1. Ethics and Dual-Use Research

Michael J. Selgelid

[W]ho is most capable of treating friends well and enemies badly in matters of disease and health?

A doctor.

…

Isn't the person most able to land a blow, whether in boxing or any other kind of fight, also most able to guard against … And the one who is most able to guard against disease is also most able to produce it unnoticed?[1]

Nuclear physics and contemporary genetics

One of the most dramatic—and ethically problematic—episodes in the history of science involved the making and use of the first atomic weapons. When key discoveries of things like atomic fission and the chain reaction were made during the revolution in physics under way during the first half of the twentieth century, the scientists involved realised that the new knowledge gained might be used for both good purposes (for example, energy production) and bad purposes (for example, weapons making).[2] The first atomic bombs were developed soon after the discovery of the chain reaction in particular. The atomic bombs dropped on Hiroshima and Nagasaki during World War II caused hundreds of thousands of deaths.[3] Since that time vast resources (which might have been used for other purposes) have been devoted to nuclear weapons production, and humanity has lived under the threat of nuclear holocaust for decades.

An implication of the contemporary scientific revolution in genetics and biotechnology is that the life sciences, at present, are in a situation analogous to that of physics when the key discoveries mentioned above were made. The very same biological discoveries that might be used to benefit humanity (for example, via advances in medicine) can sometimes also be used to cause harm (for example, via biological weapons). In some cases, the harms in question could be catastrophic. This is thus a key moment in the history of biology—and a crucial time for ethical decision-making and policymaking in the life sciences.

---

1 Plato 1992, *Republic*, G. M. A. Grube (trans.), C. D. C. Reeve (rev.), Hackett, Indianapolis, pp. 5–9.
2 Schweber, S. S. 2000, *In the Shadow of the Bomb*, Princeton University Press, Princeton, NJ.
3 Rhodes, R. 1986, *The Making of the Atomic Bomb*, Simon & Schuster, New York.

Research with potential to be used for both good and bad purposes is now commonly referred to as 'dual-use research'. While almost any knowledge and technology can be used for both kinds of purposes, the expression 'dual-use research of concern'[4] is used to refer to research that can be used for especially harmful purposes—that is, where the consequences of malevolent use would be potentially catastrophic. Of particular concern are advances in genetics that might enable development of a new generation of biological weapons of mass destruction. (In the remainder of this chapter, I will use the expression 'dual-use research' as shorthand for the expression 'dual-use research of concern'.)

Controversial cases

Such danger is illustrated by the following controversial studies that have been published during the new millennium.

The mousepox experiment

Australian scientists used genetic-engineering techniques to insert an interleukin (IL-4) gene into the mousepox virus. Their aim was to develop a strain of mousepox that would make mice infertile—and thus provide a potentially powerful new means of pest control. They unexpectedly discovered, however, that the altered virus killed both mice that were naturally resistant to and mice that had been vaccinated against ordinary mousepox. They published these findings, along with description of materials and methods, in the *Journal of Virology* in 2001.[5] A danger is that the same techniques might enable production of vaccine-resistant smallpox. Smallpox is one of the most feared biological weapons agents—and vaccine is our only defence against it.

Synthetic polio

Following the map of the polio genome (published on the Internet), American scientists (via mail order) bought and strung together corresponding strands of DNA. Addition of the synthesised genome to 'cell juice' (a solution containing cellular ingredients, but no live cells) led to production of 'live' polio virus that paralysed and killed mice. The scientists' aims were, inter alia, to show that

4 National Science Advisory Board for Biosecurity (NSABB) 2007, *Proposed Framework for the Oversight of Dual Use Life Sciences Research: Strategies for Minimizing the Potential Misuse of Research Information*, National Science Advisory Board for Biosecurity, Bethesda, Md, <http://oba.od.nih.gov/biosecurity/pdf/Framework%20for%20transmittal%200807_Sept07.pdf> (viewed 7 April 2013).

