12. Moral Responsibility, Collective-Action Problems and the Dual-Use Dilemma in Science and Technology¹

Seumas Miller

The dual-use dilemma

The so-called 'dual-use dilemma' arises in the context of research in the sciences as a consequence of one and the same discrete piece, or ongoing program of scientific research, intentionally undertaken for good ends having the potential to be intentionally used to cause great harm.² So there is a primary user who creates new knowledge or designs new technology for good—for example, discovers how to aerosolise chemicals for use in crop-dusting. But there is also a secondary user who uses the knowledge or technology for some harmful purpose—for example, uses the newly discovered process of aerosolisation to weaponise chemicals.

Many, if not most, so-called dual-use dilemmas are not really dilemmas in the narrow sense of being situations involving two options that are equally ethically problematic. In the first place, the dilemmas in question could be tri-lemmas; indeed, there could be four or five or some very large number of options, all of which are equally ethically problematic. In the second place, the options are not generally equally ethically problematic. Certainly, there are ethical considerations for and against each of the options, however, it may well be that, all things considered, one of the options is morally preferable to the others and that this is relatively obvious to any rational, morally sensitive person. The point is rather that there are at least some significant moral costs associated with each of the available options.

Naturally, many, if not most, scientific discoveries and especially new technologies have dual-use potential in the trivial sense that they could be used by someone for some malevolent purpose. Indeed, any newly designed object, such as the first baseball bat, has dual-use potential in this trivial sense. After all, baseball

¹ This chapter is a condensed version of material provided by the author in Work Package 3 (WP3) of the European Union's Seventh Framework project, Synth-Ethics. (WP3 was jointly authored by Seumas Miller and Michael Selgelid.)

² Miller, S. and Selgelid, M. 2008, *Ethical and Philosophical Consideration of the Dual Use Dilemma in the Biological Sciences*, Springer, London, ch. 1.

bats can be used to hit people over the head, as well as for the enjoyment of playing baseball; however, it is implicit in the use of the term 'dual use' in play in the academic literature that the potential harm in question is of a very great magnitude—for example, the potential of nuclear fusion to lead to the creation of the hydrogen bomb, and the potential of genetic engineering to lead to a super virus. Moreover, it is also implicit that the dual-use potential in question is not simply a repeat application of existing science or technology—as is the case with the construction of the first baseball bat—but genuinely new science or technology that can be used to provide a qualitatively or quantitatively new means of harming, for example: the invention of the first explosives, or creating a highly virulent form of an existing much less virulent pathogen. Hitting a person over the head with a solid object has been done since time immemorial, so hitting someone over the head with a baseball bat is hardly a novel means of harming. By contrast, blowing someone up with gunpowder was initially a novel means of harming. Moreover, unlike the use of a baseball bat, the use of an explosive device, such as a 10-tonne bomb, could harm on a very large scale.

Note that accidents involving science and technology, even accidents on a very large scale such as the Chernobyl disaster, are not dual use in my sense since there is no secondary evil user, although they may involve unethical behaviour such as negligence with respect to safety precautions. Nor are weapons designed as weapons—for example, guns—instances of dual-use science and/or technology. For even if their harmful use is intended to be ultimately for the good, such weapons are in the first instance designed to harm; their use to harm is not a secondary, but rather a primary, use.

One paradigmatic case of dual-use research was the biological research on the mousepox virus. The dilemma was as follows.

- Option 1: The scientists ought to conduct research on the mousepox virus and do so intending to develop a genetically engineered sterility treatment that combats periodic plagues of mice in Australia.
- Option 2: The scientists ought *not* to conduct the research since it might lead (and in point of fact did lead) to the creation of a highly virulent strain of mousepox and the consequent possibility of the creation—by, say, a terrorist group contemplating a biological terrorist attack—of a highly virulent strain of smallpox able to overcome available vaccines.

It is a dilemma since there are two options with good reasons in favour of both and it is an ethical dilemma since these options are morally significant. In essence the dilemma involves a choice between intentionally doing good and foreseeably providing others with the means to do evil.

A second and more recent paradigm of dual-use research is the biological research done on a deadly flu virus (A [H5N1]), which causes bird flu. Scientists

in the United States and the Netherlands created a highly transmissible strain of this virus, albeit, as it emerged, a strain that is not as deadly as ordinary H5N1. Crucially, the work was done on ferrets, which are considered a very good model for predicting the likely effects on humans.

As with the mousepox case, there are two options, ethically speaking.

- Option 1: The scientists ought to conduct research on the bird flu virus and do so intending to develop vaccines against similar, but deadly, naturally occurring and artificially created strains of H5N1.
- Option 2: The scientists ought *not* to conduct the research since it will lead to the creation of a virus that is transmissible to humans, and is unlikely to lead to the development of vaccines against similar, but deadly, naturally occurring and artificially created strains of H5N1.

Notice here that the scientific claim in option two that the research is unlikely to lead to the development of vaccines contradicts the corresponding claim in option one. In this respect the ‘dilemma’ is unlike that posed by the mousepox research. Moreover, if this scientific claim made in option two is correct then the justification offered in option one for conducting the research collapses.

In such dual-use cases, the researchers—if they go ahead with the research—will have foreseeably provided the means for the harmful actions of others and, thereby, arguably infringed a moral principle (albeit their infringement might in some cases be morally justified). The principle in question is the principle of what we might refer to as the ‘no means to harm principle’.³ Roughly speaking, this is the principle that rules out providing malevolent persons with the means to do great harm—a principle that itself ultimately derives from the more basic principle: do no harm.

