8:

TOWARD AN ETHICS OF NEUROSECURITY

I USE THE TERM "NEUROSECURITY" to refer both to the ways that science and technology targeted at the brain and nervous system should be managed for the public good, and the means that democratic states must develop to protect themselves from their adversaries. As in the fields of biosecurity and atomic security, neurosecurity is complicated by the problem of dual use, the underlying and inescapable fact that medical and scientific breakthroughs can also be used for purposes unrelated to the goals of the researchers. The dual use problem is as old as human ingenuity itself. Even fire and the wheel were likely applied to intertribal conflicts as soon as the opportunity arose.

In a 2003 report on biotechnology and terrorism, the National Academy of Sciences, the nation's most prestigious scientific body, defined dual use technologies as "technologies intended for civilian application that can also be used for military purposes." The report went on to note that universities, private industry, and government labs are doing important experiments to find new treatments for AIDS, cancer, diabetes, and bacterial diseases as well as neurologic disorders such as Alzheimer's and stroke. Much of this research is directed to finding ways to detect harmful microbes and chemicals in the environment and to developing vaccines. However, "weaponization" of materials in labs is a concern that is taken much more seriously now than before 9/11 and the anthrax attacks, as it is feared that a rogue scientist might make off with a dangerous agent. This is one of the theories regarding the five anthrax deaths in the fall of 2001, but the crime has never been solved.

NEUROSCIENTISTS AND NEUROSECURITY

One theme I want to highlight in this chapter is the need for the scientific community to be more engaged in dealing with the unintended consequences of their work. Michael Moodie, the former director of the Chemical and Biological Arms Control Institute, has observed that "the attitudes of those working in the life sciences contrast sharply with the nuclear community. Physicists since the beginning of the nuclear age, including Albert Einstein, understood the dangers of atomic power, and the need to participate actively in managing these risks. The life sciences sectors lag in this regard. Many neglect thinking about the potential risks of their work." My experience suggests that an increased sense of the need to be publicly involved is taking hold among life scientists, especially in the face of recent controversies about stem cell research and intelligent design. Questions of dual use also require the informed engagement of our best scientific thinkers.

The dual use issue becomes more pressing as the science becomes more powerful and as more people possess the knowledge to apply it. We've seen that the applications of neuroscience and other brain-targeting fields to national security are no longer in the realms of science fiction or paranoid fantasy, that tremendous advances have been and are being made in understanding the way the brain works and, more slowly, in modifying it. Even though some of the claims that are being made are likely exaggerated, especially by companies trying to sell their products, not all are. Separating the wheat from the chaff is a challenge, but it does seem clear that the fascinating science I've described is on a course that, although not wholly predictable, almost surely points to greater understanding of and control over brain-related processes, and from various approaches.

When I've raised questions about the dual use implications of advances in brain science and technology at scientific meetings, many neuroscientists are surprised. Although they may receive Pentagon or CIA funding, brain scientists generally don't regard themselves as contributing to warfare. Those whose research is funded wholly by civilian agencies are taken aback when I suggest that their published results might well be examined by national security agencies to assess their implications. A number of neuroscientists have told me that they have received phone calls out of the blue from security officials interested in their work in areas such as monitoring or altering neural processes. Among those researchers who do accept national security agency funding, some tend to dismiss the idea that anything of military use will come of their research. Some believe, or prefer to believe, that they can manipulate their funding sources so that they can do the work they want to do without serving the goals of their benefactors, and that their results are going to be benign no matter what others might be looking for.

Often it's true that scientists are smart enough to get their grants without delivering the goods their funders want. But in the long run, as enough knowledge is gathered, the opportunities for dual use can't be completely avoided. For those who are deeply concerned about the exploitation of science for military purposes, an obvious answer seems to be that the scientific community should simply swear off cooperation with national security agencies, including accepting research contracts. Call this the purist approach. Based on some historical experience I shall elaborate, I believe the purist answer is shortsighted. In the real world, this kind of research is going to continue, and it's best that university researchers be those who do it, rather than building top secret science fortresses with researchers who are not answerable to anyone but their commanders. It is critical for the well-being of our democratic society that the civilian scientific community is kept in the loop and that the rest of us can have at least a general idea of the kind of work that is being done, even though for legitimate reasons many of the details may not be generally available.

An important reason to keep the scientific process as normal as possible, including transparency in interactions among scientists, is that science sets an example for an open society in which secrecy is minimized. Secrecy makes it harder for our elected representatives to fulfill their constitutional responsibility of overseeing government-funded science, and for experts outside of government to contribute to sound policymaking. One way a democratic society can minimize secrecy is to keep national security agencies linked to the larger world of academic science. For the same reason, suggestions in Congress and elsewhere that DARPA should pull back on its external funding should be resisted. The link between the academic world and the national security establishment makes for a healthier society than if each were isolated from the other.