5 Jackson, R. J., Christensen, C. D., Beaton, S., Hall, D. F. and Ramshaw, I. A. 2001, 'Expression of mouse interleukin-4 by a recombinant ectromelia virus overcomes genetic resistance to mousepox', *Journal of Virology*, vol. 75, no. 3, pp. 1205–10.

such a feat would be technically possible and to demonstrate that viruses are ultimately just chemicals.[6] They published their findings, along with description of materials and methods, in *Science* in 2002.[7] A danger is that similar techniques might be used to create biological weapons agents, such as smallpox, which bioterrorists and/or state-sponsored biological weapons programs might not otherwise be able to access easily.

Reconstruction of 1918 Spanish flu

American researchers used similar techniques (to those used with polio) to reconstruct the 1918 Spanish flu virus. The virus they created was more deadly to chickens than any flu virus that had previously been studied. The purpose of this research was to yield knowledge that would facilitate protection against possible future influenza pandemics—for example, vaccine production. They published their findings, along with description of materials and methods, in *Science* in 2005.[8] A danger is that malevolent actors (for example, bioterrorists) might use the published information to create and unleash the virus in question—which was responsible for one of the worst epidemics in human history, killing 20 to 100 million people over the course of a year or two.

Transmissible H5N1

The most recent and, to date, most controversial dual-use life-science research involved the study of H5N1 (avian) influenza transmissibility among ferrets, which provide the best model for influenza in humans. While it is estimated that H5N1 kills 60 per cent of humans infected, it is not (currently) transmissible between humans. Researchers in the Netherlands and the United States thus conducted experiments that aimed to determine whether H5N1 might develop into a human-to-human transmissible strain. Genetic engineering of the virus and 'passaging' of the altered virus between ferrets led to creation of strains that were airborne and easily transmissible among ferrets—thus indicating that natural evolution of a human-transmissible strain of H5N1 might be possible. Much debate surrounded the question of whether or not this research should be published in detail. On the one hand, it was argued that publishing these studies was important because this would facilitate vaccine development and/or surveillance of relevant changes to H5N1 occurring in nature. The hope with regard to surveillance is that this would enable earlier detection of emerging

6 Selgelid, M. J. and Weir, L. 2010, 'Reflections on the synthetic production of poliovirus', *Bulletin of the Atomic Scientists*, vol. 66, no. 3, pp. 1–9.

7 Cello, J., Paul, A. V. and Wimmer, E. 2002, 'Chemical synthesis of poliovirus cDNA: generation of infectious virus in the absence of natural template', *Science*, vol. 297, pp. 1016–18.

8 Tumpey, T. M. et al. 2005, 'Characterization of the reconstructed 1918 Spanish influenza pandemic virus', *Science*, vol. 310, no. 5745, pp. 77–80.

pandemic strains and thus earlier implementation of protective measures. On the other hand, others argued that a human-transmissible strain of H5N1 could kill millions or perhaps even billions of people if produced and unleashed by bioterrorists or other malevolent actors (or in the event of accidental release from research laboratories). The US National Science Advisory Board for Biosecurity (NSABB) in December 2011 recommended that the findings of these studies be published, but that detailed description of materials and methods be omitted from the published articles.[9] After a highly publicised World Health Organisation (WHO) meeting reached the opposite conclusion in February 2012—recommending that the studies eventually be published in full[10]—the NSABB reversed its initial decision in March 2012.[11] Revised (fully detailed) versions of the papers in question were finally published in *Science*[12] and *Nature*[13] in June 2012. Part of the reason for the NSABB's reversal of decision apparently involved misunderstanding regarding whether or not the airborne transmissible strains were deadly to ferrets. While much of the initial public discussion indicated that the strains produced in the Netherlands were just as deadly as ordinary H5N1, it was later revealed that airborne strains created in the lab were not deadly to ferrets. Another cited reason for the NSABB's reversed decision is that the revised manuscripts they eventually approved provided additional data regarding ways in which the studies might have public health benefits—for example, via surveillance.