The ‘no means to harm principle’ (NMHP) is the principle that one should not foreseeably (whether intentionally or unintentionally) provide others with the means to intentionally do great harm and it assumes: 1) the means in question is a means to do great harm; and 2) the others in question will do great harm, given the chance.

As with most, if not all, moral principles, the NMHP is not an absolute principle and, therefore, it can be overridden under certain circumstances. For example, it is presumably morally permissible to provide guns to the police in order that they can defend themselves and others. Moreover, as is the case with most, if

³ This principle, or similar ones, is familiar in a variety of ethical contexts. See, for example, Scanlon, T. 1977, ‘A theory of freedom of expression’, in R. M. Dworkin (ed.), *The Philosophy of Law*, Oxford University Press, Oxford.

not all, moral principles, the application of the NMHP is very often a matter of judgment. In the case of NMHP the need for judgments depends in large part on the uncertainty of future harms.

The dual-use dilemma is a dilemma for researchers, governments, the community at large, and for the private and public institutions, including universities and commercial firms that fund or otherwise enable research to be undertaken. Moreover, in an increasingly interdependent set of nation-states—the so-called global community—the dual-use dilemma has become a dilemma for international bodies such as the United Nations. The dilemma is perhaps most acute in those areas of science and technology that operate on an ‘engineering’ or ‘construction’ model—for example, synthetic biology, nanotechnology (as opposed to a ‘description’ model, which restricts itself to the description of pre-existing entities and their causal and other relationships—for example, astronomy).

Scientific research and weapons of mass destruction

The history of science and technology is replete with examples of scientific research being used intentionally or unintentionally to create weapons, including weapons of mass destruction (WMD).

Scientists have developed biological, chemical and nuclear weapons. Such weapons include the following historical examples: the mustard gas used by German and British armies in World War I; the aerial spraying of plague-infested fleas by the Japanese military in World War II that killed thousands of Chinese civilians; the dropping of atomic bombs on Hiroshima and Nagasaki by the US Air Force in World War II; the large-scale biological weapons program in the Soviet Union from 1946 to 1992; the biological weapons program of the apartheid government in South Africa; and the use of chemical agents against Kurds by Saddam Hussein’s Iraqi regime in 1988.

In recent years there have been a number of high-profile ‘defections’ of scientists from Western countries to authoritarian states with WMD programs. For example, Dr Abdul Qadeer Khan joined and in large part established Pakistan’s nuclear weapons program after working for Urenco in the Netherlands, and Frans van Anraat (also from the Netherlands) went to Iraq to assist Saddam Hussein’s WMD program in producing mustard gas.

Moreover, in recent years there have been a number of acts, or attempted acts, of bioterrorism—notably, by the Aum Shinrikyo in Japan (they attempted to

acquire and use anthrax and botulinum toxin, and actually carried out a number of terrorist attacks using the chemical sarin gas), by al-Qaeda (they attempted to acquire and use anthrax) and by US Government employees in the so-called Amerithrax attacks (involving the actual use of anthrax).

Given that a small number of animal, human and plant pathogens are readily obtainable from nature, and that bioterrorists with some microbiological training could use these to inflict casualties or economic damage, and especially given the new possibilities provided by synthetic biology (involving the creation of pathogens *de novo*), evidently there is a non-negligible bioterrorist threat, and it is increasing.⁴ This threat is perhaps greatest in unregulated environments, including in weak or failing states in which well-resourced international terrorist groups are allowed to flourish.

Moreover, it would be naive to assume that the scientific community can be entirely trusted to regulate itself in relation to dual-use problems. After all, thousands of scientists have worked in the abovementioned and other WMD—for example, biological weapons—programs, and in doing so have had as their institutional collective end the production of biological weapons. Accordingly, these scientists are directly collectively morally responsible for the existence of those weapons and, in the case of scientists working for authoritarian governments, for enabling authoritarian regimes to possess them. Moreover, on some occasions, as already noted, WMD have actually been used; accordingly, the scientists involved in the development of these WMD are morally implicated, even if only indirectly, in the harms caused by such use.

Scientific freedom and joint action

According to scientist cum philosopher Michael Polanyi:

The existing practice of scientific life embodies the claim that freedom is an efficient form of organization. The opportunity granted to mature scientists to choose and pursue their own problems is supposed to result in the best utilization of the joint efforts of all scientists in a common task. In other words: if the scientists of the world are viewed as a team setting out to explore the existing openings for discovery, it is assumed that their efforts will be efficiently coordinated if only each is left to follow his own inclinations. It is claimed in fact that there is no other

4 Miller, S. 2009, *Terrorism and Counter-Terrorism: Ethics and Liberal Democracy*, Blackwell, London, ch. 7.

efficient way of organizing the team, and that any attempts to coordinate their efforts by directives of a superior authority would inevitably destroy the effectiveness of their cooperation.⁵

Polanyi's view is that each scientist acts freely but does so

1. on the basis of the work of past scientists
2. with constant reference and adjustment to the work of other contemporary scientists
3. in the overall service of a collective end of comprehensive knowledge (in the sense of understanding) of the scientific phenomena in question.

So his conception is one of individual scientific freedom in the overall context of intellectual interdependence in a joint *epistemic* project.

The epistemic character of scientific research is obvious: science aims at knowledge or, more broadly, at understanding. What is perhaps not quite so obvious, however, is that scientific work is, therefore, epistemic *action*. We often contrast the acquisition of beliefs, knowledge and other cognitive states with action—the notion of action being reserved for behaviour or bodily action. Action is intentional and voluntary; on the other hand, knowledge acquisition, it is often held, is unintentional and involuntary. Consider perceptual knowledge. If an object, say, a piece of litmus paper placed in a liquid solution in a test tube, turns blue before our eyes then we come to believe that it has turned blue; we have no choice in the matter, or so it seems.

