8. Responsible Dual Use

John Forge

Introduction

This chapter addresses the *moral* dimension of dual-use research. To set the scene, I will begin by explaining what I take dual use to be. I understand a dual-use item to be something that has both a good or neutral (neither good nor bad) use or application and a bad use. Three different categories of dual-use items can be distinguished: research, technologies and artefacts.[1] These are clearly different sorts of things. Research is an activity, while technology is a form of knowledge—knowledge of the techniques for the production of artefacts—whereas artefacts are objects. But it is also clear that these items are related: research aims to give us technology, which in turn produces artefacts.[2] It is the last that normally has the immediate impact on the individual, be this good or bad. For instance, research has led to methods for mass-producing cheap ammonium nitrate, which has a 'primary' use as a fertiliser and a 'secondary' use a bomb-making material. It is the substance—the artefact—ammonium nitrate that is the fertiliser/bomb component, not the technology or the research. The import of this is that when we talk of threats or risks associated with dual use, we need to distinguish between different levels of risk or threat, depending on the nature of the dual-use item, and between different ways of dealing with these. For instance, substances like ammonium nitrate can be physically contained, but the knowledge of how to manufacture it cannot.

My suggested definition of a dual-use item is therefore as follows:

> An item (knowledge, technology, artefact) is dual use if there is a (sufficiently high) risk that it can be used to design or produce a weapon, or if there is a (sufficiently great) threat that it can be used in an improvised weapon, where in neither case is weapons development the intended or primary purpose.[3]

---

1 This is the classification that I gave in Forge, J. 2010, 'A note on the definition of dual-use', *Science and Engineering Ethics*, pp. 111–18. This paper was written as a response to a challenge by David Resnik to come up with a definition of dual use that is neither too narrow nor too broad. I do not claim that my contribution is the last word, but I do think it is along the right lines. I also think that it is enough by way of a starting point for the present discussion. I would, however, note that the sense in which artefacts are said to be objects is a special one—see ibid., p. 115.

2 There are of course other kinds of research and other kinds of technologies besides those that aim to produce artefacts.

3 Forge, op. cit., pp. 111–18.

There are several comments in order. First, I equate the bad—secondary, unintended, and so on—use with weapons development. It is certainly possible to broaden the discussion of dual use to include other harmful activities, but I think these would be more controversial.[4] I should stress that by 'weapons development' I include all aspects of the design, production, testing and commissioning of both weapons themselves and all the ancillary structures, platforms and processes that are necessary for their use. Second, the reference to *risk* in the definition points to the possibility of the production of a new or improved weapon. The (by now) familiar examples of techniques of synthesing pathogens are thus risky in this sense. Threats differ from risks here in that they signify the intention to do something harmful, rather than the possibility that something harmful might happen. I have coupled risk with the design of (new) weapons but threats with improvised weapons. States threaten one another, as North and South Korea are doing at present, though the threat here is usually conditional.[5] In regard to dual use, however, I think we are more concerned at the moment with sub-state actors, such as terrorist organisations, using or adapting existing items for improvised weapons, as has been done with ammonium nitrate, than with states arming themselves with improvised weapons, as states do not usually need to improvise weaponry.[6] Finally, research that is intended to produce a weapon and which, fortuitously, has a friendly civilian application, is not dual use according to my definition: the 'bad use' is always the unintended one.

My focus here is with dual-use research, the dual-use item that appears furthest removed from any immediate impact on the individual, and with the responsibilities of those who undertake research that is, or could be, dual use. So if a particular research project has some prospect or likelihood of leading to a bad outcome then what do we expect of those who would conduct the project? Perhaps if the researcher does not intend for the bad outcome to come about, or if she does not foresee that it will, she has no responsibility for it? There are well-known discussions of the doctrine of the double effect that could have some application here.[7]

Another issue is if the likely bad effects seem to be outweighed by the prospect of good outcomes. The consequentialist who sees moral action as reducible to the production of what, on the whole, is the best outcome is likely to be sympathetic to this possibility. I shall touch on these issues in what follows, but

4 And this is very much in the spirit of the early definition of 'dual use' given by the Office of Technology Assessment in the United States and subsequent definitions given by bodies such as the National Research Council. See ibid., p. 112.

