Biodefence and the production of knowledge: rethinking the problem

Allen Buchanan,¹ Maureen C Kelley²

¹Duke University, Durham, North Carolina, USA

²University of Washington, School of Medicine, Seattle, Washington, USA

Correspondence to

Dr Allen Buchanan, Duke University, Durham, NC, USA; allenb@duke.edu

Accepted 17 November 2011

Published Online First

22 December 2011

ABSTRACT

Biodefence, broadly understood as efforts to prevent or mitigate the damage of a bioterrorist attack, raises a number of ethical issues, from the allocation of scarce biomedical research and public health funds, to the use of coercion in quarantine and other containment measures in the event of an outbreak. In response to the US bioterrorist attacks following September 11, significant US policy decisions were made to spur scientific enquiry in the name of biodefence. These decisions led to a number of critical institutional changes within the US federal government agencies governing scientific research. Subsequent science policy discussions have focused largely on 'the dual use problem': how to preserve the openness of scientific research while preventing research undertaken for the prevention or mitigation of biological threats from third parties. We join others in shifting the ethical debate over biodefence away from a simple framing of the problem as one of dual use, by demonstrating how a dual use framing distorts the debate about bioterrorism and truncates discussion of the moral issues. We offer an alternative framing rooted in social epistemology and institutional design theory, arguing that the ethical and policy debates regarding 'dual use' biomedical research ought to be reframed as a larger optimisation problem across a plurality of values including, among others: (1) the production of scientific knowledge; (2) the protection of human and animal subjects; (3) the promotion and protection of public health (national and global); (4) freedom of scientific enquiry; and (5) the constraint of government power.

Biodefence, broadly understood as efforts to prevent or mitigate the damage of a bioterrorist attack, raises a number of ethical issues, from the allocation of scarce biomedical research and public health funds, to the use of coercion in quarantine and other containment measures in the event of an outbreak, to efforts to extend international arms control regimes to biological weapons. In response to the US bioterrorist attacks following 9–11, significant US policy decisions were made to spur scientific enquiry in the name of biodefence. These decisions in turn led to a number of critical institutional changes within the US federal government agencies governing scientific research, both at government laboratories and academic research centres. Subsequent science policy discussions have focused largely on 'the dual use problem': how to preserve the openness of scientific research while preventing research undertaken for the prevention or mitigation of biological threats from being used to cause harm by non-state terrorists or aggressive dictators. On this characterisation of 'the dual use

problem', biomedical scientists must consider whether and, if so, to what extent the commitment to 'open science' ought to be compromised.

Although the term 'open science' is unfortunately broad, the main idea, as Robert Merton and others have noted, is that the scientific enterprise is characterised by a commitment to costless or low cost information sharing, understood as an element of the more basic commitment to the accumulation of knowledge through collective effort.^{1–2} The chief justification of openness is that it contributes to the production of scientific knowledge. Our aim is to join others in the bioethics literature in shifting the ethical debate over biodefence away from a simple framing of the problem as one of dual use, by making clear how a dual use framing distorts the debate about bioterrorism and truncates discussion of the moral issues.^{3–5} To advance the debate further we offer an alternative framing rooted in social epistemology and institutional design theory, better to inform policy deliberation over the full range of ethical challenges raised by the biodefence enterprise.

REFRAMING THE DUAL USE ISSUE

Framing the ethical concerns of biodefence as predominantly a problem of dual use is inadequate for at least two reasons. First, the reference to 'the dual use problem' is misleading. As others have noted there are at least two distinct dual use problems.⁶ Furthermore, measures to cope with one may be inadequate for coping with—or may even exacerbate—the other. Biodefence research might be used not only by non-state terrorists or aggressive dictators, but also by any state that has or contemplates developing an offensive bioweapons programme.

It is important to understand that even states that have no aggressive intentions may have an incentive to develop offensive bioweapons. Fear of not having offensive bioweapons when others have them can motivate a self-defensive offensive bioweapons arms race, as existed between the USA and the former Soviet Union during the cold war.^{7–8} States not intent on aggression may conclude that, as with nuclear weapons, a 'balance of terror' is necessary for their security. Scientists and ordinary citizens should thus be concerned not only that biodefence research may be used to develop offensive bioweapons by non-state terrorists or by 'outlaw states,' but also by their own governments. Furthermore, it is not enough that a country refrains from seeking to use biodefence research to develop offensive weapons. Unless other countries have adequate assurance that this is so, a self-defensive bioweapons arms race may occur. Clarity and candor would be better served if

Editor's choice
Scan to access more
free content

- ▶ <http://dx.doi.org/10.1136/medethics-2012-100526>
- ▶ <http://dx.doi.org/10.1136/medethics-2012-100545>
- ▶ <http://dx.doi.org/10.1136/medethics-2012-100923>

To cite: Buchanan A, Kelley MC. *J Med Ethics* 2013;**39**:195–204.

ambiguous talk about 'the dual use problem' were abandoned and replaced with 'the dual use problems' or by explicit references to 'dual use problem 1' and 'dual use problem 2:'

DU1: Research undertaken for prevention or mitigation of biological threats being used to cause harm by non-state terrorists or aggressive state actors.

DU2: Research used to develop offensive bioweapons by one's own government.

Second, it is not the case that measures to cope with the dual use problem(s) would be the first instance in which biomedical scientists are faced with the problem of a conflict between the values that underlie the norms of 'open science' and other important values. The norms of openness have never been absolute, nor should they be, because the values that underlie them are not absolute but instead must be balanced against other important values. Two examples should suffice to make this simple but crucial point: intellectual property and privacy protections for human research subjects. What sorts of items should count as intellectual property and how extensive the rights to control their uses should be are complex matters on which there is much disagreement; but if there is any room at all for intellectual property in the scientific research enterprise, then the norms of 'open science' cannot be absolute, because intellectual property rules constrain the dissemination of knowledge by limiting access to items (such as gene sequences) whose use is necessary for gaining knowledge. Similarly, ethical concerns about privacy quite properly limit the freedom of researchers to exchange information about human subjects. So, openness is not and has never been an absolute value. The current processes by which scientific knowledge is produced already reflect a compromise between openness and other values.