There is of course another good reason that our best and brightest scientists should not reject this relationship: in a dangerous world, we do need to be protected. Though we should worry about the proliferation of new weapons and the diversion of resources necessary to develop them, the fact is that there are some rather malevolent forces out there. I am no pacifist and I do not advocate unilateral disarmament. While we need to be vigilant about dual use and the way our liberties may be affected, we also need to acknowledge the realities of modern threats to public safety. If our brain science and technology can provide us with some advantages, within carefully crafted constraints, it would be foolhardy, and perhaps immoral, not to explore them.

NO "NEUROPREPARAT"

Readers might be surprised that I reach this conclusion after cataloging the historic and modern examples of the not always admirable interest that national security agencies have shown in the brain and mind. Why wouldn't I advocate that our growing knowledge of the brain simply be made off-limits to military exploitation? For that matter, why not urge that no new science be studied for its potential in helping wage war? While admirable as an aspiration, in practice this position would lead to baleful consequences. In designing policy, we must acknowledge that neither science nor its martial applications can ever be static. While we might wish that no new kinds of weapons were added to the terrifying and massive arsenal already at the disposal of fallible human leaders, we need to find practical ways to address the problem. Similarly, rigorous separation of military research and civilian science will only result in war planners locating all research in facilities off-limits to public view (even more than is now the case), and creating a cadre of scientists who are beholden only to their government masters. Down that road lies the danger of powerful science fully captured by the state and beyond the reach of civilian control.

The history of science shows why we need to keep a close eye on the dual use problem. Though concern about dual use in biology is relatively recent, the atomic physicists of the 1940s and their intellectual descendants have been worrying over this problem since even before the atomic bombs were dropped at Hiroshima and Nagasaki. What lessons can dual use experiences in biology and physics provide as we face similar challenges in the applications of neuroscience?

Like the brain sciences, the development of modern microbiology is closely linked to the possibilities of warfare. The pioneering nineteenth-century microbiologists Louis Pasteur and Robert Koch developed techniques to isolate and culture microbes like anthrax so that medical defenses could be devised, but in short order their discoveries were applied to making biological weapons. In microbiology as in chemistry and atomic physics, the line between knowledge for defensive purposes and knowledge for offensive purposes is vanishingly thin.

Although the idea of using germs in warfare has been around since the ancient world—generally in crude forms like flinging the carcasses of infected horses over battlements—modern microbiology makes it feasible to select and even design agents in a "rational" rather than merely empirical manner. During World War I, Germany was accused of infecting horses of enemy forces with glanders, a highly contagious bacterial disease, and Japan engaged in what can only be described as an industrial effort to make massive quantities of bacteria in Harbin, Manchuria, during World War II, accompanied by horrific human experiments. At the same time, the United States, the United Kingdom, and Canada worked together in a secret offensive bioweapons program through the Army Chemical Warfare Service, though no deaths were intended or attributed to the Allies' efforts.

After the war, many bioagents were used in military and intelligence research until the United States ended its biological weapons program in the late 1960s. Much of the work was done at Fort Detrick in Frederick, Maryland, where treatments were devised for a number of infectious diseases with patriotic volunteer soldiers largely drawn from the Seventh-day Adventist Church. In 1972, the United States signed the international treaty on bioweapons, the Biological and Toxin Weapons Convention (BTWC), which made clear that only research for peaceful or defensive purposes is allowed. The Soviet Union was also a signatory, but didn't believe that the American program had really ended. In their suspicion, the Soviets developed a huge top secret system called Biopreparat that operated under civilian cover until at least 1992. Biopreparat continued to function even as Mikhail Gorbachev was deconstructing the country.

Ken Alibek was second in command of the former Soviet Union's biological weapons program until his defection to the United States in 1992. In his book *Biohazard,* Alibek describes the mammoth secret establishment put in place to counter suspected U.S. intentions in biowarfare development.

> Our program paralleled the Soviet nuclear program in organization and secrecy. Both generated a sprawl of clandestine cities, manufacturing plants, and research centers across the Soviet Union. The atomic weapons network controlled by the Ministry of Medium Machine Building was much larger, but the production of microbes doesn't require uranium mines or a massive work force. When our biological warfare program was operating at its peak level, in the late 1980s, more than sixty thousand people were engaged in research, testing, production, and equipment design throughout the country. This included some thirty thousand Biopreparat employees.

The Soviets in effect re-created the intellectual and material resources that were equivalent to several major research universities, but instead of using institutions that were integrated into the rest of the society, they created a huge clandestine system of science. Perhaps only an authoritarian system like that of the former Soviet Union could manage to sustain such a massive covert scientific establishment, or perhaps the creation of such a system would help lead to an authoritarian state. Either way, it is clear that the isolation of national security–related science from the larger scientific community is neither a viable nor a desirable option for an open society.