This focus, however, on perceptual knowledge of ordinary middle-sized objects provides a distorted picture of the scientific *epistemic* project for reasons that will become clear in what follows. A further point to be made here picks up on the joint or cooperative nature of scientific work described above. If scientific research is epistemic action and, at the same time, it is joint action then scientific research is *joint epistemic action*.

Naturally, this being an epistemic academic project, the participants are governed by epistemic principles (for example, replication of experiments), peer-review processes to ensure quality control and uncensored publication.

The key notion in play here is what I am referring to as joint epistemic action. What is joint action?

Joint action consists of multiple individual actions performed by multiple agents and directed towards a collective end—for example, a team of workers building the Empire State Building; a team of terrorists destroying the Twin Towers,

5 Polanyi, M. 1951, *The Logic of Liberty: Reflections and Rejoinders*, Routledge & Kegan Paul, London, p. 34.

killing thousands; and a team of scientists discovering the cure for cancer.⁶ A collective end is an individual end that each of the participating agents has, but it is an end that no one agent acting alone realises on his or her own. So each agent acts interdependently with the other agents in the service of the same, shared end: the collective end. Again, consider the collective end of a security organisation such as the Federal Bureau of Investigation (FBI) whose members may be jointly working to prevent harm—notably, great harms planned by criminal organisations such as terrorist groups; or consider a team of scientists working feverishly to develop an antidote for some infectious disease that is reaching epidemic proportions.

Joint actions exist on a spectrum. At one end of the spectrum there are joint actions undertaken by a small number of agents performing a one-off simple action at a moment in time—for example, two lab assistants lifting some equipment onto a bench. At the other end of the spectrum there are large numbers of institutionally structured agents undertaking complex and often repetitive tasks over very long stretches of history—for example, those who built the Great Wall of China, or biological scientists developing vaccines.

Joint activity within institutions typically also involves a degree of competition between the very same institutional actors who are cooperating in the joint activity. For example, life scientists within a university who are engaged in joint research and teaching activity in the service of the university's institutional goals are also competing for academic status within the relevant international scientific community, and competing in part on the basis of their individual contribution to those same institutional goals.

Moreover, in many institutional settings organisations compete with one another—for example, business organisations in market settings. Here there is joint activity at a number of levels. For one thing, each competing organisation (for example, a single corporation) comprises a 'team' of individual agents who are cooperating with one another and jointly working to secure the collective ends of the organisation (for example, a biotech company trying to maximise market share). For another thing, each 'team' (for example, each corporation) is engaged in *joint* compliance with the regulatory framework that governs their competitive market behaviour—that is, each complies with (say) the regulations of free and fair competition interdependently with the others doing so and in the service of ensuring the ongoing existence of the market in question. This is consistent with the existence of a regulator who applies sanctions to those organisations which breach the regulations, including safety regulations that might be regarded as a costly and unnecessary impost on business (for example, the 'select agent rule', some mandatory safety procedures in laboratories, and

6 Miller, S. 2001, *Social Action: A Teleological Account*, Cambridge University Press, Cambridge, ch. 1.

screening all orders of DNA sequences); the last compliance mechanism is an 'add-on' to the fundamental underlying structure of interdependence of action in the service of collective ends that is constitutive of market mechanisms.⁷

Let us consider further the notion of a joint action and the correlative notion of a collective end. As stated above, joint actions involve multiple agents with the same end—for example, to build a house (the collective end) or map the human genome. Note the following points.

- First, each agent's individual action is a (possibly small) causal contribution to the collective end—for example, building the Great Wall of China, or mapping the human genome.
- Second, each agent's individual action or omission is performed on condition that others perform their contributory actions/omissions; there is interdependence of action.
- Third, each has the collective end only on condition others have the collective end; there is interdependence of ends.
- Fourth, what the collective end is and that it is being pursued are matters of mutual true belief among the participants (A and B mutually truly believe that p if and only if A believes truly that p, B believes truly that p, A believes that B believes that p, and so on).⁸
- Fifth, collective ends are purely conative states; they are not affective states such as feelings or desires. Accordingly, we need to distinguish the mental states constitutive of joint actions (that is, intentions, ends and beliefs) from the mental states that might motivate some joint actions (for example, feelings and desires).

Joint actions can realise collective ends that are also goods—namely, collective goods. Examples of such collective goods are a law-abiding society and an economically viable biotechnology sector. At an organisational level, a collective good might be the realisation of a collective end that consists in harm minimisation or prevention. Thus a biotechnology firm might have as one of its collective ends to avoid any major industrial accidents or to prevent any serious security breaches. Note that in my sense of the term 'collective good', a collective good is simply a good that is produced by joint action directed to a collective end the realisation of which consists in the provision of that good. Such joint action includes action that consists in joint compliance with safety and security procedures that has as a collective end the collective good of prevention of harm.

7 Miller, S. 2010, *The Moral Foundations of Social Institutions: A Philosophical Study*, Cambridge University Press, Cambridge, pp. 50–2, and ch. 2.

8 Miller, 2001, op. cit., pp. 56–9.

Such collective goods are not necessarily reducible to an aggregate of individual benefits. Relational goods produced by joint activity—such as social harmony and mutual scientific knowledge (each knows that *p* and each knows that each knows that *p*, and so on)—are cases in point.

Some collective goods are intrinsic goods and are jointly pursued for their own sake—that is, they are not pursued merely as a means to an individual end. Various kinds of collective interest, such as the national interest or the interests of the scientific community, are examples of this.