Recognising these two deficiencies in the dual use framing of biodefence has two important implications. First, one should not assume that policy measures crafted to cope with dual use problem 1 will be effective for coping with dual use problem 2. For example, omitting certain steps in the creation of a deadly virus from a publication might render the publication useless to a non-state terrorist group or to the relatively poorly trained or under-resourced bioweapons researchers of a so-called 'outlaw state', but the better trained, better resourced bioweapons researchers of a 'great power' might be able to fill in the gaps. What is more, some measures to mitigate the risks of dual use 1 might actually increase the risks of dual use 2. For example, a government-appointed national advisory board charged with vetting research to prevent it from being used by non-state terrorists or 'outlaw states' might officially or unofficially channel information to its own government's bioweapons researchers while increasing the value of the information to them by preventing others from getting access to it. Second, and more fundamentally, once we understand that the norms of 'open science' and the values that underlie them are not absolute, it becomes evident that the dual use problems should be reconceived as one aspect of a larger optimisation problem: how can policy, broadly understood, help shape the scientific enterprise in such a way as to give due weight both to its distinctive role in producing knowledge and to other relevant values, including, but not restricted to, the reduction of both dual use risks?

Just what values ought to be included in the optimisation project and how they ought to be weighted, are of course, difficult, contested questions. The central point is that an overly simplistic assumption that the problem is how to balance the two competing values of biosecurity and open science diverts public discussion from the other important values at stake. In the dual

use policy discussions to date, we have seen two examples of this error: (1) failure to consider adequately the impact of biodefence research on the ethical use of human and non-human animals in research; and (2) failure to account for the opportunity costs of biodefence research vis-à-vis efforts to reduce the burden of infectious disease among the world's poor.

Few would dispute that the protection of human and non-human animal subjects also ought to be taken into account in the design of the enterprise of producing scientific knowledge. Yet, when 'the dual use problem' (meaning dual use problem 1) occupies centre stage, it is the interests of only two parties that are likely to be strongly represented: scientists who fear constraints on the pursuit of knowledge, and government officials whose worst nightmare is a bioterrorist attack that could have been prevented. Therefore, one of the dangers of an overly simplistic framing of the ethics of biodefence is that it largely ignores or arbitrarily discounts values that have been central to the research ethics debate since its inception: the protection of research subjects, both human and non-human. Special attention ought to be given to the need for protecting research subjects against risk in the testing or use of experimental vaccines in the event of an outbreak, or in the process of 'emergency preparedness'. In this regard, the ethics of research ought to be nearer the centre of the biosecurity debate.

Similarly, as May has argued, it is important to consider the opportunity costs of investments in biodefence research.⁴ In particular, it can be argued that concerns about distributive justice ought to be given some weight in policies affecting the production of scientific knowledge, for example, by devising policies to provide greater incentives for research that is likely to yield results (such as a vaccine for malaria) that will help meet the special needs of the world's worst-off people.^{9 10} In biodefence discussions, if 'the dual use problem' is treated as central, consideration of this value, if it occurs at all, tends to be almost an afterthought.

To counter this tendency, some have appealed to yet a third sense of the term 'dual use', what might be called the 'dual use opportunity': the prospect that research undertaken for biodefence may contribute, or might be made to contribute, to the alleviation of the burden of disease among the world's worst-off people. This possibility was discussed, for example, at the Bioethics and Biodefence Meeting, 5 February 2007, at the Johns Hopkins School of Advanced International Studies. This meeting was sponsored by the Southeast Regional Center for Excellence for Emerging Infections and Biodefence and co-sponsored by the Johns Hopkins University Berman Bioethics Institute, the University of Minnesota Center for Bioethics, and the University of Washington Department of Medical History and Ethics. The idea is that knowledge for responding to bioterrorist attacks may also be valuable for responding to naturally occurring infectious disease outbreaks, many of which disproportionately affect poor populations, and that biodefence policy should take this fact into account.

Unfortunately, concerns about distributive justice have not been incorporated into the biodefence debate in any serious or systematic fashion. For example, in recent debates concerning the US government investments in global health, including HIV and the President's Emergency Plan for AIDS Relief, at no time were the trade-offs vis-à-vis renewed investments in biodefence research funding mentioned.¹¹ For example, monies that are allocated to anthrax studies are not available for developing new antimalarial drugs.^{12 13} Keeping the biodefence allocation decisions out of transparent debate has masked the opportunity costs of the massive biodefence effort. It is critical to ask,

however, what research or health investments might we forego in order to continue funding biodefence research?

The point is not that concerns of distributive justice or the protection of research subjects ‘trump’ security concerns, nor is it to deny that under exceptional circumstances they should be accorded less weight than they ordinarily have. Instead, it is that there should be a vigorous debate about the ethical justification for reducing the threshold for acceptable risk in the process of consent for experimental vaccines, and for increasing the use of non-human primates in biodefence research. Such a debate requires discussion of multiple values, each of which has substantial weight. An ethically responsible policy approach cannot simply assume that in effect the only two values at stake are ‘open science’ and biosecurity, because efforts to reconcile these two values may have serious consequences for the pursuit of other important values. To summarise, it is not simply that there are two dual use problems, not one (as well as a ‘dual use opportunity’); the more fundamental conclusion is that the dual use problems (and ‘the dual use opportunity’) are only aspects of a larger optimisation problem.

The idea of optimisation is crucial because it emphasises that the task is not to maximise the realisation of any one value (such as protection against bioterrorism), or to achieve an acceptable trade-off between just two values (such as ‘open science’ and biosecurity), but rather to achieve an overall outcome that gives due weight to all relevant values. The optimisation framing opens the door to discussions of values, such as giving some priority to a more equitable distribution of the benefits of scientific research or the protection of research subjects, that otherwise might be ignored or indefensibly discounted as a result of focusing exclusively on the trade-off between ‘open science’ and biosecurity. Notice that we use the term ‘optimisation’ here in a broad sense; there is no assumption that all competing values can be fully quantified and subjected to a definitive maximising calculation. Rather, the point is that there are multiple values that must each be given due consideration in an attempt to make an all-things-considered judgement about what to do. In many cases, optimising will require judgement, not just calculation.

The optimisation framing is also useful for dispelling the view, promoted by the political rhetoric of the ‘war against terror’ (as in all putative national emergencies), that the goal is to maximise risk reduction, that is, to reduce the risk of harm (in this case harm due to the rapid spread of infectious disease) to zero. Maximal, as opposed to optimal risk reduction is irrational and the attempt to achieve it is unethical because efforts to achieve it come at the expense of other important values.

It might be objected that in times of national emergency, such as the so-called ‘war on terror’, the goal is not to optimise across a plurality of values, but to seek a proper balance of only two dominant values: biosecurity and ‘open science’. The idea here would be that in current conditions other values can and ought to be ignored, because the stakes are so high. The unargued and highly problematical assumptions behind this objection are: (1) that in circumstances of extraordinary risk of bioterrorism, biosecurity and ‘open science’ are values of much greater value than all other relevant values combined; (2) that the only way to secure those two values is to proceed as if no other values existed; and (3) that the circumstances of extraordinary risk—risk sufficient to justify such an abandonment of the optimisation approach—can be reliably ascertained. Not one of these three assumptions has been explicitly defended by those who place ‘the

(first) dual use problem’ at centre stage of the debate on biodefence.