Moreover, Biopreparat's lack of accountability to Soviet society, repressed though that society was, arguably undermined Soviet quality control systems, resulting in threats to the public's well-being, such as the accidental release of anthrax from the biological weapons facility in Sverdlovsk in 1979. The exact death toll is unknown, but estimates range from sixty-six (the official Soviet figure) to over a hundred. A thorough KGB cover-up of the actual events—a defective filter on fermenting equipment was removed by a worker but never replaced—fooled even distinguished American scientists sent to investigate. Internally, military and government officials were in a conflicted position, as no one wanted to point fingers of blame. "The determination with which Soviet officials set about concealing the Sverdlovsk leak from their own people as well as the world was, under the circumstances, not surprising," Alibek wrote. "The truth would have severely embarrassed the nation's leaders, many of whom were not even aware that biological arms production was under way, and caused an international crisis." In representative democracies, both legislative oversight bodies and independent watchdog organizations play a significant role in keeping responsible parties accountable.

There are also security reasons to resist locating all potentially sensitive research in top secret institutions walled off from the rest of the scientific world. Secrecy will encourage states that feel threatened by our intentions to engage in exactly the sort of proliferation we don't want. Even the mere suspicion of secret activities is enough to encourage those who champion proliferation in their own systems, as was the case in the Soviet Union. Rather, we need to reassure others through the most transparent verification programs possible. Other countries may also use our secret neuroscience as an excuse to cloak their own programs. And if our superior defensive capabilities are known to potential adversaries, they will be less likely to be interested in probing for weaknesses. Progress in neuroimaging techniques that provide evidence that an individual is familiar with a certain place (a terrorist training camp, for example) could prove helpful in distinguishing between more and less likely sources of information among captives. Similarly, many scientists argue that the best way to ensure there are defenses that can neutralize new weapons of any kind is to allow the scientific community at large to learn about them—within reasonable limits of course—as scientists can then also learn what their weaknesses are and what countermeasures can work. Secrecy about science is not necessarily good for our security. It can have exactly the opposite effect.

The shadow of Biopreparat survives. Fears persist that some of the vicious new biological agents that were being attempted in those Soviet-era labs may be stored somewhere, in spite of vigorous attempts to identify and clean up the sites. Certainly the expertise that was created is stored in the heads of many ex-Soviet scientists, and the United States and other countries have made strenuous and apparently successful efforts to ensure that this two-edged sword of scientific knowledge is not hired out to the highest bidder.

LESSONS FROM ATOMIC PHYSICS?

My view that it would be bad for science and for our society for the neuroscience community to insulate itself from support from national security agencies does not imply that research should be unconstrained. Rather, I believe that the neuroscience community needs to be part of discussions about the conditions for entering into relationships with security agencies and the guidelines that would govern the research. Many of these rules, such as the standards that govern experiments involving human subjects, are already in place but are of a technical nature. There are larger philosophical questions about the social obligations of scientists.

The most dramatic case of a scientific community grappling with its moral responsibilities regarding military applications of its work is that of the atomic physicists and nuclear weapons. Although differences of opinion appeared from the very earliest days of work on the superweapons, for many physicists the turning point came after the development of the hydrogen bomb, a weapon of far greater destructive capacity than the original atomic weapons. A number of them turned decisively toward the antiwar activists and struck up an alliance. The leader of that group was the British philosopher Bertrand Russell, a longtime socialist. Einstein decided that one of his last acts would be to join his prestige with that of Russell, and together they drafted a statement known as the Russell-Einstein Manifesto. Published on July 9, 1955, the manifesto expressed the gloom that had descended upon many of the original Manhattan Project scientists:

> No doubt in an H-bomb war great cities would be obliterated. But this is one of the minor disasters that would have to be faced. If everybody in London, New York, and Moscow were exterminated, the world might, in the course of a few centuries, recover from the blow. But we now know, especially since the Bikini test, that nuclear bombs can gradually spread destruction over a very much wider area than had been supposed. . . .
>
> There lies before us, if we choose, continual progress in happiness, knowl-

> edge, and wisdom. Shall we, instead, choose death, because we cannot forget our quarrels? We appeal as human beings to human beings: Remember your humanity, and forget the rest. If you can do so, the way lies open to a new Paradise; if you cannot, there lies before you the risk of universal death.

The manifesto resolved that only the abandonment of war by governments, and especially by the United States and the Soviet Union, could rescue humanity from imminent catastrophe. Although war has surely not been renounced, the Russell-Einstein statement did catalyze an international anti–nuclear weapons movement that arguably helped create the test-ban treaties of the 1960s and the drawdown of armed missiles in the 1980s. Gradually, the atomic physics community shifted its position to a "no first use" philosophy, which seemed a far more achievable constraint to which all nuclear nations could, at least rhetorically, commit themselves. The manifesto and the movement it stimulated showed that scientists, especially Nobel laureates, could employ their prestige to influence the political use of their science.