Moreover, a belief in the value of collective goods can motivate action irrespective of individual self-interest—for example, a soldier giving his life in the national interest, or a whistleblower blowing the whistle on a secret biowarfare program.

It will be evident from the above that I am distinguishing between self-interested reasons (or motives) for individual and joint action—for example, so-called ‘sticks’ and ‘carrots’—and moral reasons (or motives) for action (including joint action), and claiming that the latter are not reducible to the former (and vice versa).

It will be further evident that I hold that moral reasons—for example, a belief that biological warfare is wrong, or a belief in the common good—can motivate in and of themselves. So individual self-interest is not the only motive for action. Moreover, even when the motive of self-interest is present, which it obviously typically is, it is not necessarily the dominant motivation.

Armed with this general characterisation of joint action, collective goods and so on, let us now turn to joint epistemic action—the salient form of joint action involved in scientific activity.

Joint epistemic action

As we have seen, epistemic actions are actions of acquiring knowledge. Here we can distinguish between so-called ‘knowledge-that’ and ‘knowledge-how’, the former being propositional knowledge (knowledge of the truth of some proposition), the latter being practical knowledge (knowledge of how to undertake some activity or produce some artefact). The definition of propositional knowledge, in particular, is philosophically controversial, but let us assume for our purposes here that someone, *A*, has knowledge that *p* if and only if *A* has a true belief that *p* and *A* has a justification for believing that *p*, which does not rely on some other false belief.⁹

9 See, for example, Moser, P. 1989, *Knowledge and Evidence*, Cambridge University Press, Cambridge.

The methods of acquiring propositional knowledge are manifold but for scientific knowledge they include observation, calculation and testimony. Moreover, the acquisition of these methods is very often the acquisition of knowledge-how—for example, how to calculate, how to use a microscope, how to ‘read’ an x-ray chart.

In the case of the engineering sciences there is an even more obvious and intimate relationship between propositional and practical knowledge, since both are in the service of constructing or making things. Thus in order to build an aeroplane engineers have to have prior practical (‘how-to’) knowledge and that practical knowledge in part comprises propositional knowledge—for example, with respect to load-bearing capacity. Moreover, this engineering model has increasing applicability in new and emerging sciences such as synthetic biology and nanotechnology. In the case of synthetic biology, for example, scientists can develop new vaccines, enhance the virulence and transmissibility of existing pathogens and even create new pathogens (albeit, presumably using elements of existing pathogens as building blocks).

What counts as sufficient evidence for the possession of knowledge varies from one kind of investigation and one kind of investigative context to another. Thus a scientist seeking a cure for cancer would need his or her experimental results to be replicated by other scientists and the putative cure (say, a new drug) would need to be subjected to clinical trials before being made widely available. A detective investigating a series of murders—for example, by the Yorkshire Ripper—will be focused not only on physical evidence but also on motive (a mental state) and opportunity. Moreover, the evidential threshold for being found guilty is beyond reasonable doubt.

Whereas the acquisition of practical knowledge is readily seen as emanating from action and, indeed, as being a species of action (‘knowledge-in-action’), the acquisition of propositional knowledge is a different matter; however, coming to truly believe that *p* on the basis of evidence—that is, propositional knowledge acquisition—is action in at least three respects.

First, the agent, *A*, makes a decision to investigate some matter with a view to finding out the truth; the action resulting from this decision is epistemic action. For example, a detective intentionally gathers evidence having as an end to know who the serial killer of prostitutes in Yorkshire is—that is, who the Yorkshire Ripper is. Thus the detective gathers physical evidence in relation to the precise cause and time of death of the Ripper’s victims; the detective also interviews people who live in the vicinity of the attacks, and so on. Here *A* has decided that *A* will come to have a true belief with respect to some matter, as opposed to not having any belief with respect to that matter—for example, a true belief with respect to who the Yorkshire Ripper is. *A*’s decision is between

coming to have true belief and being in a state of ignorance, and in conducting the investigation A has decided in favour of the former. Similarly, a scientist seeking to discover the genetic structure of some organism makes a decision to come to have a true belief with respect to this matter rather than remaining in ignorance.

Second, the agent, A, intentionally makes inferences from A's pre-existing network of beliefs; these inferences to new beliefs are epistemic actions. For example, a forensic scientist might infer the time of death of a murder victim on the basis of her prior belief that rigor mortis sets in within 10 hours after death.

Third, in many cases A makes a judgment that p in the sense that when faced with a decision between believing that p and believing that not p, A decides in favour of p; again, A is performing an epistemic action. For example, our detective, A, intentionally makes an evidence-based judgment (mental act) that Sutcliffe is the Yorkshire Ripper (as opposed to that Sutcliffe is not the Yorkshire Ripper) and does so having as an end the truth of the matter. Here, A is deciding between believing that p and believing that not p; but A is still aiming at truth (not falsity). A is not deciding to believe what he thinks is false. Similarly, our forensic scientist makes an evidence-based judgment in relation to the cause of death of the victim having as an end the truth of the matter. Here the scientist is deciding between believing that the cause of death was x and believing that the cause of death was not x (but was, say, y).

As is the case with non-epistemic action, much epistemic action—whether it is propositional or practical epistemic action or, more likely, an integrated mix of both—is joint action, that is, joint epistemic action. Joint epistemic action is knowledge acquisition involving multiple epistemic agents seeking to realise a collective epistemic end. For example, a team of scientists seeking knowledge of the cure for cancer is engaged in joint epistemic action.¹⁰

In cases of joint epistemic action there is mutual true belief among the epistemic agents that each has the same collective epistemic end—for example, to discover the cure for cancer. Moreover, there is typically a division of epistemic labour. Thus in scientific cases some scientists are engaged in devising experiments, others in replicating experiments, and so on. So, as is the case with joint action more generally, joint epistemic action involves interdependence of individual action, albeit interdependence of individual epistemic action.