It may be difficult to ascertain when conditions justify abandoning the optimisation approach and disregarding values we otherwise agree are of great importance. This point warrants elaboration. Institutions, preeminently, government institutions, shape beliefs about what constitutes an emergency and about when a state of emergency exists. Institutional agents sometimes have strong incentives to encourage a blurring of the line between preparing for an emergency and the occurrence of an emergency. Political leaders, whose roles give them opportunities for shaping public perceptions, have incentives to foster the belief that an emergency exists, because it is generally assumed that emergencies require extraordinary powers and reduce the requirements of transparency as a condition for the legitimacy of political authority. In brief, once people become convinced that we are in an emergency, they are more willing to accept the view that ordinary moral norms and the standard checks and balances of democratic constitutional government do not apply, or apply with less force—that the government should be given a ‘free hand’, and that criticism of the government is inappropriate, dangerous and even disloyal.¹⁴

So whether we are in fact in an emergency is a matter of great importance. Presumably scientific knowledge should play some role in determining the magnitude and probability of the risks that are judged to constitute an emergency and therefore in determining whether a state of emergency exists.

Although good facts are relevant to determining whether an emergency exists, a ‘state of emergency’ is not a natural fact to be discovered by empirical methods. The statement that a state of emergency exists is a political act, grounded in an evaluation of how serious certain risks are, with the added implication that the ordinary moral, political and legal rules do not apply. If this is the case, then a thorough investigation of alternative institutional arrangements for achieving biodefence at acceptable costs—when all relevant moral costs are considered—cannot take the distinction between emergency and non-emergency situations for granted, but must consider the possibility that scientific institutions can play an important role in providing a check on the tendency of government leaders to be too ready to declare an emergency. Furthermore, there is a tendency, as we have seen in the USA since the 9–11 attacks, for institutions implemented in a state of emergency to become permanent; arguably we have remained in a chronic state of emergency, or heightened alert, for a decade. So, once again we come to the same conclusion: it is a mistake to think that the only values to be balanced are biosecurity and ‘open science’. Reduction of the risk of erroneous judgements about the state of emergency, and more generally the risk of abuse of government power, should also be taken into account.

There has been another unclarity in the policy discussions over biodefence policy, particularly concerning the dissemination of scientific findings. Sometimes the assumption is that the solution is to formulate guidelines to help individuals engage in risk–benefit assessments regarding the dissemination of particular research results, when the assumption is that the risk is that of ‘dual use’ (ie, dual use 1) and the benefit is ‘open science’. Those who advocate such risk–benefit assessment also propose that a number of different parties, who in fact occupy quite different roles, including the scientific researchers themselves, scientific journal editors and perhaps government officials as well, should follow the same risk–benefit assessment guidelines and apply them to the same thing, namely, the dissemination of particular research results.^{15–19}

Such proposals overlook the importance of the division of labour in a reasonable response to the optimisation problem. Better outcomes might be achieved if different agents, depending on their institutional roles, engage in different activities, following different guidelines. For example, it could be argued that government officials should not engage directly in the risk–benefit assessment of particular research results, but instead should be responsible for ensuring the accountability of the risk–benefit assessment procedures of other agents, including editors of scientific journals. According to this way of thinking, government officials might well employ some form of risk–benefit analysis, but they would apply it to the evaluation of risk–benefit assessments of particular research results by other agents, not to the act of disseminating or withholding particular research results. Similarly, it could be argued that scientists could assess the risks and benefits of disseminating their research more accurately if they did not attempt actual risk–benefit assessments of it, but instead employed guidelines that include reliable proxies for risk–benefit calculations. The idea that the best way of achieving a favourable balance of benefits over costs is not always to act on the maxim ‘maximise benefits over costs’ is familiar from discussions of indirect utilitarianism.²⁰

While it is correct to say that a proper response to ‘the dual use problem(s)’ will include a role for risk–benefit analysis, determining which agents should apply such analysis to which actions is a complex matter. More precisely, it is a problem of institutional design.

The role of institutions

Institutional solutions to the problem of balancing ‘open science’ with protection against ‘the dual use problem’ have been proposed, but they have typically been defective in two ways. First, they have been based on uncritical assumptions about the role of government—not just by neglecting dual use problem 2, but also by a more general failure to take seriously the conflicts of interest to which government officials are often subject. Although the fact that government involvement brings risks has sometimes been acknowledged in the US biodefence debate, the chief risk has been assumed to be interference with the production of scientific knowledge. There has been no systematic exploration of the full range of risks involved or the sorts of institutional arrangements that may either magnify or reduce them. The fact that US institutional proposals have failed to distinguish the two dual use problems and to acknowledge that solutions to the former may exacerbate the latter are clear indications that the risks of government involvement have not been taken seriously, much less systematically explored.

Second, discussions that do assign an important role to institutions frequently assume that a particular division of labour among institutions and agents is appropriate, without providing good reasons for why this is so and without considering alternative arrangements. For example, some have advocated voluntary ‘self-policing’ of the dissemination of information by researchers or by researchers working with scientific journal editors, claiming that government oversight should either be avoided or kept to a ‘minimum’.^{21–24} Such proposals provide no evidence for the efficacy of ‘self-policing’, show little awareness of the conflicts of interest and limitations of knowledge about the risks of harmful misuses of research to which researchers and journal editors may be subject, and ignore the fact that the admonition to keep government involvement to ‘a minimum’ only makes sense within the context of an account of optimisation that they have not begun to provide. What is needed is a more critical and systematic exploration of solutions to the optimisation

problem, one that first applies cost–benefit analysis, broadly construed so as to accommodate moral values as well as efficiency, not to the choices of individuals as to whether to disseminate research particular results, but to the design of institutions, with the goal of developing an institutional division of labour whose overall result will achieve a proper balancing of biodefence with other values, including, but not limited to, the value of ‘open science’.

This institutional optimisation task is exceedingly complex, as others have acknowledged.⁵ To make headway on it we identify two key conceptual resources to advance the current debate over biodefence: the idea of social epistemology and that of institutional design.

SOCIAL EPISTEMOLOGY AS A RESOURCE FOR CONCEPTUALISING THE OPTIMISATION PROBLEM

The relevance of social epistemology

Social epistemology has been defined as the comparative assessment of the efficacy of alternative institutions creating, transmitting and preserving true or justified beliefs.²⁵ Institutions are understood broadly to include formal and informal norm-governed, relatively stable patterns of organisation that typically include an internal division of labour characterised by roles.