LIMITS OF THE BOMB ANALOGY

Unfortunately, "no first use" doesn't fit well in the context of neuroscience and national security. There are a number of important differences between the atomic physicists' experience and that of the neuroscientists; most neuroscience is aimed at either healing directly or understanding the brain well enough to do so. But the original atomic physicists, after some initial uncertainties about the technical possibilities of nuclear fission, knew that they were developing a weapon of unprecedented destructive capacity. It was their intention to do just that, out of fear that the Germans would get there first. In fact, the war effort's funding and support for the Manhattan Project paved the way for the knowledge and technologies that made possible the peaceful uses of atomic energy, such as reactors to produce electricity. One could almost say that atomic energy was dual use turned on its head, as military research and development created the conditions for the civilian use.

Another important difference between nuclear and neuroscience weapons is that, as became clear to war planners in the 1950s, especially after development of the hydrogen bomb, nuclear weapons are musclebound. Although “improvements” have been made that modify the nuclear bomb’s fission release, in general the bomb is too powerful and its effects too uncontrollable to provide tactical advantage on a battlefield. What nuclear bombs mainly afford their owners is strategic advantage by influencing the behavior of potential adversaries. Thus “nuclear blackmail” and “mutual assured destruction” became familiar concepts during the cold war. The utility of neuroscience-based weapons, however, is mainly tactical, in that they might provide short-term and relatively targeted advantages such as disrupting an enemy patrol or disabling a terror cell. In that sense neuroweapons are much more manageable than nuclear weapons. Not weapons of mass destruction, they are better considered weapons of selective deception and manipulation.

Yet another difference is that it’s much easier to distinguish between the offensive and defensive use of atomic bombs than offensive and defensive applications inspired by or based on neuroscience. As I’ve noted, innovations that focus on the brain and nervous system may be applied to enhancing the training, selection, and prospects for survival of troops long before they are deployed, if ever. An imposing advantage could discourage adversaries or at least create various advantages in small confrontations. And, according to some, the atomic weapons industry creates harrowing environmental problems that have caused critics to conclude that the bomb is a curse with few long-term benefits for humanity.

Finally, the neurosciences and related fields may well lead to measures that both give us an advantage over our adversaries and are morally superior to other tactics, a combination of considerations that doesn’t so easily apply to nuclear weapons. An example is interrogation. Naively, torture might seem like the easiest way to get someone to talk. But a brain bombarded with painful stimuli is going to have activated neural systems associated with fear and survival rather than cognition. I spoke about the shortcoming of violent interrogation techniques with Michael Grodin, a psychiatrist and bioethicist at Boston University who has worked with torture victims for years. “I’ve got experience with eight hundred torture victims and no one has ever been able to show that any critically useful information comes from torture,” he told me. “And the more severe the torture, the more problematic the so-called information from the victim.”

Scientifically informed techniques and devices seem likely to provide far more subtle ways to obtain information that is not polluted by the stress of talking under torture. Functional MRI studies of people who are playing a trust game with money have found increased activity in the caudate nucleus as one player learned to trust the other to invest his money. These signals appeared more quickly in the brain as the game went on and one player gained more confidence in the other. The caudate nucleus is linked to the brain's reward pathways so that it is more activated when there is an expectation of a positive event, such as being given some juice or, in this case, the socially rewarding experience of feeling good about someone else. If these signals could not only be monitored during interrogation procedures but the relevant pathways deliberately stimulated—setting aside for the moment ethical and legal concerns about peering into neural processes—it could advance the conditions for fruitful interrogation.

All in all, it doesn't look as if neuroscientists can take the same position about the often unintended fruits of their labor as the atomic physicists, that a "no first use" policy should prevail and they should only be used for defensive purposes. That policy makes sense for weapons of strategic value that might stimulate an uncontrolled exchange of weapons of mass destruction such as nuclear bombs, but it doesn't seem to work for tactical weapons that might head off more violence by, say, compromising a terrorist group holding civilian hostages. Nor is it clear how to apply such a policy to devices used in screening or training soldiers, or pharmaceuticals designed to improve a fighter pilot's cognitive capabilities. It seems that grappling with the ethical issues raised by the applications of neuroscience to national security will require an entirely different approach from that taken by the atomic scientists.

LESSONS FROM BIODEFENSE?

The field of biodefense is spawning an ethical debate of its own, and some of that work is relevant to neuroscience. For example, Ronald Atlas, a professor of biology and biosecurity expert at the University of Louisville, has developed a code of ethics for biodefense research. His code requires scientists to avoid doing anything to facilitate bioterrorism and obligates them to call the public's attention to any such activities, to restrict access to information that could lead to dual use, to ensure that the benefits of the research outweighs the risks, to respect the rights of human research subjects, and to respect the conscientious objections of those who decline to participate in the research.

These are admirable standards, and several apply directly to neurodefense research. One important practical difference is that research involving bioagents often results in potentially dangerous materials being kept in laboratories, so the physical security of those labs poses concerns that are not so common in neuroscience. But the emphasis on the social responsibilities of researchers, the obligation to consider risks and benefits and to protect the rights of those who might be used in experiments, is a useful precedent for neuroscience and national security.