5 *If* North Korea shells South Korean islands again *then* the South will retaliate.

6 The extent to which we should therefore be cautious about undertaking dual-use research becomes important. The relevance of the precautionary principle, which would seem to have application to dual use, is the subject of other chapters in this volume.

7 See Suzanne Uniacke's Chapter 10, in the present volume, for a discussion of this.

my main aim is to apply a particular way of understanding the responsibilities of researchers that I proposed in a recent work, *The Responsible Scientist*,[8] to the question of dual use. These responsibilities are both 'backward looking' and 'forward looking', having to do with what researchers are responsible for in the past and what we want to hold them to in the future. Both sorts of responsibility are relevant to dual use. The upshot of the account that emerges is a cautious and restrictive response to dual-use research; other accounts will have other implications. I will not assume much familiarity with the topic of responsibility or with moral philosophy, so much of what follows in section one will be exposition. I hope the implications of this will become clear in sections two and three. As a final comment by way of introduction, what follows is put forward as a contribution to the *ongoing* discussion of dual use. The topic is still, I think, at a stage where a variety of different approaches are worth considering, in regard to both definitions and the moral and policy dimensions. I hope the account of responsibility that I have developed can stand as one such approach.

Responsibility and research

There are two kinds of responsibility: 1) responsibility for past actions, actions that an agent has already completed, or omitted to do—namely, backward-looking responsibility; and 2) responsibility for future actions or omissions—namely, forward-looking responsibility. As we shall see, these concepts are related—and both are relevant to scientific research. It is easiest to begin with a simple example that has nothing to do with science: suppose it is the job of a railway guard to make sure everyone is on the train and the doors are closed before signalling to the train driver to leave the station. If the guard fails to realise that a disabled person will take extra time to get on the train or if he omits to make sure that a carriage door is closed then he will be held responsible for any harm that subsequently occurs. This would not be true of a member of the public. We would hope that a member of the general public would try to assist the disabled person or close the carriage door, but, unlike the guard, she has no 'special responsibility' to do so. In this example, the special responsibility in question can be termed *role responsibility*. It is (a kind of) forward-looking responsibility—that is, one that obliges the person in the role to conform to certain expectations about his or her future actions in a certain sphere: train guards, for example, are supposed to make sure trains are safe. There are other jobs and professions that define roles that are such that the performers are supposed to safeguard others from harm. These are not all of a piece. For instance, those of the medical and legal professions are more complex than those of train guards, lifeguards, and so on, which is one reason these roles have an

8 Forge, J. 2008, *The Responsible Scientist*, Pittsburgh University Press, Pittsburgh.

associated professional ethic. An important point here is that certain failures to foresee or omissions to act that result in harm to others can only be attributed to the agent if there is a corresponding forward-looking responsibility. I am not obliged to keep an eye on toddlers at my local pool, but the lifeguard is, and if she omits to do so then she is open to blame.

In my book, I argued that scientists have a kind of role responsibility, even, and perhaps especially, scientists who claim to be doing pure research.[9] I used this to ground a wide account of backward-looking responsibility for science, one that held that scientists could be held responsible for failures to anticipate where their work might lead, and for omissions to act in certain ways, in addition to being responsible for what they intended to do and for what they foresaw would be outcomes of their work. These are, of course, complicated issues.[10] For instance, it might be asked how could a scientist know where his research might go, what kind of applications it could have in the future? How could Newton know of all the applications of universal gravitation or Einstein of relativity? I did not try to argue that all scientists can foresee all applications of their work, for that would have been a hopeless task, but I did show that there are circumstances where this is possible. For example, had all the information about the properties of the uranium nucleus available in 1939 and 1940 been published, there was a significant probability that this would have led to a more vigorous effort on the part of the Nazis and others to develop an atomic bomb. And as the gap between pure research and applied research narrows, as it is surely doing more and more in an age when funding for science is increasingly tied to applications, it is easier to look ahead. Indeed, to continue the theme of funding, research proposals are very often obliged to list possible applications as outcomes. I will come back to this point.