Social epistemology is grounded in three simple but powerful ideas. (1) Knowledge generally, including scientific knowledge, is largely a social, not a purely individual accomplishment. (2) Institutions (broadly understood) play a vital role in the social production of knowledge. (In this broad sense we can speak of ‘the institutions of science’ meaning the totality of persisting patterns of norm-governed interactions that constitute the scientific community.) (3) The institutionalised social production of knowledge requires a complex division of cognitive labour, but does not require any overall central authority to direct the process of knowledge production.²⁶ (In that sense, social epistemology proceeds on a very loose analogy with the ‘invisible hand’ explanations of market economics. Note: this is not to say that knowledge is best produced ‘in the private sector’). Peter Railton uses the term ‘the invisible mind’ here and provides a valuable discussion of the implications of a social epistemology approach for current debates about the objectivity of science.²⁷

Thus far, social epistemology has concentrated chiefly on the institutions of science, attempting to identify their ‘epistemic virtues’, the features of these institutions that contribute to the production of scientific knowledge (or, on some more cautious formulations, justified empirical beliefs). The task of optimisation with which we are concerned is more complex: to try to ensure that other important values, over and above the production of scientific knowledge, including biosecurity, are properly accommodated with the least detriment to the epistemic virtues of the institutions of science.

Nonetheless, a focus on the epistemic virtues of the institutions of science is a logical place to begin the more complex task. If the problem is to balance protection of the ‘norms of open science’ against other values, including biosecurity, then it will be important to know what role various ‘norms of open science’ actually play in the production of scientific knowledge; but to know this we need a social epistemology of scientific institutions. In brief, before we modify the knowledge-producing institutions of science in the name of biosecurity (or, more accurately, to solve an optimisation problem in which biosecurity is one value), it would be useful if we had some idea how the institutions of science produce knowledge. Saying that they do so through the operation of ‘norms of openness’ is

hardly adequate. Current work on the social epistemology of science indicates that many other factors besides 'norms of openness' play a role in the production of scientific knowledge and that there are discrepancies between the putative 'norms of openness' and how science actually works. It is remarkable that the current debate about biodefence and 'open science' has proceeded without even acknowledging the relevance of the social epistemology of scientific institutions.

Instead, various parties to the public policy discussion have made assumptions about what can and what cannot be changed in the enterprise of scientific knowledge production without undercutting its effectiveness, in the absence of any basis for making these assumptions. If, as current work in social epistemology indicates, the production and dissemination of scientific knowledge depends upon much more than 'norms of openness', then this complicates the policy response. In particular, it will not be sufficient to show that a particular policy does not unduly erode the 'norms of openness'. A policy that scored well on this count might nonetheless have the unintended effect of damaging some of the scientific knowledge-producing enterprise. This would be the case if, for example, the cooperation of scientists with government for the sake of biodefence diminished the credibility of scientists in the eyes of the public and called the objectivity of their findings into question.

From one standpoint we might say that the current policy debate suffers from a lack of awareness that there is systematic work in social epistemology that is directly relevant to it. From another we might say that the problem is that the current debate unwittingly operates with a very primitive, unarticulated and empirically unsupported 'folk' social epistemology according to which the (largely unspecified) 'norms of openness' are the only significant epistemic virtues of the scientific enterprise.

The limitations of social epistemology

As powerful as its key ideas are, mainstream social epistemology, although necessary for tackling the optimisation problem, is insufficient for several reasons. First, its theorists have tended to concentrate on efficacy, neglecting efficiency, in the production of scientific knowledge. In other words, they have focused on whether one institutional arrangement is better at producing knowledge rather than another, without taking into account differences in the costs of knowledge production. Even from a purely epistemic standpoint, setting aside for the moment the need to accommodate moral values, institutional arrangements that produce knowledge at lower cost are preferable and efforts to prevent dual use problems that needlessly raise the costs of producing scientific knowledge would be unacceptable. As our characterisation of the optimisation problem makes clear, however, the costs of producing scientific knowledge must not be restricted to financial costs or time costs, but must also include the risks of harm from accidents and malicious use.

Second, investigations of the epistemic virtues of scientific institutions have frequently proceeded on the highly idealised assumption that the scientific enterprise, as a knowledge-generating process, is largely free from government interference. All theorising requires idealisation, but this particular idealisation is extremely problematical in the perceived emergency situation in which the problem of devising ethically sound biodefence policy arises.

There is another sense in which mainstream social epistemology of scientific institutions neglects the political: it tends to view the relationship between the scientific enterprise and the public exclusively in epistemic terms, chiefly from the standpoint of investigating how institutional practices such as

educational credentialing and peer review of publications can help non-scientists identify genuine 'epistemic authorities', meaning especially reliable sources of true or justified empirical beliefs. This vantage point, although of great value, overlooks important issues concerning institutional legitimacy. An exception is Philip Kitcher's book, 'Science in a democratic society' (forthcoming, 2011), which discusses an apparent erosion of trust in mainstream scientific expertise regarding global climate change.

On the one hand, by identifying certain individuals as scientific authorities, the institutions of science create opportunities for government or other institutions to try to convince the public that their policies are legitimate by presenting them as scientifically informed. For example, government leaders may cite scientific estimates of the harm that would be done by a bio-terrorist attack to justify the claim that a state of emergency exists and to try to convince the public that the infringements of civil liberties that its response to the putative emergency entail are legitimate. On the other hand, whether the public regards scientists as genuine epistemic authorities can depend on whether the institutions of science are themselves viewed as legitimate. If the institutions of science are thought to be unduly influenced by government or by religion or ideology, then the credibility of scientists as epistemic authorities may decline and science may suffer a 'legitimacy crisis'. If this occurs, scientific knowledge may still be produced, but it will not be recognised as such by the public. The widespread denial of anthropogenic climate change may be an illustration of this phenomenon.

In addition, if the public comes to believe the legitimacy of the institutions of science has been seriously compromised, it may refuse to provide the resources needed to support them and this too may reduce their efficacy in producing knowledge. Legitimacy is an important value to be taken into account in thinking through the optimisation problem, then, with regard to both the legitimacy of scientific institutions and the role of science in contributing to the legitimacy of political institutions.

The concept of legitimacy is relevant in yet another, more fundamental way. The presumption should be that the overall policy effort to cope with the dual use problems must be compatible with the legitimate exercise of political power. The requirement of political legitimacy does not automatically disappear whenever there is a state of emergency and it certainly does not vanish simply because the government says that a state of emergency exists. 'Legitimate' as applied to political institutions is generally understood to mean 'having the right to rule' and the state is said to have the right to rule only if it operates within certain moral constraints, often specified in terms of individual rights.