As we saw above, knowledge of the cure for cancer, for example, is joint epistemic action that involves a collective epistemic end and also involves a division of epistemic labour and interdependence of epistemic action. The further point to be made here is that there is interdependence in relation to such collective

¹⁰ See Miller, 2010, *op. cit.*, ch. 11.

epistemic ends. This is because, given the need for replication of experiments by others, each can only know that *p* is the cure for cancer—to continue with our example—given that others also know this, that is, there is interdependence in relation to the collective end of knowledge.

A collective epistemic end can be both a collective intrinsic good—and thus an end in itself—and the means to further ends. Knowledge of the cure for cancer is a case in point. Such knowledge consists of propositional and practical knowledge: knowledge of the cure for cancer and knowledge of how to produce it. This knowledge, however, has as a further (collective) end: the actual production of the cure (say, a drug). And this end has in turn a still further end—namely, to save lives.

If knowledge of the cure for cancer is a collective end in itself then it is not simply a means to individual ends—namely, each having as an end that he/she knows the cure for cancer. Rather it is mutually believed that knowledge of the cure for cancer is a collective intrinsic good. In my view, however, moral beliefs can have motivational force.¹¹ In that case the mutual belief that knowledge of the cure for cancer is a collective intrinsic good can have motivational force.

It follows from this that—as we saw with joint action more generally—joint epistemic action can be collectively self-motivating and does not necessarily have to rely on prior affective states such as desires.

Layered structures of joint epistemic action

Organisational action typically consists of what elsewhere I have termed a *layered structure of joint actions*.¹² One illustration of the notion of a layered structure of joint actions is an armed force fighting a battle. Suppose at an organisational level a number of joint actions ('actions') are severally necessary and jointly sufficient to achieve some collective end. Consider an army fighting a battle. Here the 'action' of the mortar squad destroying enemy gun emplacements, the 'action' of the flight of military planes providing air cover and the 'action' of the infantry platoon taking and holding the ground might be severally necessary and jointly sufficient to achieve the collective end of defeating the enemy; as such, these 'actions' taken together constitute a joint action. Call each of these 'actions' level-two 'actions', and the joint action that they constitute a level-two joint action. From the perspective of the collective end of defeating the enemy, each of these level-two 'actions' is an individual action that is a component of a (level-two) joint action: the joint action directed to the collective end of defeating the enemy.

¹¹ See, for example, Benn, S. 1988, *A Theory of Freedom*, Cambridge University Press, Cambridge, ch. 2.

¹² Miller, 2001, op. cit., ch. 5.

Each of these level-two ‘actions’ is, however, already in itself a joint action with component individual actions; and these component individual actions are severally necessary and jointly sufficient for the performance of some collective end. Thus the individual members of the mortar squad jointly operate the mortar to realise the collective end of destroying enemy gun emplacements. Each pilot, jointly with the other pilots, strafes enemy soldiers to realise the collective end of providing air cover for their advancing foot soldiers. Further, the set of foot soldiers jointly advances to take and hold the ground vacated by the members of the retreating enemy force.

At level one there are individual actions directed to three distinct collective ends: the collective ends of (respectively) destroying gun emplacements, providing air cover, and taking and holding ground. So at level one there are three joint actions—namely, the members of the mortar squad destroying gun emplacements, the members of the flight of planes providing air cover, and the members of the infantry taking and holding ground. Taken together, however, these three joint actions constitute a single level-two joint action. The collective end of this level-two joint action is to defeat the enemy; and from the perspective of this level-two joint action, and its collective end, these constitutive actions are (level-two) individual actions.

Importantly for our purposes here there are layered structures of joint *epistemic* action. Consider a crime squad, comprising detectives, forensic scientists and so on, attempting to solve a crime.

At level one, a victim, A, communicates the occurrence of the crime (say, an assault) and description of the offender to a police officer, B. But A asserting that p to B is a joint epistemic action; it is a cooperative action governed by conventions—the convention that the speaker, A, tells the truth and the hearer trusts the speaker to tell the truth.¹³

Also at level one, a couple of detectives interviews the suspect to determine motive and opportunity; the detectives are cooperating with one another in the performance of a joint epistemic action the collective end of which is to discover motive and opportunity.

Finally, at level one, a team of forensic scientists analyses the available physical evidence—for example, the DNA of the blood samples of the offender found on the victim are matched to the suspect’s DNA; the forensic scientists are engaged in joint epistemic action to determine whether there is or is not a DNA match.

These three level-one joint epistemic actions are constitutive of a level-two joint epistemic action—namely, the level-two joint epistemic action directed towards

13 Miller, 2010, op. cit., ch. 11.

the collective end of determining who committed the crime. Accordingly, when each of the level-one joint epistemic actions is successfully performed then the level-two joint epistemic action is successfully performed—that is, the crime squad solves a crime.

Now consider an example of a large scientific project conducted by a number of cooperating organisations and hundreds of scientists over many years—namely, the human genome project. The project involved multiple connected goals—collective ends—and multiple layered structures of joint action, including joint projects in publishing, undertaken to realise those goals.