The basic reason all scientists have a kind of role responsibility is simply because science affects people (SAP) and does so in profound ways. I contend that these effects are not always good. I believe that all weapons of mass destruction, especially nuclear weapons, have done us no good, are terrifyingly dangerous, and that it would have been better had the Manhattan Project never been set up and no-one else had attempted to make atomic bombs.[11] Like all these issues, this is somewhat (but I think in this case not very) controversial: some think that the dropping of the atomic bombs on Japan was a good thing because it ended

9 I understand 'scientist' in a broad sense, not restricted to the natural sciences. Most of the examples in the book are from physics, which is my area of expertise.

10 So complicated that I needed to devote two chapters to deal with them. One of the important steps in the argument here was to show that we can be responsible for what we don't know—we can be responsible for being ignorant, which is at odds with the classical account of the matter by Aristotle. See Forge, 2008, pp. 107–9.

11 I agree with the assessment that it was only in wartime that the vast expenditure would have been justified for the first nuclear weapons project, so without the Manhattan Project there may well have been no nuclear weapons.

World War II, and that nuclear weapons prevented the Cold War from becoming a hot one. As we learn more about the intentions of the former Soviet Union, the latter proposition is becoming increasingly less plausible. (I am also willing to argue, though not here, that *all* scientific research for the ends of weapons innovation is wrong.[12]) Whatever one makes of these particular examples, even someone with blind faith in science would be hard put to maintain that *everything* scientists do has turned out for good. It should be stressed that it is *not* part of the present position that when something has gone wrong then this has been done intentionally or that scientists are always to blame. And it is far from the present position that science either has no effect on us or is bad overall. Clearly, science has done a great deal of good, as is clear when we review the history of medical science.[13]

Once we accept SAP, it seems clear that we would like science to somehow maximise the good outcomes and minimise the bad ones. The first step, or maybe the preliminary step, is to get scientists to accept SAP; by this I mean all scientists and not just those working in applied areas. Surely this should not be too difficult and is probably already true of the majority of scientists. That all scientists should be aware that their work can affect others can be argued with reference to what I call the changed context of science. Indeed, since World War II, all working scientists really should know that no scientific research is pure in the sense that it cannot in principle affect people. Once this is established, we get on to the hard choices: just what kind of research should scientists be encouraged to do, and what to avoid? Philosophers can have some influence on the policy process—at least they can in theory—by suggesting some general guidelines. What I have tried to do is to provide a set of moral principles that could be used to inform the choices made by scientists themselves, though policymakers could also use them. There are different ways in which this can be done, given what I have said so far about responsibility, and in the rest of this section, I will outline the way I prefer.

The reader may be aware that there are two traditions in moral philosophy—traditions that seem sharply opposed. According to the consequentialist tradition, the moral import of an action or choice resides solely in its consequences. The nature of the act itself, or of how we describe or characterise it, counts for nothing. Evidently, there will be issues here for the consequentialist, resembling those discussed above, in regard to how consequences stretching into the future can be determined and weighted with respect to their moral significance. Leaving these aside, a consequentialist normative ethic will specify some

12 See my *Designed to Kill: The Case against Weapons Research*. Dordrecht: Springer, 2012.

13 But even here there is no universal agreement: some strange folk even think that the triple antigen vaccination is wrong. My general point here is that science is a 'mixed blessing' and that is why it is important to determine just what the responsibilities of the scientist are, Forge, 2008, op. cit., pp. 28–31.

property, to be taken as 'the good', and then require moral agents to maximise this property by their actions. Famous nineteenth-century consequentialists like Bentham and Mill took happiness to constitute the good, and thus moral action was such as to maximise happiness, the agent's included. The calculus has to include consequences where some were made unhappy, and this means that all sorts of scenarios can be constructed in which happiness is maximised but some small number of unfortunates is made utterly and forever miserable. Some of these scenarios are portrayed as counterexamples to the theory. My own account is non-consequentialist in that it does not hold that consequences are all that matter.[14] I won't make any further criticism of the consequentialist tradition here.