Third, thus far social epistemologists concerned with understanding the epistemic virtues of scientific institutions have not explored in any depth the fact that these institutions, like institutions generally, are not only norm-governed, but are also venues in which existing norms are contested and new norms are developed. Thomas Kuhn, for example, focuses on conceptual change in the form of shifts to new paradigms of scientific explanation, not on norm contestation and change *per se*.²⁸ The issue of norm change is important for the optimisation problem under consideration here in two ways. First, in considering alternative institutional arrangements we must try to determine what role norms, understood as internalised rules, should play in the overall process of balancing biosecurity with the creation and dissemination of scientific knowledge and other important values. A workable solution to the optimisation problem might require modifying some of the norms that have until now

characterised the scientific community, perhaps by using new educational strategies and role modelling to try to instil a clear sense of responsibility for helping reduce the two dual use risks. In particular, we cannot assume that scientists will adopt new norms regarding responsibility for possible uses of their research simply because a new code of ethics says they should. The new norms may be weak unless they are reinforced by or are at least consistent with the incentives to which scientists are subject. Second, strategies for optimisation must also consider the possibility of unintended norm change. Institutional changes to cope with dual use risks might unwittingly erode some of the most valuable norms that constitute the institutions of science. For example, by encouraging the idea that scientists should play a key role in defending the nation, government, perhaps supported by the media, may encourage biases on the part of scientists that compromise the validity of their research.

Principles of institutional design: a primer

So far, we have argued that the central insights of social epistemology and institutional design are crucial for a sound public policy response to biodefence issues, properly conceptualised as a complex optimisation problem. Our purpose here is not to articulate a theory of institutional design but rather to sketch some of the elements of such a theory and, in this section, to show how they can be used to evaluate some key aspects of current biodefence policy. The following principles will be familiar to students of institutional design in the social sciences, but are remarkably absent from current biodefence discussions.

Successful institutions typically rely not just on norm-governed behaviour, but on a plurality of role-differentiated, indirect norms of action

Different agents, occupying different roles, can contribute to the achievement of institutional goals by acting on different and even sometimes conflicting norms. These norms do not direct agents to ‘achieve institutional goal G, G1, etc’, but instead prescribe specific actions or processes, which, taken together in the overall operation of the array of institutions, tend to promote institutional goals. Here an analogy with market economy explanations is useful: under ideal conditions markets produce efficient states in equilibrium, but not because various agents in the market follow the norm ‘produce an efficient state’. Instead, individual agents follow other norms—such as ‘price your goods so as to maximize your profit’—thereby producing behaviour whose aggregative effect is an efficient state. Similarly, the best way to balance biosecurity with other relevant values may not be to encourage scientists (or government officials, or members of a national science biodefence advisory board) to follow the norm ‘try to strike a reasonable balance between the values of “open science” and “biosecurity”’ as they decide whether particular research findings ought to be disseminated.

If institutional goals are to be achieved, norms are important, but incentives also matter

Institutional effectiveness depends on incentive compatibility (the absence of perverse incentives, ones that encourage agents to act in ways that thwart institutional goals). Institutional effectiveness also depends upon the complementarity of norms and incentives; in particular, incentives should be aligned so as to make an agent’s compliance with appropriate norms rewarding or at least not excessively costly to her. The idea of norm/incentive complementarity is perhaps less familiar and obvious than that of incentive compatibility, but it is equally powerful. Norms, understood as internalised rules, play a crucial role in

achieving desired institutional outcomes generally, but the power of norms can be either augmented or diminished, depending upon whether institutionally generated and extra-institutional incentives support or compete with them.

Institutional systems can be locally inefficient but globally efficient

What appears to be wasteful or even dysfunctional behaviour narrowly considered may make a positive contribution to overall efficiency or at least may not be eliminable without reducing the efficiency of the system as a whole. The judgement that some aspect of institutional performance is inefficient may rest on either a failure to see how it fits into the larger whole or on an overly narrow characterisation of the optimisation problem the institution is designed to solve. For example, allocating funds among a plurality of research teams might seem less efficient than simply funding the team that is best qualified to do the job, but spreading the funds may be more likely to produce competitive pressures that result in the best team performing even better.

In well-functioning institutions, the relationship between the motives of agents and desirable outcomes may be complex and even counterintuitive

For example, within the right sort of overall institutional context, it may be highly beneficial for scientists to be motivated not just by the commitment to producing knowledge, but also by the desire for prestige and financial reward. In the production of knowledge, as in many other institutional endeavours, self-interested motivations can, under the right circumstances, contribute to the greater good through what might be called constructive competition. Here, too, an analogy with markets is instructive, although one need not go as far as Mandeville did, when he proclaimed that ‘private vices’ are ‘public virtues’.²⁹

Just as it cannot be assumed that good collective outcomes require the absence of competition among agents or the moral purity of motivations, it cannot be assumed that optimisation is to be achieved through thorough-going intra-institutional or inter-institutional harmony

Conflict, including not only the clash of opposing ideas, but also the clash of interests, can be productive overall. Call this the Madisonian idea. The most obvious application of the Madisonian idea is to consider the role that a system of ‘checks and balances’ can play in an overall institutional optimisation strategy. For example, we should not assume that the best system for reducing the risk of dual uses (in either sense) is one that exhibits a thoroughly hierarchical structure of authority, with one entity at the top and all others subordinate to it. Instead, a degree of institutional competition and even some ambiguity about the ultimate locus of authority might be superior. In times of perceived emergency, there is a pronounced tendency to demand total harmony and cooperation; the Madisonian thesis emphasises that acceding to this demand can sometimes be self-defeating.

Because we can expect institutions to perform imperfectly and because institutional goals may need to be re-assessed in the light of new developments, sound institutions will include provisions for the critical revisability of both means and ends. Good information and effective incentives for utilising it properly are both essential to critical revisability. Other things being equal, institutional arrangements that insulate key actors from

criticism and that limit their sources of information to others who share the same institutional interests, ought to be avoided.

Thus far we have: (1) distinguished two dual use problems; (2) demonstrated that unless both are considered, solutions to the first may exacerbate the second; (3) shown that the real issue is a complex optimisation problem whose solution requires applying cost-benefit analysis broadly conceived at the level of institutional design; (4) explained how social epistemology provides valuable conceptual resources for tackling the optimisation problem; (5) identified shortcomings of conventional social epistemology that limit its usefulness in this context; and (6) offered a list of principles of institutional design for employment in constructing a solution to the optimisation problem. In the next section we illustrate the fruitfulness of this more comprehensive analytical framework—which we call the institutionalist optimisation approach—by using it to evaluate several current US biodefence policies.