The rules governing the biological weapons field present some interesting dilemmas when applied to neuroscience weapons. For instance, the Biological and Toxin Weapons Convention prohibits the use of bioweapons to manage civil disturbances such as riots. It makes sense not to allow law enforcement agencies to use anthrax, but perhaps not to ban the use of fMRI for "lie detection" or even hypersonic sound to root out terrorists, for instance. Whether these are acceptable remains to be debated, but at least they raise different issues than bioweapons.

Potentially undermining any measures intended to police biodefense research and development is the intensified secrecy surrounding biodefense research. Again, this experience should serve as a warning in the area of neurodefense. The Federation of American Scientists' bioweapons expert, Barbara Hatch Rosenberg, has written in the journal *Disarmament Diplomacy* that for about the first fifteen or twenty years after the Biological and Toxin Weapons Convention came into effect in 1975, the Defense Department kept its program unclassified, except for results that could reveal "U.S. military deficiencies, vulnerabilities, or significant breakthroughs in technology." After that, perhaps as a result of the Gulf War and the revelations about the Soviets' Biopreparat, the policy seemed to change. Just one week before 9/11, the *New York Times* reported on three secret biodefense projects that strain the treaty. One was a plan to create a vaccine-resistant anthrax strain that it is thought the Soviets had produced, ostensibly to learn how to defend against it. Although U.S. Western allies were disturbed by these revelations, they muted their reactions due to the 9/11 attacks.

The secrecy problem returns us to the need to retain connections between civilian science and the national security establishment. How can citizens make judgments about the policies their leaders are carrying out on their behalf if the information they have is so limited? It can't be assumed that universities are fonts of openness, however. They often participate in private financing arrangements that protect intellectual property. Universities also have differing policies for handling classified research. These policies should be the subject of public discussion and standardization so that the academic world does its part by requiring the greatest possible transparency.

For some kinds of science, greater use should be made of inspection regimes, but recent political decisions have hampered these arrangements. As Rosenberg writes in her article, "Secrecy is particularly corrosive, especially when combined with rejection of international monitoring. Suspicions would be largely dispelled if threat assessment projects were openly declared and subject to international inspection. There would be no need to disclose project results that impinge on national security." Obviously, finding the balance between national security and democratic openness is easier said than done, but the right start seems to be to make transparency rather than opacity the default position.

ETHICALLY REGULATING NEURODEFENSE

There's no easy fix to these issues, no bumper sticker solutions. They will have to be considered carefully and dispassionately by neuroscientists, agency officials, and representatives of the public. Many government advisory committees are already structured to provide and receive input from various sources. Rather than sweeping policies, the diversity of the neuroscience and its applications will challenge our ability to craft policies that tie familiar ethical concepts such as respect for personal autonomy to specific neuroscience-based techniques.

The mechanism for many of the decisions about the appropriate research and use of brain-based national security measures is already in place. The Department of Defense and the CIA are included in a federal regulatory framework called the Common Rule designed to protect human research participants; it requires review of proposals by a research ethics committee and the informed consent of the volunteers. The FDA regulates the licensing of drugs and devices. Though this system is far from perfect, it at least creates obstacles and some measure of accountability for the dissemination of much of the technology I've described.

But, clearly, the current regulatory system doesn't automatically answer all the questions about proposals for experiments with new drugs and devices or applications for licensure that might come before it. In some instances the technology is already licensed and can simply be applied to novel uses without going through the FDA process, as in the case of calmative drugs licensed as anesthesia for surgery. However, if experiments are to be done to see how well they work for some new purpose, such as managing a hostage situation, then informed consent and prior review are generally still required. All this could be bypassed if military authorities request a waiver of informed consent on national security grounds, as happened during the first Gulf War when agents thought to be protective against nerve gas and biological weapons were offered to troops even though the agents had not been approved for that purpose. Some policies and procedures will need to be in place for those experiments and applications of the new neuroscience and related fields that can't be captured in the routine regulatory process because of national security needs.

Often, the variety and potential usefulness for national security of innovations intended to affect the brain will require close ethical analysis. I asked Bill Casebeer, the terrorism expert and neuroethicist, how decisions are going to have to be made. His answer was partly framed in the technical language of ethics:

> Consideration of virtue theoretic [what are the moral potentials in the situation], deontic [what duties are at stake and to whom], and utilitarian [likely empirical consequences] aspects of the problem will be a useful starting point; I'd be surprised if there's anything that's entirely new in the field of military ethics that would be posed by consideration of neuroscientifically informed military operations. Considerations of virtue and vice, human functionality, rights and duties, consent, innocence, involvement in the causal and logical chain of agency required to do harm to another, outcomes in both the act and rule sense, etc., will all be operative. Some emerging technologies may make these considerations more pointed (e.g., does marketing informed by neuroscience somehow diminish the agency of those involved in marketplace transactions?), but the "entirely new" issues will be few and far between, I suspect.