Non-consequentialist moral philosophy is usually expressed as a set of norms or *rules*. What these are, how these are understood and how they are justified distinguish different non-consequentialist accounts. Modern accounts do not usually interpret the rules as absolutely binding, and here there is a break with the Kantian tradition. Kant famously—or notoriously—said that one should not lie, not even to mislead a maniac looking for one's best friend with malicious intent. The rules contained in modern accounts are thus often said to be prima facie, or to lay down prima facie duties: they are to be obeyed 'in the first instance', unless there are very good reasons not to obey them. As for the content of these rules, I follow those like Bernard Gert who think that they should forbid harming but not require any positive or benefiting action. Here there is evidently a considerable difference with the traditional consequentialist. One good reason to accept our view is that it is possible to impartially refrain from harming everyone, but one cannot impartially benefit everyone. If I have a limited amount of help to give, I must favour some, but moral action is supposed to be impartial. Gert's moral philosophy, which I follow closely here, does not exclude trying to do good; however, good action, which is understood to be equivalent to the prevention of harm, is taken to be the content of moral *ideals*, and hence agents are not bound by rule to prevent harm. It has to be said that these are quite subtle matters and it is hard to do justice to them here. Another way to put the difference between the rules and the ideals is as follows: agents need to justify breaking the rules—recall that these are not absolutely binding—but only encouraged to act in accord with the ideals. Gert refers to his account as *common morality*.

To conclude this discussion of the moral basis of responsibilities of the scientist, we should first note that this way of doing moral philosophy makes a lot of sense. No person in her right mind wants to be harmed: *everyone*, wherever

14 Note that non-consequentialism is not the contrary of consequentialism, which would be that consequences do not matter at all.

they come from, can agree with that.[15] Thus everyone can accept a moral code that proscribes harming, and agree that everyone would be better off with less harming. Doing good in some positive sense is much more difficult to universalise. Moreover, striving not to harm others is surely within the reach of everyone, while looking to maximise the good consequences of one's actions looks just too hard. Now scientists are in a special position because of the far-reaching implications of their work, primarily through technology.[16] Therefore it seems fair to attribute to them special responsibilities, and in line with the proposed moral position, these should enjoin scientists in the first place not to do research that has harmful outcomes. I have taken this to mean above all (and at the very least) that scientists should not engage in weapons research;[17] however, I acknowledge that scientists should be encouraged to do research that will have beneficial outcomes. In *The Responsible Scientist*, I therefore maintained that the responsibilities of the scientist are two tiered. The first 'tier' comprises this demand to avoid research that can have harmful outcomes, while the second encourages scientists to do research that prevents harm. Research in the biomedical sciences could well fall into the latter category, as could research that breaks down racial and sexual stereotypes and biases. These two tiers correspond to the moral rules and the moral ideals of common morality. Most of this section has been devoted to the forward-looking responsibility of the scientist, and I will come back in the final section to see how this applies to dual-use research. The next section, on the other hand, will be concerned mainly with backward-looking responsibility.

Dual use and responsible research

Our interest here is with dual-use research, and not with other dual-use items, although of course dual-use research is worrisome precisely because it can lead to the latter. To begin with, then, suppose that a given line of research, R_1, is already established as dual use. By this I mean that it is *known* that its implications are such that there are bad uses as well as good or neutral ones. It may seem strange that anyone would work on such a project, and indeed it may be hard to find many real-life examples; however, the enrichment or reprocessing of spent nuclear fuel rods to make weapons-grade uranium and plutonium, and also to make new, and high-power, fuel rods counts as an example, and thus

15 Gert holds that common morality gets purchase on all those who are rational in the sense of accepting certain basic beliefs, such as that I am mortal, I can be harmed, I have interests, and so on. See Gert, B. 1998, *Morality: Its Nature and Justification*, Oxford University Press, Oxford.