PUTTING THE INSTITUTIONALIST OPTIMISATION APPROACH TO WORK

To illustrate better the virtues of the optimisation framework, we turn now to an overview of some of the key institutional changes implemented in the USA in response to the bioterrorism threat. Our aim is not to make an all-things-considered, thumbs-up or thumbs-down evaluation of the policy alternatives, but rather to show how the institutionalist optimisation approach can contribute to such an evaluation. Above all, our remarks are designed to show how this approach provides protections against the tendency to omit from consideration certain factors that ought to be prominent in public policy deliberations, but that were largely absent in the public discussions of US biodefence policy when these institutions were created.

NSABB and BARDA

In March 2004, in response to the anthrax attacks in the USA, the Department of Health and Human Services announced the formation of a new government entity, the National Science Advisory Board for Biosecurity (NSABB); the charter was renewed in March 2010. NSABB was implemented under 42 USC 217a, section 222 of the Public Health Service Act, as amended and Pub L 109-417, section 205 of the Pandemic and All-Hazards and Preparedness Act. NSABB is governed by the provisions of the Federal Advisory Committee Act, as amended (5 USC app). According to its charter the purpose of the board is to ‘provide, as requested, advice, guidance, and leadership regarding biosecurity oversight of dual use research, defined as biological research with legitimate scientific purpose that may be misused to pose a biologic threat to public health and/or national security’. In its first year NSABB worked to develop a definition of dual use research in order to inform the responsibilities of scientists conducting such research. They identified ‘dual use research of concern’ as research that ‘based on current understanding, can be reasonably anticipated to provide knowledge, products, or technologies that could be directly misapplied by others to pose a threat to public health and safety, agricultural crops and other plants, animals, the environment, or material’.³⁰ The board is charged with: (1) recommending strategies and guidance for those conducting dual use research, or those with access to select biological agents and toxins; (2) providing recommendations for educating and training scientists, laboratory workers, students and trainees about dual use research issues; (3) advising on policies governing publication, communication and dissemination of dual use research methodologies and results; (4) recommending strategies for promoting

international engagement on dual use research issues; and (5) advising on the development of codes of conduct for life scientists engaged in dual use research.³¹ The purpose and activities of NSABB take for granted the above definitions of ‘dual use’ and ‘dual use research of concern’, which while somewhat vague, track the definition of dual use 1 as we have discussed and implicitly present the problem as a trade-off between open science, and preventing the misuse of otherwise beneficial science for malevolent ends. The absence of attention to the risks that ‘good’ governments will misuse biological science indicates that the focus of attention has been primarily limited to dual use 1.

In December 2006 US Congress passed a biodefence bill, Project BioShield, and created a second new institution, the Biomedical Advanced Research and Development Authority (BARDA), whose sole purpose is to oversee funding for the development and purchase of vaccines, drugs, therapies and diagnostic tools in response to public health medical emergencies, including bioterror and pandemic agents.^{32–34} The rationale behind the agency and its management of Project BioShield is to speed up the procurement and development of potential countermeasures for chemical, biological, radiological and nuclear agents as well as medical countermeasures for pandemic influenza and other emerging infectious diseases that fall outside the scope of Project BioShield.³⁵ BARDA also manages the Public Health Emergency Medical Countermeasures Enterprise, which represents an attempt to offer ‘a central source of information regarding research, development, and acquisition of medical countermeasures for public health emergencies, both naturally occurring and intentional’. Comments regarding the purpose of the Public Health Emergency Medical Countermeasures Enterprise reflect an implicit stand on the overriding value a scientific enterprise that is able to respond quickly to a terrorist threat or unintentional pandemic: ‘Our nation must have a system that is nimble and flexible enough to produce medical countermeasures quickly in the face of an attack or threat, whether it’s one we know about today or a new one. By moving towards a 21st century countermeasures enterprise with a strong base of discovery, a clear regulatory pathway, and agile manufacturing, we will be able to respond faster and more effectively to public health threats.’³⁶

The overall strategy adopted by BARDA has been to channel funding to earlier stages of drug and vaccine development—the stage referred to in the drug industry as the ‘valley of death’, as companies are left to pay for research and development until vaccines and drugs are ready for use and government purchase. The agency received an initial budget of US\$1 billion over 2 years. Most recent budget figures for civilian biodefence in 2010–11 totalled US\$6.48 billion. Of that total, US\$5.90 billion (91%) has been budgeted for programmes that have both biodefence and non-biodefence goals and applications, and US \$577.9 million (9%) has been budgeted for programmes that deal strictly with biodefence.³⁷

NSABB has been criticised for its lack of transparency, but the deeper issue is appropriate accountability. As an illustration of concerns regarding NSABB transparency and accountability to stakeholders, see the public comments. See also the series from the activist group, The Sunshine Project, whose mission has been to shed light on the conduct of government-sponsored biodefence research.^{38–39} Transparent processes provide necessary but not sufficient conditions for holding institutions and members accountable. There are no provisions for holding NSABB members accountable, either as individuals or collectively, beyond their accountability to the federal government that

appointed them. Such an arrangement is suspect, to say the least, given that here, as elsewhere, the interests of the government and those of the public and other relevant parties, including the scientific community, are not perfectly congruent. In particular, under sustained conditions of the continued 'war on terror' government officials are under incentives that may lead them to exaggerate the risk of 'dual use 1' to the detriment of a proper accommodation of other relevant values. In brief, the federal government may be systematically biased towards the avoidance of 'type 1 errors' (in this case, a 'type 1 error' would be the failure to take adequate precautions against bioterrorist threats; a 'type 2' error would be taking more extensive protection than is necessary, at the expense of other values). If the federal government's standards for holding NSABB accountable reflect such biases, then to that extent accountability is inadequate.

The structure of BARDA raises similar concerns. Two features of the agency's structure are likely to increase the bias towards type 2 errors: First, the amount of funding earmarked for biodefence research as summarised above increases the incentives for individual scientists, research programmes and industry to join in bioagents research. This huge infusion of funds may actually increase the risk of bioterrorist attacks by increasing the number of individuals who have the knowledge and means to weaponise biological agents. The strong focus on dual use protection since the increased funding indicates some awareness of exactly this risk. The more funding goes to US scientists to investigate select agents, the greater the risk that published results and findings fall into the 'wrong hands'. What has not been openly discussed, but is also implicitly apparent in the need for international partnerships in addressing dual use risk, is that this influx of funding may well spur other governments to join in a global biodefence research race, which may in turn increase the risk of the accidental or deliberate use of bioweapons. The nuclear and offensive bioweapons programmes of the cold war give ample historical reason to believe the defensive race might well evolve into an offensive race. Another feature of BARDA that not only undermines accountability but may also exacerbate the risk of type 2 error is the exemption from the US Freedom of Information Act. While requests for information that are deemed non-threatening to national security may be honoured, such determinations will not be subject to judicial review, but rather, made internally by the agency. Given BARDA's primary goal—to expedite the research and development of bioterror countermeasures in preparation for a possible attack while protecting against malevolent uses of our own research products—the presumption will probably be one of non-disclosure. Without the possibility of independent judicial review a tendency towards type 2 error is built in to BARDA's very structure.