I agree with Casebeer that in dealing with emerging neuroethical dilemmas in the national security context we can learn from previous ethical quandaries, especially in terms of the conditions under which we have those discussions. I also agree that in many cases, the ethically acceptable course of action will be a matter of weighing and balancing rather than appeal to an overarching moral doctrine, though basic guidance from some principles is going to be needed. For instance, a number of the scientists, lawyers, ethicists, and advocates with whom I spoke in the course of writing this book agreed that there had to be vigorous protection of at least one nonnegotiable premise when considering the appropriate security applications of neuroscience. In the law, this principle might be expressed in terms of the protections afforded in the Fifth Amendment of the Constitution regarding self-incrimination: "to be a witness against himself." Philosophically, this can be expressed as the proposition that no one else should be able to decide what goes into my brain or who "reads" it.

Like any philosophical principle, this one admits exceptions if they can be justified. One exception might be the example frequently cited by proponents of at least some limited torture option: the terrorist in our custody who is aware of a ticking bomb that could kill and injure many civilians. In that sort of case, advanced neuroscience-based technology might be helpful and its application justified. But justifiable exceptions to principles do not undermine their general validity. Suppose we can agree that, in any national security question regarding the brain, the presumption is that cognitive liberty is guaranteed in the absence of an overwhelming counterbalancing argument. The difficulty, of course, is that the need for the sovereign state to defend itself can easily be used as a trump card by legitimate political authorities. Under those circumstances, we have to rely upon some legal process to constrain state power. Again, the maximum possible transparency and accountability will have to apply.

One idea is to create the neurosecurity equivalent to the National Science Advisory Board for Biosecurity that was established in 2004. This new board is administered by the National Institutes of Health, but advises all cabinet departments, including the Defense Department, and "others as appropriate." Its mission is "to provide advice to federal departments and agencies on ways to minimize the possibility that knowledge and technologies emanating from vitally important biological research will be misused to threaten public health or national security." Focused particularly on the problem of dual use of biological agents, its charge includes developing guidelines for research and professional codes for scientists. The government also needs advice on what the most likely biological threats are and what sorts of countermeasures should be developed. Another panel, the Committee on Biodefense Analysis and Countermeasures, on which I serve, was created by the National Academies in response to a request by the Department of Homeland Security.

An advisory committee on neurosecurity could cover analogous problems of dual use and the implications of countermeasure development. If it was created by the National Academies, it could draw on talent from a range of disciplines. Diversity of scientific input is crucial for addressing national security issues. Scientists tend to work in silos, concentrating on their own discipline or subdiscipline or target of interest such as a certain organ or system, gene, or protein. Normally, this focused way of working is productive, but when novel real-world problems emerge, it can be overly limiting. National Academies committees are able to overcome these disciplinary boundaries.

An example of the need to apply several disciplines to neurosecurity problems is the way that genetically engineered biological weapons can become very scary neuroweapons. Bioweapons such as viruses have a payload: the genetic content of the virus; a delivery system, the outer viral coat; and a target, such as an organ system of the human body. All three components of the weapon system can be manipulated by pathogen genetic engineering. For example, certain viral and bacterial pathogens can be engineered by insertion of foreign or synthetic genes with properties not naturally found in the virus or bacterium to become advanced neuroweapons targeting the brain and nervous system. Based on work already done in the offensive biological weapons program of the former Soviet Union, scientists who are expert in biological weapons defense have worried aloud to me about the threat of technological surprise posed by advanced viral neuroweapons carrying synthetic genes coding for short peptides (short strings of biologically active amino acids with biological activity) into the central nervous system.

Inside the central nervous system, the technological surprise stems from designer peptides produced from synthetic genes that have effects quite distinct from those normally associated with the pathogen. For example, when produced in the brain, they could function as malign neuromodulators, disabling brain functions by modifying the relationships and communications between neurons. In such advanced neuroweapons, the infectious pathogen is really just a Trojan horse, selected for its ability to get the synthetic gene quickly into a target it cannot otherwise reach.

The advanced neuroweapon does not necessarily have to enter the brain and nervous system to modulate function. An example is *Francisella tularensis* genetically modified to produce beta-endorphin. This bacterium, a well-known biological weapon in its native form, is the cause of tularemia, also known as rabbit fever because it's often found in rodents and can be passed to humans by direct contact or by inhaling the particles. Normally, tularemia can easily be treated with antibiotics, but if the bacterium has been engineered to generate a potent neurochemical, the damage would already have been done before the infection became a clinical problem. Various kinds of disabling reactions, from intense fatigue or confusion to the loss of sensation, could be attempted that would neutralize enemy forces.

I talked at length about this problem with a twenty-year U.S. biodefense expert who preferred not to be identified by name. He argues that rapid-onset, brain-targeted biological weapons are something we need to worry about now, not sometime in the future. "There is no point on the battlefield in exposing the opposing troops to a synthetic gene that is going to give them liver cancer in fifteen years or make them incapacitated next week," he told me over a pleasant lunch in Rosslyn, Virginia.