16 Science as a body of ideas has had, and still has, a huge impact on us; witness the work of the likes of Galileo and Darwin.

17 See Forge, J. 2007, 'What are the moral limits of weapons research?' *Philosophy in the Contemporary World*, vol. 14 (Spring), pp. 79–88; and Forge, 2008, op. cit. In *Designed to Kill* I argue that no one should engage in weapons research.

experiments or even theoretical work on this topic count as dual use. There are two reasons a scientist might think it is morally acceptable for her to work on this kind of project in applied nuclear physics. She might think that in fact there is no bad use, or that if she only *intends* to work on enrichment or reprocessing insofar as it has a benign use in civilian power reactors and that it deals with the disposal of hazardous waste, that is okay. The first of these reasons raises an important general issue about dual use—namely, the basis on which the uses in question are classified.[18] The second raises issues that concern attributions of backward-looking responsibility. I will address this second matter first.

The second issue can be dealt with by what I call the wide view of backward-looking moral responsibility, given in *The Responsible Scientist*. The wide view holds that scientists are responsible for what they foresee as well as for what they intend (and they can also be responsible for what they should have foreseen). The view expressed by our scientist above—that one is only responsible for what one intends—is what John Mackie called the straight rule of moral responsibility and which I call the standard view. The latter term is appropriate, as the view seems to have slipped in as a kind of orthodoxy, being endorsed by Peter Strawson and John Austin, both, like Mackie, renowned Oxford philosophers, but without convincing argument. At first sight it seems plausible, because intentions are given as reasons for action: 'Why do you work on this project?' 'To make better nuclear fuel rods and remove hazardous nuclear waste.' The obvious question, then, is why we should only be responsible for actions that we have a reason for doing.[19] For those who do think that responsibility must be tied to reasons for action, and I am not one of this number, it is easy to see how a foreseen but not intended outcome or side-effect of an action can be included in a reason for action as an explicit qualification. In this example, it can be done by adding 'although I am aware that reprocessed fuel can be used in weapons and this did not make any difference to my choice'. By incorporating foreseen outcomes in this way, we get what I have called the modified standard view.[20] How we might manipulate or change the standard view is really beside the point, for it is simply a commonplace that we *do* hold people responsible for what they foresee as well as what they intend, in spite of what Mackie, Strawson and Austin might

18 I raised the issue of the role of values in the definition of dual-use items and how different sets of values could lead to different judgments about what counts as a dual-use item in Forge, 2010, op. cit., p. 117. I have more to say about this question below.

19 In some ways it would be clearer if this were expressed in terms of interests. Our researcher has an interest in disposing of hazardous waste, but no interest in making material for nuclear weapons; her reason for undertaking the project cites the former as her reason for action. Nevertheless, she sees that her work contributes to the latter end.

20 The modified standard view, and why it is not enough, is discussed at length in Chapter 5 of Forge, 2008, op. cit. Also, it is one thing to recognise that something is a commonplace and quite another to give it a convincing philosophical rationale—hence the length of Chapter 5 and the chapters before and after it.

have thought. And hence the response of our scientist that any other outcomes of this work are beside the point because they are not what she intends will not wash.[21]

