highlight some of the issues affecting sound risk assessment prior to research: some labs in the same institution published papers arguing that the outcomes ought to have been predicted. In light of the fact that the aim was to produce an agent for deliberate release, it can also be argued that the researchers could have considered in the programme the implications of releasing an organism they could not control.

22. In Denmark, the government responded to the mousepox case with new regulations, and now issues licenses permitting companies to work on dual use research. To obtain a license the company must have a legitimate purpose and be able to demonstrate a need to work with controlled materials. Licensees need to submit vulnerability plans and the Government must then appoint a biosecurity officer for mandatory training. Regulation is not aimed at any sector, but anyone who works with materials with both civilian and military personnel included. It is important to stress that the security requirements are not voluntary, but mandatory. They do not just cover physical materials but also intangibles which can be used in biological weapons production without further modifications, with serious potential for misuse, or with more general dual-use potential.

23. Decisions are thus solely related to whether research has weapons potential. It is impossible to check all research activities, and the responsibility for screening lies with the individual scientists in charge. They are required to determine whether there is potential for misuse and, if so, need to get in touch and develop a risk management plan. This of course requires trust and dialogue between communities, as well as an appreciation of weapons potential which is often quite difficult to assess. In this regard it may be useful to look to institutions like the Centre for Biosecurity and Biopreparedness (CBB), which develops and maintains knowledge about weaponisation projects to inform this discussion. Another way is placing biosecurity in a broader context of social responsibility.

24. Elsewhere, however, standards vary and there are no internationally accepted ones. At the heart of the dual-use dilemma is the transfer of knowledge, and thus far there are no internationally valid standards on this. The issue needs careful handling, and it functions differently to standards on biosafety and biosecurity that cover products and technology. Neither is there a common understanding on how to conduct sound risk/benefit analysis; this is an issue between different states but also between different communities (scientific, security etc). That in turn makes it likely that high-risk and low-benefit experiments will be conducted in the future, unless a way of addressing the problem can be found. The more intangible standards that apply to knowledge may be established via a bioethical code, that could be internationally agreed upon. Such a code could set out scientists' responsibilities, beginning with the basic principle of 'first,

do no harm', and then categorise experiments of concern, advise on risk/benefit assessment and risk mitigation, and advise on supporting structures such as ethical review committees.

25. The WHO has conducted a multi-stakeholder discussion on various aspects of DURC, including the potential for codes. The general opinion that has emerged is that a clear set of principles will be useful and potentially helpful, but should not be too detailed. It would be desirable to develop overarching principles and tools but also to develop locally accepted codes and practice. On concordats and codes of conduct, the IAP framework of ethical principles of 2005 was signed by 68 academies; moreover there are already many codes and frameworks. Perhaps the issue is promulgation and adoption in a sustainable manner.

26. Such codes do not have to be fixed; they can be living documents based on the founding principle of 'do no harm'. Cultural definitions and nuances need to be considered for example across different entities within a federal system or in circumstances where universities maintain independence. Even when the cultural understanding of risk is the same there will be differences in the process of implementation, particularly in economically poor countries, where it may be difficult. It could be useful to develop general ethical guidelines that can then be used to develop specific regulations. The measures should be tailored to the lab they are being used in.

27. When trying to generate academic and private sector support and commitment, it is important to understand that both universities and companies want to be seen as socially responsible, and as entities that put value into society. Moreover, the consequences are serious for those discovered to be assisting proliferation; one company lost its American market because it was suggested they transferred dual-use technology to a rogue country. Losing markets is a strong incentive to adhere to established rules. Nonetheless, this remains a challenge, since large pharmaceutical companies often have interests and clients in multiple countries worldwide: how then do you control across different regulatory systems? Part of the answer lies in self-regulation: if large companies are farming out research to smaller countries, they need to check who they are outsourcing to and what standards are in place.

" For many, the point of publication is well past the point at which intervention is either desirable or likely to be effective"

## Issues around publication

28. The recent case of H5N1 gain-of-function research and its publication in Nature and Science raises the issue of what action can/should be taken where the real issues arise around communication of findings. How to assess desirability and conditions of publication? Who makes the final call? Is the point of publication simply too late for meaningful intervention? Does publication pose only security questions, or should safety issues (e.g. from replication of an experiment on a wider basis) also be considered?

29. This is the stage at which intervention becomes very tricky. Prior to this point, decisions can be made not to fund, to modify research and papers, to classify, or to go into the research process on the understating that it will not be disseminated. All of those options are off the table at the point of publication; we are left only with the question of whether to communicate, and if so how. Circulation can be limited, results redacted, or extra information on countermeasures published. Other options are delaying publications, where the rationale is to buy time to address vulnerability.

30. For many, the point of publication is well past the point at which intervention is either desirable or likely to be effective. The information has probably already been shared through the peer review process that reputable journals employ. Moreover, most science is incremental rather than revolutionary, and so one of the questions you need to ask is what crosses the line of concern. Individual papers rarely create huge discontinuity with existing knowledge, but instead are building on the work of others. This, of course, is a general problem in assessing research of concern.