Even more obviously, making NSABB and BARDA accountable only to the federal government is inadequate from the standpoint of the dual use 2 problem. Indeed, it is difficult to imagine a more favourable arrangement than the current one, from the standpoint of those interested in developing a US offensive bioweapons programme. Both NSABB and BARDA are positioned to pass information relevant to government officials who might relay it to bioweapons research on to government agencies; to prevent the dissemination of this information to others; and to not reveal the fact that they are doing so.

We do not claim that NSABB or BARDA are engaged in such activities; nor are we claiming that either is making substantively wrong decisions. The point is that from the standpoint of institutional design both are deeply flawed, because they lack

appropriate accountability, lack safeguards against bias towards type 2 errors, and do nothing to reduce the risk of dual use 2.

NSABB does allow limited access to its publicly announced meetings, but retains the power to determine what the public should and should not know, without any acknowledgement of the need for safeguards against abusing this power. Under these conditions, the mere presence of outsiders during portions of the NSABB's public meetings is not worthy of being called an accountability mechanism. The point is not that the public should be allowed to determine which of the NSABB's proceedings it should be included in; rather, it is that there should be some provision for helping to ensure that the NSABB does not, wittingly or unwittingly, abuse its power to make this decision. Similarly, judicial review of freedom of information requests is essential to maintaining the accountability of BARDA research.

Accountability includes three elements: (1) adequate standards of performance for evaluating the behaviour of institutional agents; (2) appropriate 'accountability holders' to apply these standards to evaluate the behaviour of institutional agents; and (3) adequate capacity and willingness of some designated agent or agents to impose costs on agents for failure to perform according to the standards. Unless the standards of performance to which NSABB is held accountable reflect a clear awareness that its operations are one element in an overall response to a complex optimisation problem that includes efforts to reduce dual use 2 as well as 1, those standards will not be adequate.

Adequate accountability holders are those who can be relied upon to represent all the values relevant to the complex optimisation problem, not just the most pressing current concerns of the federal government. Presumably, the federal government has the capacity to hold NSABB and BARDA accountable, but it is unclear whether the government is willing to hold them accountable to standards that reflect the plurality of values relevant to the optimisation problem rather than those that mirror its own current most pressing concerns, including the desire to avoid a bioterrorist attack at all costs. The public currently has no good reason to believe that NSABB and BARDA satisfy any of the three elements of appropriate accountability.

The predictable reply to these criticisms will no doubt be this: accountability requires transparency, but under current conditions, transparency is not compatible with either NSABB or BARDA doing its job. This reply is not adequate. If contemporaneous transparency is too risky, then provisions could be made for ex post transparency, under more favourable future conditions, when the bioterrorist threat has abated somewhat. Arguably, we have been in that state for 10 years, but the maintenance of a chronic, heightened state of alert has contributed to the sense that transparency is permanently too risky in this area of research and policy. To our knowledge, NSABB has not even raised the possibility that there may be other ways of assuring accountability than full contemporaneous transparency. BARDA's exemption from judicial review on freedom of information requests is a clear case of pre-empting an essential mechanism for assuring accountability through both contemporaneous and ex post transparency (as it shuts off the possibility that judicial review might allow a delayed or selective release of information).

The only alternatives for appropriate accountability are not full contemporaneous transparency, ex post full transparency, or the current unsatisfactory lack of accountability. A fourth alternative is suggested by one of the key principles of institutional design listed above, the Madisonian idea. For example, a committee or subcommittee from the legislative branch or a special panel of individuals could be formally charged with reviewing

periodically explanations by the NSABB as to why it had excluded the public from its proceedings. The reviewing body would be chosen so that members would have interests and be under incentives that were not unduly congruent with the interests and incentives of NSABB members. Yet another alternative would be a formal ex post review of NSABB's performance after it its work is done, with concrete costs attached to a negative review.

Concerns have also been raised about BARDA's relationship to US federal research agencies, such as the National Institute of Allergy and Infectious Diseases and the Centers for Disease Control and Prevention.³³ However, the central worry is not that BARDA may create redundancies or inefficiencies, adding an 'extra layer of complexity' to biodefence research. Rather, it is that BARDA-funded countermeasures research may be exempt from the more rigorous ethical review processes for human and non-human research subjects.

At present no clear and publically available information is available about how BARDA proposals are reviewed and how expertise for the review process is determined. An example of the lack of clarity and, indeed, mystery surrounding the federal review and oversight of dual use research is reflected in the debate among top scientific publishers. Both are key factors for accountability necessary to the ethical conduct of research. In the post-2001 rush to respond to the bioterrorism threat in the USA, these institutions were created with no discussion or consideration of any other institutional alternatives. Ten years later, the institutions have become an accepted part of the US biodefence research enterprise, while still lacking mechanisms to address serious ethical concerns beyond the narrow understanding of dual use that shape the institutions' creation and ongoing activities. This is one more indication that the conceptual framework offered by social epistemology and institutional design can begin to elucidate the more complex ethical issues at stake, well beyond even the more sophisticated understanding of dual use.

CONCLUSION

We have argued that the ethical and policy debates regarding 'dual use' biomedical research ought to be reframed as a larger optimisation problem across a plurality of values including, among others: (1) the production of scientific knowledge; (2) the protection of human and animal subjects; (3) the promotion and protection of public health (national and global); (4) freedom of scientific enquiry; and (5) the constraint of government power. We have also argued that a fruitful response to the optimisation problem will employ the tools of social epistemology as well as sound principles of institutional design.

Our goal has not been to resolve any policy issue in the area of biodefence. Instead, our focus has been methodological but at the same time eminently practical. Given the preoccupation with the protection of science as a knowledge-producing enterprise, it is remarkable that there has been so little attention to identifying more precisely those features of the institutions of science that might be adversely affected by this or that policy initiative. Instead, participants in the debate have treated the institutions of science as a kind of black box, whose mysterious operations are achieved through something vaguely called 'the norms of openness'. In other words, they have rested content with a sparse 'folk' social epistemology of science, without asking whether anything more useful is available.