On what basis do we decide whether something has both good and bad outcomes or uses, and hence whether it is a dual-use item? And if there are alternatives that lead to different judgments, which do we chose? These are important questions, ones that I cannot address fully here. But for a start we can all agree that bad outcomes are those that are, or have the means to be, harmful. Common morality of course would have this implication. If we think of our moral philosophy as expressing values then common morality elevates the value of not harming above all others, and would not, for example, allow the harming of a few to benefit many. Again, this is surely something that most of us would accept: if we could benefit at the expense of others being harmed then most of us would not accept such a deal if it were offered to us (at least I hope we would not). So I assume again that something like the values expressed by common morality provide a plausible and realistic value system. We will see in a moment how this can affect classifications of items as dual use. First, however, note that there is a difference between harmful acts, acts that directly harm moral subjects, and outcomes of research—the topic of the present discussion—which are such that they could be the means to harm. Provision of the means to harm is not the same as harming, and if we are to try to attach the kind of strong moral prohibitions about harming suggested here to activities that do not directly harm, it seems that we need some further argument.

In fact I don't think that this is too hard to supply in cases of research whose *explicit* aim is to design new and better weaponry.[22] Scientists who engage in this kind of work need *justification* for what they do: they need to give reasons why it is acceptable for them to engage in an activity whose objective is to provide tools whose purpose is to harm others. The typical justification is that we need weapons to defend ourselves from others, from the bad guys. This is similar in form to the typical justification that we give when we feel we need to harm others—namely, that we need to do so to *prevent* harm to others and ourselves. The acceptability of that justification will depend on the circumstances: clearly, inflicting a great deal of harm to prevent a small amount of harm will not do. The harm inflicted should in some way be commensurate with that prevented. But this judgment is more difficult to sustain when it comes to deciding whether making new and improved weapons—namely, making the means to harm rather than harming directly—is justified. Are they really necessary? Will others

21 The challenge, of course, is, as in the previous footnote, to provide the philosophical argument.

22 Not too hard, but quite lengthy: in Forge, 2008, op. cit., pp. 155–8, I argue for a means principle that is such as to transfer responsibility for bad outcomes using artefacts—weapons, for instance—to the designer of the artefact.

acquire them and use them for evil purposes? What alternatives are there to weapons research? I think these questions are extremely hard if not impossible to answer in ways that support weapons research (namely, yes, no, none), and so I think that weapons research is an activity that should be undertaken only in exceptional circumstances.[23]

In this way we can see how to go about addressing questions about matters that are morally suspect, such as weapons research. When the topic is dual use, one of which involves weapons, for instance, the questions become more complicated, but I think we can now see just where the complications arise. The assumption is that R1 is such that we know it can have bad uses. Any research that can be harmful needs to be justified, and so far we have focused on justifications that do not take into account any good use (that is not a product or function of the bad use). A good use, on the present system of values, is one that prevents harm; in the case of R1 this is the disposal of hazardous waste. The role this will play is thus as part of the justification of the project. What the present system of values and justifications requires for dual-use research is that the bad use is 'offset' by considerations that appeal to the good use, and any other relevant considerations, and that these good uses must involve the prevention of harm. Where the 'good' use does not prevent harm, where it is 'neutral' as evaluated on the prevent system, it cannot figure in the justification of the project. In regard to a project like R1, it may therefore even be useful to proceed as follows: justification needs to take the form as if it were the bad use that is the main object of the investigation—that this is foreseen but not intended is irrelevant here—and the good use is to be cited as a reason *not* to abandon the project.

Dual use, responsible research and uncertain outcomes

Now we need to introduce a further complication. Suppose project R2 is such that it might have a bad use, if it were to fall into the wrong hands, if it turned out in a certain way, if certain people used it as a point of departure for other research, and so on—in other words, R2 is risky. Consider a response here resembling that made above about intention—that it is not the researcher's job to look ahead and consider other possible uses of her research. None of these potential outcomes needs to be taken into account; only the project itself and its 'scientific outcomes' need concern the researcher. As we have seen already, this attitude is at odds with the responsibilities of the scientist in this day and age. Science affects us and our environment, whether it is set up as pure