> There must be an immediate purpose, such as to disrupt the characteristics that enable men to protect themselves and to fight as an organized force, not a rabble with weapons. There is some interesting history here. Before the French Revolution, soldiers fought for pay or because they were given lands and had a vested personal interest. Astonishingly, and for the first time, the French Revolutionary Armies fought enthusiastically and very successfully for abstract concepts—liberty, equality, and fraternity. Clausewitz and others studied this phenomenon, and a whole body of doctrine evolved about what it takes to employ a citizen army of millions of enlisted or drafted soldiers who are fighting for patriotism and loyalty with no real personal gain or stake. The ability of units to function is all in the mind. So if one can disrupt unit loyalty

> through fear or another emotion, the army would cease to exist as a fighting force. Claustrophobia would make soldiers tear off their protective face mask. Fear, thirst, accelerated heart rate, hypermotility of the gut—these would be the desired peptide effects.

Clearly, any attempts to engineer neuroweapons would be a violation of international law governing biologic and toxin warfare. But even with an inspection regime, there are no guarantees that the treaties will be honored. Reports that the Soviet Union was working on tularemia and viruses as candidate neuroweapons demonstrate that the genie is already out of the box. "A literal handful of people have been thinking about these issues for a decade or more, including in neuroscience," the expert told me, "but these are not the people who are now 'biosecurity experts' with lots of funds! We have got to think about these pathogens as weapons delivery systems and the critical scientific inputs we need are not those of infectious disease physicians but of those who can anticipate what might be inside the Trojan horse before it opens—because very soon after it does open, as in Troy, it might be impossible to mount a defense."

A national science advisory board for neurosecurity should have the clout to put all these neuroweapons issues on the front burner by combining a practical orientation with diverse scientific, legal, and ethical expertise. Composed of the most highly regarded scientists, its mission could be to advise all federal agencies that fund, apply, or regulate research that affects the brain and nervous system and could be used or misused for national security purposes. To give the problems the attention they deserve, perhaps this board would report to the National Security Council and review all the portfolios of those agencies engaged in neuroscience research. That would admittedly be a huge task, so the board might deal only with specific cases referred to it. In either case, such an expert board would perhaps be concerned both with monitoring dual use and helping develop policies on the applications of technologies that are targeted to the brain and nervous system in a national security context. The committee could also advise on the development of countermeasures. Considering the great variety of possible technologies to be covered, the membership would need to have correspondingly diverse expertise, including neuroethics.

A ROLE FOR NEUROETHICS

A government advisory committee on neurosecurity could benefit from including people who spend their lives trying to think clearly about ethical issues in science. Some of the best thinkers in both neuroscience and neuroethics have told me they would be interested in helping security agencies look ahead and develop policies. There are some early examples of members of the bioethics community being called upon to work on these problems by defense agencies and their contractors. Granted, you don't have to be much of a cynic to suspect that the agencies' interest in ethics is all show with no substance, but it seems to me there's more to it than that.

For instance, I am fairly confident that my friend Laurie Zoloth, a Northwestern University professor, is so far the only bioethicist to have been asked to give a lecture to DARPA. That the agency would care about ethics enough to listen to an academic presentation surprised no one more than Zoloth, who is a self-described left-wing peacenik bred in the sixties. Her biography includes dropping out of college to help save the world and working as a community organizer. Today she is an expert on Jewish moral philosophy as well as the ethics of science. A colleague on another project asked her to attend a DARPA retreat in 2002. The meeting included people from veterinary and medical schools from around the country as well as agency staff.

Zoloth gave her talk, with no honorarium, and was so fascinated by the lucid and serious discussion about the ethical theories she presented that she decided to stick around to hear the science lectures, which included presentations on improving the safety and protection of troops, wound healing, ways to better deliver health care in an extreme situation, self-care when injured, and so forth. She was especially struck that the agency officials appreciated that the implications of their efforts to make health care more deliverable when there are no health care professionals and few resources could as easily be applied to people in impoverished countries as to isolated soldiers, though that wasn't their immediate intent. Overall, Zoloth came away impressed, as she told me, by "their intelligence, their optimism and creativity, and their positive reception to thinking about ethics." She didn't expect to like them as she did, and later took some predictable criticism from colleagues who thought she was selling out to the masters of war.

As if one such experience wasn't enough, the DARPA chiefs invited her a second time (again she didn't accept compensation), on this occasion to a conference at their innocuous headquarters in suburban Virginia. Zoloth's topic this time was ethical issues that arise when technologies such as computation, nanotech, and genetics converge, leading to neural implants and worries about mind control. Again, the hundred-plus attendees raised fascinating questions about how we think about science as compared with the way we think about ethics, about cultural differences and their role in morality, even about theories of justice and access to scarce resources. Again, Zoloth found herself impressed. They "obviously weren't in it for the money," she reflected. "They were thoughtful about their power, open to new ideas, and in some ways were a refreshing change from many left-wing academic audiences, more open. They were patriotic, sure, but in a nonjingoistic way I found heartening."