Levels of governance

The dual-use phenomenon requires ethical decision-making by various actors at different levels of the science governance hierarchy. Individual scientists (insofar as they are at liberty) must decide what research to conduct and/or publish. Research institutions (insofar as they are at liberty) must decide how to regulate potentially dangerous research within their confines; how to educate

9 US National Institutes of Health 2011, 'Press statement on the NSABB review of H5N1 research', <http://www.nih.gov/news/health/dec2011/od-20.htm> (viewed 7 April 2013).

10 World Health Organisation (WHO) 2012, 'Public health, influenza experts agree H5N1 research critical, but extend delay', <http://www.who.int/mediacentre/news/releases/2012/h5n1_research_20120217/en/> (viewed 7 April 2013).

11 National Science Advisory Board for Biosecurity (NSABB) 2012, 'March 29–30, 2012 meeting of the National Science Advisory Board for Biosecurity to review revised manuscripts on transmissibility of A/H5N1 influenza virus', <http://oba.od.nih.gov/oba/biosecurity/PDF/NSABB_Statement_March_2012_Meeting.pdf> (viewed 7 April 2013).

12 Herfst, S., Schrauwen, E. J. A., Linster, M., Chutinimitkul, S., de Wit, E., Munster, V. J., Sorrell, E. M., Bestebroer, T. B., Burke, D. F., Smith, D. J., Rimmelzwaan, G. F., Osterhaus, A. D. M. E. and Fouchier, R. A. M. 2012, 'Airborne transmission of influenza A/H5N1 virus between ferrets', *Science*, vol. 22, pp. 1534–41.

13 Imai, M., Watanabe, T., Hatta, M., Das, S. C., Ozawa, M., Shinya, K., Zhong, G., Hanson, A., Katsura, H., Watanabe, S., Li, C., Kawakami, E., Yamada, S., Kiso, M., Suzuki, Y., Maher, E. A., Neumann, G. and Kawaoka, Y. 2012, 'Experimental adaptation of an influenza H5 HA confers respiratory droplet transmission to a reassortant H5 HA/H1N1 virus in ferrets', *Nature*, vol. 486, pp. 420–8.

researchers working there (regarding dual use and/or ethics); what laboratory security measures to put into place, and so on.[14] Professional societies must make decisions about the development, promulgation and/or enforcement of ethical codes of conduct for scientists—and/or decisions about relevant (ethical) education of their members. Publishers must make decisions regarding processes of review of papers posing dual-use dangers—and they must ultimately decide which papers to publish. National governments must decide what research to fund, and the extent to (or manner in) which things like research review, publication review and/or relevant education of scientists will be mandated—and they must make decisions about the extent to which controls should be placed on access to potentially dangerous materials. International governance bodies (such as the WHO), finally, must make relevant decisions concerning global policy—for example, whether or not there should be international guidelines regarding dual-use research, or oversight thereof, and/or what the content of such guidelines should be.

Ethical dilemmas

In all cases the decisions will be difficult. On the one hand, responsible actors will want to take actions that will promote the development and use of beneficial science. On the other hand, they will want to take actions that will prevent the malevolent use of science (which might sometimes require avoidance of generation and/or publication of potentially dangerous information). An implication of the dual-use phenomenon, however, is that it is inherently difficult to achieve both goals at the same time in the cases where the very same research that is likely to be beneficial might also be used to cause harm. In the case of governmental decision-making, for example, a laissez-faire approach to scientific governance might facilitate scientific advance and the benefits thereby enabled—but it might also lead to especially dangerous research getting done and/or published. A more restrictive approach, on the other hand, might prevent generation and/or publication of dangerous information—but it might also stifle beneficial scientific advance at the same time. Hence the expression 'dual-use dilemma'.