23 So exceptional as to never obtain in fact, though this is not the occasion to try to make that case.

or applied, and all scientists are obliged to try to look ahead and see where their research is going. Everyone has responsibilities when they take on the role of the scientist, and in this way scientists resemble our train guard.[24] But again there are difficulties: how are scientists able to look into the future and see where their research will lead, for are not outcomes unpredictable? This matter was raised above, where I conceded that some research is unpredictable but also maintained that some is not. Moreover, in the present context where science is heavily sponsored and funded, it can only operate where it promises to give useful outcomes. Were this not true then institutions such as defence departments would not spend more money than anyone on scientific research.[25] Science is now far too big and expensive to be left undirected. Of course, when scientists propose research projects and say what they hope to achieve, they maintain that these are desirable outcomes, that their methods are proper and so on. They do not, as a rule, list undesirable outcomes; these will not normally be the focus of attention of researchers, but I do not see any reason they could not be. I conclude that the present account of the responsibilities of scientists has the implication that they are obliged to do their best to look ahead and try to see if their work will lead to any bad outcomes.

These kinds of assessments involve judgments about risk and threat, and are therefore more or less uncertain, as indeed are good outcomes, because all research has some uncertainty associated with it. The present account implies a cautious and conservative approach—and this may not be welcomed by all researchers. Consider the following familiar scenario: a research project is aimed at uncovering the structure of a pathogen with the aim of finding a preventative therapy or cure for the condition that it causes; however, this knowledge could also be used for making the pathogen more virulent and resistant to the very measures that the project was designed to put in place. The present account, as we saw in section one, sees the prevention of harm as something to be encouraged. This is the second tier of the two-tiered account of the forward-looking responsibility of the scientist, and we can all agree that the prevention of harm is a good thing. But the first tier forbids scientists from doing harmful research, and that strongly proscribes any bioweapons research, including bioweapons research that is unintended. Where does this leave the research project? It does not follow that it should not be undertaken, but it does follow that justification must be given. And here an assessment of the costs versus the benefits is needed, as well as control of the results.

---

24 I discuss the case of Frederic Joliot-Curie, who published nuclear data in 1938, against the urging of Szilard. Fortunately, the data were incorrect and it did not seem to interest the Nazi scientists; in Forge, 2008, op. cit. Joliot-Curie's actions were irresponsible even back then.

25 For many years the US Department of Defense has spent more money on scientific research than anyone else.

If there is some risk that the project will be used for bad ends then it seems clear that not only should the results not be published, but also they should be tightly controlled. Michael Selgelid[26] has raised the question of the censorship of research in regard to dual-use questions, and has advocated a moderate and balanced position. The present account will certainly suggest stricter controls (perhaps another unwelcome implication). From the perspective of common morality, agents are free to do whatever they wish, as long as this does not harm others. Thus, scientists are free to do whatever research they like, but the freedom to research, or research itself, is not given special value or status. When there is a risk that research will cause harm then it is firmly proscribed. And when research also has benefits then it should be conducted with effort to minimise harms, and if this entails censorship then that is what is required for the project to go ahead. Indeed, it seems that this is precisely the view that makes the most sense. The response that the whole purpose of publication is to make ideas and results open to all members of the community who can then build on them will not do here. Networks of respected colleagues can be informed, but the risks of open publication may be too great with dual-use research.

Conclusion

I conclude that the account of science and responsibility put forward in *The Responsible Scientist* has relevance for the ethics of dual-use research, and in two sorts of ways. In the first place, the theory of responsibility ties the scientist to the outcomes of his or her work, and does so more tightly than other viewpoints. For instance, an account that incorporated the straight rule and did not see any special sort of forward-looking responsibility for scientists would have quite different implications; however, I think that the topic of dual use shows that we want a wide-ranging account of responsibility, and we want researchers to think carefully about what they do. The second way in which the account is relevant is in the way it sees responsibility as being discharged or cashed out. This is with respect to the no-harming ethic of common morality. I suspect that this will be the more controversial aspect of the present proposals.

26 Selgelid, M. 2007, 'A tale of two studies; ethics, bioterrorism, and the censorship of science', *Hastings Center Report*, vol. 37, no. 3, pp. 35–43.