31. Intervening at the publication stage has other drawbacks. Firstly, there is a generic interest in new scientific research being published, including that with potential for concern. Over-active intervention may drive researchers towards science that is completely non-controversial. Secondly, publishing is professionally essential for researchers, in science as for all academic disciplines, and a steady flow of published, peer-reviewed work is part of progressions of an individual's career. There is particularly great pressure on early career researchers who have to 'publish or perish'. The powerful drive to publish means that journals, particularly high-status ones, have heavy backlogs of articles waiting for peer review and/or publication, and to avoid this more researchers are turning to lower-ranked but less over-subscribed journals where the peer review process can be less formal. This in turn means that it can be less easy to spot potentially worrisome research at the peer-review stage. Publishing is also linked to funding requirements and institutions that fund research often obligate scientists to put results and methodologies into open access databases.

32. A further problem here is with the definition of fundamental research. Work on fundamental research is generally exempt from export controls, but the space between fundamental and applied research is not always clear. Scientists have to declare why their work has important applications and this is perhaps one of the tests that can be applied in assessing DURC implications. There are also voluntary options wherein scientists agree to redact or limit the distribution of their work.

33. Regulations tend to focus on applied research that is published. One consequence of this is that the US government can impose conditions on certain publications, and this can be extraordinarily difficult for US institutions, or institutions that receive US government funding. It is only because of fundamental research exemption that they don't have to get a license; if they do need a license, it may take 90 days to obtain one. This also applies to discussion at conferences and restrictions upon foreign nationals. Colleagues from other countries indicate that there are a number of US universities that have long standing policies which say they will not accept funding if there are restrictions on publishing. Stanford University has turned down money in these areas because of that policy. An interesting question is whether US universities will say they will not permit faculty to participate in that research. A second order question is whether this results in work not being done by the best researchers.

## Outreach, engagement and generating support

" Risk assessment prior to research, security management during it, and resolving concerns around publication all rest on better recognition of security concerns among scientists, more effective assessment and articulation of security concerns, and securing better buy-in by scientists"

34. Risk assessment prior to research, security management during it, and resolving concerns around publication all rest on better recognition of security concerns among scientists, more effective assessment and articulation of security concerns, and securing better buy-in by scientists. This is increasingly urgent as we also face the broad dissemination of knowledge and equipment, and the emergence of DIY biotechnology. We thus need to find a middle ground as we cannot curtail research progress, but can place certain checks and balances in the way to avoid huge vulnerabilities. To do this, though, requires the buy-in of the scientists themselves, and to recognize the scientific community as part of the solution.

35. That this is necessary was underlined in 2004-05 by research carried out by Malcolm Dando and Brian Rappert, who conducted some 90 biosecurity seminars in 13 different countries around the world and found a pervasive lack of awareness on DURC issues. It seems clear that these issues are not well embedded in the education of life scientists. Despite regular calls for the development of a 'culture of responsibility' and 'ethics training' among scientists, it may be that this misrepresents the problem. The scientific community does not suffer from a deficit in responsibility or ethics, but it does appear to suffer from a deficit in awareness. Understanding this is important in designing an effective engagement policy; without it, even a successful engagement risks generating a culture of compliance without understanding.

36. What is required is a system for sensible voluntary compliance, and this will be greatly

assisted if the security policy community can develop a mechanism for expressing and quantifying their concerns in a way that allows them and scientists to address this together. Thus a co-operative effort, done in as transparent a way as possible, to refine, review and better judge the risks should be a priority. Increasing the interaction between different sectors could go some way towards getting to a common language, which would facilitate the development of compatible methodology and mutually accessible risk assessments.

37. Risk can be thought of as the function of likelihood of occurrence and the consequence of occurrence. Quantitative assessment sounds attractive because it feels evidence-based and hence more dependable and less open to counter-argument. However, the chances are that firm statistical data will be hard to come by, and that the sort of risks inherent in dual-use biological research cannot be quantified easily (which is not to say that they cannot be quantified at all).

" Subsequent interventions should also have stronger legitimacy if the preliminary risk assessment has the right collaboration between sectors."

## Conclusion

38. In very broad terms, three thematic conclusions stand out. Firstly, the complexity of the issue is daunting, multi-layered, and likely to grow not shrink. Secondly, the security risk surrounding dual-use biological research appears to be generally recognised by those who are aware of the issue, but widening and targeting such awareness is still a challenge, as is reducing scepticism. Thirdly, this issue can be managed and ameliorated, but not resolved. Risk and hazard, and perhaps threat, is something that will have to be lived with not discounted; the price to be paid for the benefits of biological research.