The current debate is equally remarkable for its lack of attention to the most rudimentary principles of institutional design. Even if the problem was simply that of 'balancing biosecurity with open science', institutional design would still be relevant.

Once it is seen that the real issue is a much more complex optimisation problem, the case for thinking explicitly about institutional design becomes all the more compelling.

The biodefence policy choices we make now may have profound effects, not only on the enterprise of science, but on the relationship between science and government, for years to come. The issues are difficult enough; there is no need to make them more intractable by poorly conceptualising them.

Competing interests None.

Provenance and peer review Not commissioned; externally peer reviewed.

REFERENCES

- 1 Merton RK. *On Social Structure and Science*. Chicago, IL: University of Chicago Press, 1996.
- 2 Merton RK. *The Sociology of Science: Theoretical and Empirical Investigations*. Chicago, IL: University of Chicago Press, 1973.
- 3 King NB. The ethics of biodefense. *Bioethics* 2005;19:432–46.
- 4 May T. Funding agendas: has bioterror defense been over-prioritized? *Am J Bioeth* 2005;5:34–44.
- 5 Miller S, Selgelid MJ. *Ethical and Philosophical Consideration of the Dual-Use Dilemma in the Biological Sciences*. New York, NY: Springer, 2008.
- 6 Atlas RM, Dando M. The dual-use dilemma for the life sciences: perspectives, conundrums, and global solutions. *Bio Secur Bioterror* 2006;4:276–86.
- 7 Alibek K, Handelman S. *Biohazard: The Chilling true Story of the Largest Covert Biological Weapons Program in the World*. New York, NY: Random House, 1999.
- 8 Miller J, Engelberg S, Broad W. *Germ: Biological Weapons and America's Secret War*. New York, NY: Simon and Schuster, 2001.
- 9 Pogge T. Human rights and global health: a research program. *Metaphilosophy* 2005;36:182–209.
- 10 Pogge T. World poverty and human rights. *Ethics Int Aff* 2005;19:1–7.
- 11 Kaiser Global Health. *Experts Discuss PEPFAR Funding in Congressional Hearing*. <http://globalhealth.kff.org/Daily-Reports/2010/March/15/GH-031510-PEPFAR-Funding.aspx> (accessed 13 Jun 2011).
- 12 DeFrancesco L. Throwing money at biodefense. *Nat Biotechnol* 2004;22:375–8.
- 13 Enserink M, Kaiser J. Has biodefense gone overboard? *Science* 2005;307:1396.
- 14 Buchanan A. Institutions, beliefs, and ethics: eugenics as a case study. *J Polit Philos* 2007;15:27.
- 15 Risks and benefit of dual-use research. *Nature* 2005;435:855.
- 16 Fischer JE. *Dual-Use Technologies: Inexorable Progress, Inseparable Peril*. Washington, DC: Center for Strategic and International Studies Press, 2005.
- 17 National Academy of Sciences. *Biotechnology Research in An Age of Terrorism: Confronting the 'Dual Use' Dilemma*. Washington, DC: National Academies Press, 2003.
- 18 Petro JB, Relman DA. Understanding threats to scientific openness. *Science* 2003;302:1898.
- 19 Shea DA. Balancing scientific publication and national security concerns: issues for congress. *CRS Report for Congress*. Washington: National Council for Science and the Environment, 2006. (RL31695).
- 20 Rawls J. Two concepts of rules. *Philos Rev* 1955;64:3–32.
- 21 Falkow S, Kennedy D. Letters: science publishing and security concerns. *Science* 2003;300:737–9.
- 22 Atlas R, Campbell P, Cozzarelli NR, et al. Statement on scientific publication and security. *Science* 2003;299:1149.
- 23 Keel BA. Protecting America's secrets while maintaining academic freedom. *Acad Med* 2004;79:333–42.
- 24 Sharp PA. 1918 Flu and responsible science. *Science* 2005;310:17.
- 25 Goldman AI. *Knowledge in a Social World*. New York, NY: Oxford University Press, 1999.
- 26 Kitcher P. The cognitive division of labor. *J Philos* 1990;87:5–22.
- 27 Railton P. Truth, reason, and the regulation of belief. In: Villanueva E, ed. *Truth and Rationality*. Atascadero, CA: Ridgeview Publishing Company, 1994:71–93.
- 28 Kuhn TS. *The Structure of Scientific Revolutions*, 3rd edn. Chicago, IL: University Of Chicago Press, 1996.
- 29 Mandeville B. (1723) *The Fable of the Bees, and other writings*. In: Hundert EJ, ed. Indianapolis: Hackett Publishing Company, 1997.
- 30 Office of Biotechnological Activities, National Institutes of Health. *Does Your Research Have Dual Use Potential: An Overview for Investigators*. NIH, OBA, 2010. <http://oba.od.nih.gov/biosecurity/pdf/EducationalBrochureDualUseResearch.pdf> (accessed 13 Jun 2011).
- 31 US Department of Health and Human Services. *Charter: National Science Advisory Board for Biosecurity*. 2004. http://oba.od.nih.gov/biosecurity/about_nsabb.html (accessed 13 Jun 2011).
- 32 Wadman M. Agency set up to tackle bioterror. *Nature* 2006;444:796–7.
- 33 Kaiser J. Critics question proposed countermeasures agency. *Science* 2005;310:755.

- 34 US Department of Health and Human Services. *Biomedical Advanced Research and Development Authority*. <http://www.phe.gov/about/barda/Pages/default.aspx> (accessed 13 Jun 2011).
- 35 US Department of Health and Human Services. *Public Health Emergency Medical Countermeasures Enterprise Review*. <http://www.phe.gov/Preparedness/mcm/enterprisereview/Pages/default.aspx> (accessed 13 Jun 2011).
- 36 US Department of Health and Human Services. *Countermeasures Review Rollout*. 2010. <http://www.hhs.gov/secretary/about/speeches/sp20100819.html> (accessed 13 Jun 2011).
- 37 Franco C, Sell TK. Federal agency biodefense funding, FY2010–FY2011. *Biosecurity* 2010;8:129–49.
- 38 Grant RW, Keohane RW. Accountability and abuses of power in world politics. *Am Polit Sci Rev* 2005;99:1–15.
- 39 The Sunshine Project. *Biodefense Blackout: Texas BSL-4 Lab Keeps Records Secret UTMB Resists Attorney General's Ruling, Case Moves to the Courts*. 2007. <http://www.sunshine-project.org/publications/pr/pr280207.html> (accessed 11 Dec 2011).