In fact most organisations are hierarchical institutions comprising task-defined roles standing in authority relations to one another, and governed by a complex network of conventions, social norms, regulations and laws. Consider a science department in a university or the forensic laboratory in a police organisation: both comprise heads of department, scientists, laboratory assistants and so on, and the work of both is governed by scientific norms of observation, replication of experiments, and so on. So most layered structures of joint action, including joint epistemic action, are undertaken in institutional settings, and scientific joint epistemic action is not an exception.

Institutions have *de facto* purposes/strategic directions—that is, collective ends, such as to maximise shareholder profit (corporations), find a cure for cancer (university research team), or build an atomic bomb (military organisation). Institutions also have specific structures (hierarchical, collegial and so on) and they have specific cultures (for example, a competitive, status-driven ethos).¹⁴ In this connection consider scientific activity—for example, biological research, undertaken in three different institutional settings: that of the university, the commercial firm and the military biodefence organisation. Some of the principal purposes/strategic directions (collective ends) of commercial firms—for example, to maximise shareholder profits—are quite different from, and possibly inconsistent with, those of universities (for example, scientific knowledge for its own sake), and quite different again from those of military research establishments (for example, to save the lives of *our* military personnel). Again, the hierarchical structures within a military research establishment are quite different from the more collegial structures prevailing in universities; and the structure of commercial firms is quite different again. The general point to be made here is that scientific activity is not only a form of complex joint activity (a layered structure of joint epistemic action); it is activity that is inevitably shaped by the non-scientific institutional setting in which it is conducted—that is, by the specific collective ends, structure and cultures of particular institutions.

14 Miller, 2010, *op.cit.*, 'Introduction'.

Here we also need to stress the distinction between the *de facto* institutional collective end, structure and/or culture and what it *ought* to be; cultures, for example, can vary greatly from one organisational setting to another, notwithstanding that the type of institution in question is the same or very similar.

In the light of the above, we can distinguish the normative account of science as a joint intellectual activity—for example, aimed at knowledge for its own sake—from science as a means to broader social ends (for example, vaccines to save lives). Moreover, we can distinguish both from the normative account of specific institutions in which science exists principally as a means—for example, commercial firms (vaccines to make profit) and a military biodefence organisation (vaccines to save the lives of our military personnel). Importantly, in the context of discussion of dual-use concerns, we can distinguish within the normative account of science (both at the level of joint intellectual activity and at the level of specific institutions) between its positive ends (for example, knowledge for its own sake and knowledge as a means to combat disease) and negative ends (for example, harm prevention). Earlier we discussed the no means to harm principle (NMHP)—namely, the principle that one should not foreseeably (whether intentionally or unintentionally) provide others (for example, bioterrorists) with the means to do great harm. Clearly the means in question is essentially scientific or technological knowledge and this knowledge is the product of joint epistemic action (indeed, joint epistemic action undertaken in the context of multiple layered structures of joint epistemic action). Accordingly, harm prevention in relation to dual-use concerns is, or at least morally ought to be, a joint enterprise with respect to joint epistemic action; it is something that biological scientists as members of their scientific community (or communities) and specific institutions (for example, biology departments in universities, biotech companies) morally ought to jointly address including (presumably) by way of education and regulation of their potentially harmful joint epistemic action. But let us be clear on the moral responsibilities in play here.

Collective moral responsibility

As we have seen, dual-use issues in science are inherently ethical, and scientific activity is essentially joint epistemic activity undertaken in various institutions with possibly divergent normative underpinnings. Given the multilevel, cooperative character of scientific activity and given its ethical or moral significance, a key theoretical normative notion in play here is that of collective moral responsibility.

Let us distinguish between natural, institutional and moral responsibility and, in respect of responsibility, between individual and collective responsibility. I note that the notions of natural, institutional and moral responsibility are not mutually exclusive. I also note that agents responsible for some outcome might have both individual and collective responsibility for that outcome.

An agent, A, has *natural* responsibility for some action, x, if A intentionally did x for a reason and x was under A's control. Bench scientists engaging in routine scientific research—for example, replication of experiments—have natural responsibility for their actions. Moreover, such actions might not have any obvious moral implications.

Agent A has *institutional* responsibility for action x if A has an institutional role that has as one of its tasks x. Thus, for example, laboratory assistant, A, has the institutional responsibility to clean the test tubes; moreover, A has this responsibility even if A does not in fact do this.

What of moral responsibility? Roughly speaking, agents have moral responsibility for natural or institutional actions if those actions have moral significance. So if A is naturally or institutionally responsible for x (or for some foreseeable outcome of x, O) and x (or O) is morally significant then, other things being equal, A is morally responsible for x (or O) and, other things being equal, can be praised/blamed for x (or O).

Note that other things might not be equal if, for example, A is a psychopath (and, therefore, incapable of acting in a morally responsible fashion) or if A does something wrong but has a good excuse (and, therefore, ought not to be blamed).

Note also that if O involves some intervening agent, B, who directly causes O, then A may have diminished moral responsibility for O.

Let us now consider collective moral responsibility. In essence, the account of collective moral responsibility mirrors that of individual moral responsibility, the key difference being that the actions in question are joint actions, including joint epistemic actions.

Accordingly, if agents A, B, C and so on are naturally or institutionally responsible for a joint (including epistemic) activity, x (and/or some foreseeable outcome of x, O), and x (and/or O) is morally significant then, other things being equal, A, B, C and so on are collectively (that is, jointly) morally responsible for x (and/or O) and, other things being equal, can be praised or blamed for x (and/or O).