DARPA and its counterparts aren't about to become departments of moral philosophy or Franciscan monasteries. But targeted ethical analysis on specific issues could and should help guide policy on acceptable areas of research. An example of the potential role of ethical analysis in the world of national security neuroscience was a proposed presentation at the first AugCog International Conference in Las Vegas in July 2005, held in conjunction with the Eleventh International Conference on Human-Computer Interaction. This was a sort of expo of applied neuroscience, with themes and topics such as "human performance engineering," "human-computer interaction," "engineering psychology and cognitive ergonomics," "multisensory interfaces," and just about every other bit of jargon that could have made its appearance in an Isaac Asimov novel. Once again, the work being presented was admirable, including numerous papers on developing assistive devices for people with disabilities. The range of work represented was staggering, a potent exhibition of the potential for sophisticated, practical neuroscience to change the way we deal with our world and the way we see ourselves in it.

How interested were the AugCog attendees in the ethical issues associated with their work? It's hard to know, but one piece of information is not encouraging. Two bioethicists, Judy Illes from Stanford and Paul

Wolpe from the University of Pennsylvania, were scheduled to be on a panel called "Exploring Ethics in Augmented Cognition Research," in which they offered to begin a dialogue between AugCog and ethics, two fields that they believe need to understand each other better. Both are exceptionally well qualified. Wolpe is a sociologist who works in both bioethics and psychiatry and is also NASA's bioethics adviser; Illes is recognized as one of the founders of neuroethics and has been asked by several DARPA-funded companies to assist with ethical issues. At the AugCog convention, Wolpe was to discuss the history and ethics of the use of human volunteers in research, Illes the particular issues that arise in neuroscience research. Unfortunately, the panel was canceled after only three people signed up for the session. "We really wanted to get this exchange of views on the ethical questions going," Illes told me, adding that "the apparent lack of interest in ethical concerns by DARPA-funded scientists is a challenge."

Is it pie-in-sky idealism to expect that entrepreneurs who make their living as defense contractors will take a nanosecond to worry about neuroethics? I might have thought so, but then I was introduced to Don DuRousseau, a veteran who owns a small company called Human Bionics LLC. DuRousseau had been talking with the Pentagon about developing a brain-machine interface system that puts together several of the concepts I've talked about in earlier chapters: a handheld device that wirelessly records and analyzes all sorts of biological data in real time for measuring the general health of the brain and body and how these systems respond to target stimuli. Such capabilities would let the computer side of the brain-machine system predict whether a soldier was losing a high level of alertness on a particular task and, if so, redirect his or her attentional resources to the significant environmental elements and improve arousal level. Thus, in addition to simple heart and respiration rate information, twenty-first-century brain-machine technologies will include ongoing assessment of working memory and attentional cognitive systems to improve overall job performance, train individuals and teams faster, and even extend cognitive capacities through external means.

The concept looks impressive, and DuRousseau has some top computer scientists working with him. Of particular interest to me is the fact that he has had ethics advice from the very beginning. His start as a neuroscientist working on basic research questions provided the rigor of the scientific method and exposure to the process of peer review and oversight by coworkers and seniors, he told me. This environment laid the groundwork for the way he views the importance of immediately and clearly establishing the ethical principles on which his company operates, particularly with respect to developing new products that expose aspects of the human side of the brain-machine interface. DuRousseau says that he's always been interested in learning about the brain: the source of our mind and human uniqueness.

> Of all the organs in our bodies, it's our brains that separate us from the apes and from each other, and those thoughts have been the driving influence behind my business. That being said, it places my business at the heart of passionate debate, not only over issues of our origins but also of the technologies now able to peer in at our uniqueness. Without a strong ethical foundation with clear operating policies, my business would fail as soon as it became known to the public that we can run machines with only our brains. No matter what the application, someone will find it unethical, so I have to be able to define the greatest good from my technology and weigh that against the costs of defending the ethics of my business.

CODA: BEYOND MIND WARS

Dual use is a two-way street. In this book I have primarily considered the military applications of neuroscience and brain-related technologies. But cooperation is as much a part of the human condition as conflict and its material manifestation constitutes the bulk of the evolved human brain. As with all successful species, our ancestors survived because they spent more time helping than hurting one another, which made it possible for the human cerebrum to find ways to thrive. We should be able to learn and apply the lessons of the new brain science for peaceful purposes. As the national security implications of neuroscience become more apparent, the pressing need to examine how our brains dispose us to peace as well as war should gain currency. The fields of conflict resolution and peace studies could enrich and be enriched by information from the neurosciences. Future interventions into international and civil conflicts may benefit from greater sophistication about the human brain.

The long-term trajectory of humanity combines a growing capacity for indiscriminate destruction along with vast increases in constructive methods and techniques for solving problems that inhibit human flourishing. Somehow, these seemingly contradictory traits must be neurologically linked. Perhaps, understanding more about this excruciatingly complex system, we can turn ourselves from the wars of the mind to the peace of the soul.