39. More specific conclusions can be grouped under those three headings. To start with the first of the list, the complexity of the issue, this is a recurrent theme in any discussions about dual-use research. In terms of scientific/technical aspects, the growing multi-disciplinary character of biological research means that any regulatory management system that is developed to concentrate on work involving certain dangerous pathogen may fail to capture much future DURC. This also has implications for outreach and education: if the security community's engagement with dual-use biology, and the researchers who work on it, is focussed solely at, say, microbiologists, then it will be only a partial engagement with the field and thereby will be limited in effectiveness. Increasingly 'biology' is being done by engineers, physical scientists, mathematicians and information technologists.

40. In terms of sectorial aspects of the complexity, industry and the commercial sector have thus far played a limited or non-existent role in debates about dual-use research. Initiatives such as the proposals for an international concordat must therefore involve industry, and take account of the range of companies, small to large, involved in the biotechnology industry. It will also need to involve the foundations and councils who fund biological research, who play a pivotal role in determining the criterion for receiving funding and steering the research (and potentially the methodology employed) as well as monitoring outcomes and unexpected results. Similarly there may also be a role for social scientists, ethicists, and the legal and law enforcement communities.

41. This all means that dealing with DURC cannot be done by any one community, and no single set of actors can or should (or would want to) own this problem. An assessment of what such a process of knitting together disparate communities would look like should thus be a priority.

42. Turning to the second theme of how to frame the issue, debate continues as to what definitions or understandings of DURC are appropriate to the discussions. It can be defined in relatively narrow terms as work that involves a limited set of pathogens and toxins in a few categories of experiments; or in wider terms, such as those employed by the US government, as encompassing all "life sciences research…that could be directly

misapplied to pose a significant threat". The choice is between a narrow definition that has the benefits of clarity of focus and relatively manageability, but may risk losing sight of some research of concern; or being more all-encompassing at the risk of losing clarity. This is a debate about the RC in DURC: if most or all biology can be used for dangerous purposes, at what point do you become concerned? At one level this is about risk assessment which is discussed below, but at another it is a definitional issue with important implications.

43. A review of the three case studies examined at the conference indicates that national approaches to management of DURC differ widely, even within the EU: varying from regulation and licensing of research identified as of dual-use concern in some countries, to reliance largely on scientific self-regulation and shared responsibility in others. However, the organisation and scale of the biological research effort also differs in different countries, and therefore a global one-size-fits-all approach to DURC management may not be appropriate. Thus, in considering international approaches, the development of a set of guiding principles may be more appropriate, not to mention more feasible, than agreement on a set of global standards.

44. The third theme for conclusions is how to develop a coherent and effective response to the challenges of DURC; who owns the decision to intervene and where and how can the intervention be made. To restate an earlier point, it seems very clear no single sector or individual can own the decision; therefore the final call will need to be through a collaborative process, except in very clear-cut cases such as where there is definite evidence of mal-intent.

45. On the issue of where to focus DURC risk assessments and interventions, the publication stage appears to be the least desirable place. This is partly because it is reasonable to expect significant risks to have been identified well before the point of dissemination. Moreover, practical difficulties also make this very problematic as a place to intervene: it is increasingly difficult to restrict publication on the internet, and in some countries there may be legal obstacles to preventing publication. Moreover, the use of dual-use export controls to prevent publication of DURC is problematic due to differing interpretations of what constitutes intangible technology subject to control. However, it is necessary to review manuscripts submitted to journals for security concerns: to take expert advice where concerns are raised and, in some cases, employ redaction of sensitive information before publication.

46. Collaborative risk assessment and the acknowledgement of inherent uncertainties and different costs benefit analysis at the start of research is potentially the most fruitful and effective stage, and should involve security input where appropriate. If this can be got right at the outset, with the requisite involvement and hence approval from the different sectors, then a great deal of the potential for dispute can be removed, and potential security problems can be addressed before they arise. Subsequent interventions should also have stronger legitimacy if the preliminary risk assessment has the right collaboration between sectors.

47. So sound and cooperative risk assessment at the outset of research, and (drawing on that) codes of conduct/concordats during it, seem to offer a productive way for the problems of dual-use biology to be effectively managed. Scientists recognise the need for a greater understanding of the security risks and threats, but largely favour a quantitative approach to assessment of the risks of dual-use research. While development of quantitative tools would be useful, ultimately the final call will require qualitative assessment and informed judgement, also taking into account the benefits of the research. A 'through-life' risk assessment process could be established as part of the experimental design to ensure that any unexpected results are addressed.

**James Revill and Mark Smith**
Wilton Park | December 2013

Wilton Park reports are brief summaries of the main points and conclusions of a conference. The reports reflect rapporteurs' personal interpretations of the proceedings – as such they do not constitute any institutional policy of Wilton Park nor do they necessarily represent the views of the rapporteur.

Should you wish to read other Wilton Park reports, or participate in upcoming Wilton Park conferences, please consult our website www.wiltonpark.org.uk

To receive our e-newsletter and latest updates on conferences subscribe to https://secure.wiltonpark.org/en/secure/subscribe?action=subscribe