The 'other things being equal' clauses function here as they did in the above account of individual moral responsibility. Moreover, as was seen to be the case with individual moral responsibility, if there are additional intervening

(individual or joint) actions then those jointly responsible for the joint action in question, and its outcome, may have diminished moral responsibility. Scientists who engage in dual-use research that is subsequently used in the construction of WMD may well have diminished responsibility for the harm caused by those WMD. Diminished responsibility is not, however, necessarily equivalent to no responsibility. Further points to be made here are as follows.

First, each agent may have full or partial moral responsibility for x jointly with others for the joint action x and/or its outcome. If, for example, five men each stab a sixth man once, killing him, each is held fully morally (and legally) responsible for the death even though no single act of stabbing was either necessary or sufficient for the death. In some cases each agent might have full moral responsibility (jointly with others) for some outcome, O —notwithstanding the fact that each only made a very small causal contribution to the outcome—in large part because each is held to have prior full institutional (including legal) responsibility (jointly with others) for O .

On the other hand, each agent might have partial and minimum moral responsibility jointly with others if each only makes a very small and incremental contribution as a member of a very large set of agents performing their actions over a long period—for example, the scientists who worked on the human genome project.

Second, we need to distinguish cases in which agents have collective moral responsibility for some joint action or its outcome from cases in which agents only have collective moral responsibility for failing to take adequate preventative measures against O taking place. Many untoward dual-use cases are of the latter kind.

On the other hand, agents may not have any collective moral responsibility with respect to some foreseeable morally significant outcome, O , if O has a low probability, takes place in the distant future and involves a large number of intervening agents.

The collective moral responsibilities of scientists are multiple. Scientists have a collective institutional (professional) and moral responsibility as scientists to acquire knowledge for its own sake.

Scientists functioning in universities also have a collective institutional and moral responsibility to acquire knowledge for the good of humanity—for example, vaccines for poverty-related diseases.

Scientists functioning in commercial firms might have a collective institutional and (contractually based) moral responsibility to acquire (say) knowledge of vaccines for rich people's diseases—since that is a commercial imperative of their employer and they are being paid to do just that.

Scientists functioning in biodefence organisations have a collective institutional (and moral?) responsibility to acquire knowledge of vaccine-resistant pathogens if this is a national security imperative of their employer—that is, the government.

As human beings, scientists have a collective moral responsibility not to provide the means for others to intentionally do great harm—for example, the means to allow others to drop atomic bombs on Hiroshima and Nagasaki or engage in biowarfare.

Moreover, these various collective institutional and moral responsibilities may be inconsistent with one another—notably, the collective moral responsibilities scientists have as human beings and the institutional responsibilities that they might have as members of military research organisations.

Regulatory frameworks

In the light of the dangers stemming from the dual uses of science and technology there is a need to consider a range of regulatory measures to reduce the risks. Such measures include the imposition of limits on dual-use experiments and on the dissemination of potentially dangerous information resulting from dual-use discoveries. These measures themselves exist on a spectrum ranging from the least intrusive/restrictive to the most intrusive/restrictive.¹⁵

Some specific regulatory measures that might be considered include the following ones.

1. Mandatory physical safety and security regulation: Should there be regulations providing for physical safety and security in relation to the storage of, the transport of and physical access to samples of pathogens, equipment, laboratories and so on? And should compliance with such regulations be mandatory? In both cases the answer is presumably in the affirmative.
2. Licensing of dual-use technologies/techniques: Should there be mandatory licensing of dual-use technologies/techniques/pathogen samples? Only certain laboratories in the public sector and the private sector might be licensed to

¹⁵ For a detailed treatment of these and other options, see Miller and Selgelid, 2008, *op. cit.*

engage in research involving the use of certain dual-use technologies and licences for DNA synthesisers might be required.

3. Mandatory education and training: Given the potential harms arising from, for example, 'experiments of concern', it is clear that some process of education and/or training for relevant researchers and other personnel is called for.
4. Mandatory personnel security regulation: Doubtless it is prudent, indeed it is a moral requirement, that access to virulent pathogens be disallowed to a researcher diagnosed as a psychopath or to a known member of a terrorist organisation.
5. Censorship/constraint of dissemination: The question of whether research findings ought to be freely disseminated, censored or their dissemination in some lesser way restricted is an extremely difficult issue and it is by no means obvious who the ultimate decision-maker ought to be. A relevant important distinction here is that made above between first-tier and second-tier dual-use research. For example, first-tier research findings might need to be disseminated in such a way that anyone being informed of these findings would not be able to replicate the experiments that enabled the results reported in the findings. Indeed, this is precisely what was recommended at one point (for example, by the US National Science Advisory Board for Biosecurity) in relation to the dual-use bird flu research mentioned at the outset of this chapter.

Let us assume that a range of regulatory options ought to be pursued, at the institutional (university, commercial firm, government research laboratory), national and international levels (for example, Biological Weapons Convention verification processes). There remain some in-principle obstacles to the establishment of adequate measures to deal with the dual-use problem in science and technology and a number of these stem from various perverse incentive structures that derive from collective-action problems. In the final section of this chapter, I will briefly discuss some of the more salient of these.

Collective-action problems

Thus far we have characterised the scientific enterprise as essentially a joint epistemic one: the emphasis has been on intellectual cooperation to achieve common scientific (epistemic) goals in an institutional context of minimal restrictions on the freedom of individual scientists. This picture, however, while acceptable as far as it goes, is an oversimplification. Specifically, it obscures the competitive dimension of scientific activity and, in particular, it masks various

collective-action problems arising from such competition. This is important for our purposes here, not least because it casts the dual-use problem in a somewhat different light.