In any case, it is important to recognise that key decisions posed by dual-use research are inherently *ethical* in nature. The decisions faced by the various actors enumerated above largely concern: 1) the responsibilities of the actors in question (for example, to what extent would a scientist be responsible if her research is used to cause harm?); 2) issues of how one should go about promoting

14 Research institutions, of course, already do such things to varying degrees—and numerous relevant measures (for example, regarding biosafety) are required by law. The point here, however, is that additional new measures are required to address dual-use research in particular.

benefits while avoiding harms or reducing risks (for example, should a paper be published if this has a good chance of promoting a significant amount of human wellbeing but a small chance of causing disaster?); and/or 3) questions about values and value conflict (for example, how should governments strike a balance between the goal to promote scientific freedom/progress and the goal to promote security?). Issues regarding responsibilities, harms, benefits and values—and, ultimately, what ought to be done—are exactly the kinds of things that ethics is about.

In the meantime, however, much of the debate about dual-use research has involved scientists and security experts rather than ethicists in particular.[15] While much bioethical discussion has focused on research ethics and the ethical implications of genetics/biotechnology, it is ironic that relatively little bioethical discussion has, to date, focused on dual-use research in particular. This is unfortunate because, given the potential for catastrophic consequences, dual-use research is surely one of the most important ethical issues regarding research and genetics/biotechnology.

This volume

The chapters in Part I of this volume, 'Dual Use in Context', map the terrain of dual-use issues emerging in various areas of life-science research—that is: nanotechnology (Chapter 2, Jim Whitman); neuroscience (Chapter 3, Valentina Bartolucci and Malcolm Dando); synthetic biology (Chapter 4, Alexander Kelle); agriculture (Chapter 5, Simon Whitby); and tuberculosis (Chapter 6, Nancy Connell).

The chapters in Part II, 'Ethical Frameworks and Principles', explore the relevance of various existing philosophical frameworks and/or tools of ethical analysis to the dual-use problem (in cases where their relevance to dual-use problems has received relatively little previous exploration). In particular, they examine the impact of environments and institutions on moral development and ethical reasoning (Chapter 7, Judi Sture); the ethics of weapons research in general (and the relevance thereof to dual-use research in particular) (Chapter 8, John Forge); application of 'rational decision theory' to the dual-use problem (Chapter 9, Thomas Douglas); the relevance of the 'doctrine of double effect' to the dual-use problem (Chapter 10, Suzanne Uniacke); implications of uncertainty for dual-use decision-making (Chapter 11, Michael Smithson); collective-action problems associated with dual-use research (Chapter 12, Seumas Miller); the

15 Selgelid, M. J. 2010, 'Ethics engagement of the dual use dilemma: progress and potential', in B. Rappert (ed.), *Education and Ethics in the Life Sciences: Strengthening the Prohibition of Biological Weapons*, ANU E Press, Canberra, pp. 23–34.

relevance of just-war theory to the dual-use problem (Chapter 13, Koos van der Bruggen); and the relevance of the 'precautionary principle' to dual-use decision-making and policymaking (Chapter 14, Steve Clarke).

Parts III and IV, 'Ethical Practices' and 'Ethical Futures', consider existing, developing and future practices and policies regarding ethics and dual-use life-science research—that is, the prospect of self-regulation by scientists (Chapter 15, David Resnik); lessons learnt from the history of nuclear physics (Chapter 16, Nicholas Evans); dual-use governance in developing countries (Chapter 17, Louise Bezuidenhout); the responsibilities of individual scientists in the context of incapacitating chemical and toxin agents (Chapter 18, Michael Crowley); and the WHO's project on 'Responsible Life Sciences Research' (Chapter 19, Emmanuelle Tuerlings and Andreas Reis). The Conclusion (Chapter 20) offers a wide-ranging and agenda-setting summary by Brian Rappert.

This volume is a product of a workshop (funded by the Wellcome Trust) on 'Promoting Dual Use Ethics' (organised by Michael Selgelid and Brian Rappert) held at the Centre for Applied Philosophy and Public Ethics (CAPPE) at The Australian National University in Canberra in January 2010.