On the purist (as we might call it) model of scientific activity as joint epistemic action performed under conditions of scientific freedom, the dual-use problem arises only because scientific research undertaken for the benefit of humankind can be misused by others for harmful purposes. Accordingly, there is a need to monitor dual-use research and erect safeguards against misuse by malevolent individuals and groups—for example, ‘lone wolf’ malcontents, nihilistic terrorist groups, ‘rogue states’, and so on.

Notice, first, that there is here an implicit additional assumption—namely, that scientific activity will be undertaken in the first place in order to benefit humankind. This is, as we have seen in relation to WMD programs, not necessarily the case. On the other hand, WMD research is not dual use in our sense (unless military research undertaken for protective purposes yet having the potential, as is probable, to be misused for aggressive purposes is to be regarded as dual use).

Notice, second, that much scientific work, including in the biological sciences, is not undertaken under conditions of scientific freedom. Consider research undertaken in the private sector or for various government laboratories in, for example, authoritarian states such as China. Which research is undertaken, and whether or not it is published, is not necessarily or even typically a decision made by individual scientists or, indeed, by groups of scientists. Rather these are commercial decisions made by managers or they are decisions made by government officials in the national interest (presumably). Accordingly, it is simply not true that scientific work, including scientific work in the biological sciences, is necessarily, or even typically, conducted under conditions of scientific freedom.

But to return to the main point at issue—namely, competition and, relatedly, collective-action problems: in the biological sciences, as elsewhere, there is competition between individual scientists, between scientific institutions and between nation-states.

It is self-evident that there is competition between, for example, biotechnology companies in the private sector. Moreover, governments compete insofar as they have an interest in promoting their own biotech industries and, more generally, insofar as they want to ensure that they do not fall behind in research and development (R&D) in the biological sciences (not least for military reasons). Further, even in the case of scientific work undertaken under conditions of

scientific freedom (for example, in universities), there are important elements of competition—for example, between rival teams of scientists in competition for status and (relatedly) for scarce funding.

As suggested above, competition in these various sectors gives rise to a variety of collective-action problems that have important implications for the dual-use issue.

First, in the private sector there are collective-action problems arising from commercial competition. As already noted, many scientists work in commercial firms in which there is an imperative to maximise profit. In such a context of fierce commercial competition restrictions on dual-use research may handicap an organisation. This is a collective-action problem insofar as an organisation—all things considered and given appropriate weight—ought to choose *not* to perform a particular dual-use experiment on the grounds that the potential harm to humankind resulting from this kind of experiment might outweigh the potential benefits to humankind. All things might not be considered, however, or given appropriate weight. Specifically, the firm might give excessive weight to its commercial interests, especially if it believes that some other competing firm is likely to be less scrupulous and go ahead with the experiments in question. In short, in dual-use cases where discretionary judgment is called for, the judgment might be skewed by considerations of commercial self-interest in a fiercely competitive commercial environment. I note that commercial self-interest may well be dominant in such cases, notwithstanding the commitment of individual scientists to the NMHP. For one thing, it is not necessarily the decision of the scientists, who are, after all, mere employees; and for another, their self-interest as employees might align them with the firm's commercial interest, especially given the relative ignorance of scientists of security issues.

Second, in the university sector there are collective-action problems arising from competition for status. As already noted, many scientists working in the university sector are engaged in a competition for status (and for scarce funds to undertake projects by means of which they can achieve status), both for themselves as individuals and on behalf of the institutions they work for. Accordingly, there is an analogue of the above-described collective-action problem in the private sector. In dual-use cases where discretionary judgment is called for, the judgment might be skewed by considerations of individual or institutional self-interest in a competitive environment, albeit the competition in question is primarily for status (and scarce funds to achieve status).

Third, in the government sector there are collective-action problems arising from competition among nation-states. As noted, in the past and, indeed, in the present there has been a variety of arms races—for example, the nuclear arms race—in which scientists played a central role. The problem here is that

national self-interest is pitted against humanity's collective interest in a context in which there is no enforceable international law; evidently nation-states cannot effectively collectively self-regulate—hence the WMD programs of the United States, Iran, North Korea and so on. The inability or unwillingness to collectively self-regulate exists, however, at least potentially, in relation to the dual-use problem and does so independently of any inherent desire on the part of nation-states to maintain WMD programs. We saw above that the self-interest of individual scientists and the institutions in which they work (for example, commercial firms and universities) can under conditions of fierce competition lead to collective-action problems in relation to dual-use research. Nation-states are, however, themselves in competition with one another, and it is in the economic, military and so on interests of nation-states to support their own R&D in science and technology—that is, to support the work of their commercial firms and universities, and to do so in the face of 'foreign' competition. Accordingly, we cannot necessarily look to individual governments to regulate adequately the scientific research in their own institutions, at least if 'adequately' in this context refers to an all-things-considered, morally informed decision made in the long-term interests of humankind—as opposed to a decision made in the (possibly short-term) national interest.

A fourth collective-action problem arising from competition is of a somewhat different kind: it is a species of the generic problem of freeriding. It is the possibility of the untoward consequences of scientific freeriding, so to speak. Let us assume that Polanyi's scientific freedom model—for example, no censorship—is in fact the best model to acquire new knowledge; those operating entirely outside the model cannot compete. Accordingly, so the argument runs, the 'good guys' (for example, the scientists making vaccines within the framework of scientific freedom) stay ahead of the 'bad guys' (for example, the scientists weaponising pathogens outside the framework of scientific freedom); the 'bad guys' are always playing catch-up. Contrary to this argument, however, it might be claimed that a well-qualified national cohort of 'bad guys' can always free ride but then get ahead of 'good guys'—for example, scientists in an authoritarian state with biodefence projects benefit from the work of those in the scientific freedom model but don't share their own work.