Koepsell, David. "Nanotech Nightmares." *Innovation and Nanotechnology: Converging Technologies and the End of Intellectual Property*. London: Bloomsbury Academic, 2011. 176–195. *Bloomsbury Collections*. Web. 26 Apr. 2023. <<http://dx.doi.org/10.5040/9781849664783.ch-008>>.

Downloaded from Bloomsbury Collections, www.bloomsburycollections.com, 26 April 2023, 09:12 UTC.

Copyright © David Koepsell 2011 2011. You may share this work for non-commercial purposes only, provided you give attribution to the copyright holder and the publisher, and provide a link to the Creative Commons licence.

CHAPTER EIGHT

Nanotech Nightmares¹

What is the regulatory environment that will best encourage the development of nanowares? Clearly, I believe that significant changes need to be made in our IP regimes, or in how we choose to relate to existing regimes, to best achieve the full promise of the technology. But the potential for physical harms must also be fairly considered, and some balance between private and public action should be achieved. We touched briefly already on the visions and likelihood of various nightmare scenarios emerging from converging technologies. These concerns, while likely still remote, are real and implied by past trajectories of science, technology, and industry, as well as by history's noteworthy lapses in human judgment and ethics. Someday, someone will create new means to cause harms, and these harms may threaten either individuals or populations. They may even threaten the species, or all species. Our technologies, while inherently neutral, are often used for ends we neither predict nor wish, and eventually perverted to serve what might well be termed *evil* ends. We should take seriously the duties of scientists and innovators to avoid the unintended but evil potential uses of their work, as well as to confront publicly these possibilities both to educate and to prepare.

Central to the discussion of the dangerous or deadly potentials from converging technologies is the question of whether and to what extent scientists and engineers owe duties to the public at large, or to individual consumers. Moreover, do governments have duties to restrict the development and uses of certain technologies, to regulate them, or simply to provide oversight? Bill Joy's grey-goo scenario stands at one extreme of the dangers posed by true molecular nanotechnology (MNT), but there are more subtle, potentially harmful uses to which nanowares might be directed. A whole new realm of intrusion could be ushered in by miniature machines capable of monitoring us without our notice. An arms race, of the type envisioned by Neal Stephenson in *The Diamond Age*, where nations, corporations, or even individuals constantly attempt to defeat the defenses or offenses of another, could easily fill our environment with competing nanomachines. Or more benign forms of surveillance conducted by invisible machines might threaten

our current notions of privacy. Do scientists and engineers owe us some duty to foresee and prevent these possibilities?

Scientific Duties and Dangerous Technologies

Consider this: You are a nanotech researcher, trying to wed computing with materials. Your dream is to create molecule-sized robots that will do our bidding, constructing items atom-by-atom, capable of user-generated, customized alterations, and fully recyclable. As you near your goal, you realize the full potential of your new creation, which could just as easily disassemble a human being as construct a cup. You consider not only whether and how evil uses of this breakthrough technology might be prevented, but also whether the potential harms outweigh all potential benefits.

Usually, when discussing the moral implications of various technologies and sciences, we take for granted certain presuppositions, such as ‘we can’t put the genie back in the bottle’, and that ethics is the realm of technologists, not scientists, since scientists have a duty to explore all questions, but technologists have no duty to release every technology. Is it conceivable that these presuppositions are erroneous? Do scientists have duties, regarding especially dangerous aspects of nature, *not* to pursue certain fields of research? Do they share responsibility with technologists who eventually release dangerous technologies? Does this responsibility involve moral culpability whether or not there are any harms that result?

These questions and assumptions deserve a second look. They were the focus of a number of thinkers, including scientists such as Einstein and Oppenheimer, at the beginning of the nuclear age when scientists who had been involved in the development of nuclear weaponry came to oppose the tools they had helped develop. It’s an age-old angst, born in *Faust* and *Frankenstein*, involving the inevitable clash of unbounded, unfettered scientific exploration and deadly consequences that sometimes result. Too often, scientists have plodded or lunged along, investigating new means of engineering more destructive technologies, insulated by the concept that science should never be stifled, and that liability for the tools eventually developed because of their investigations rests solely on technologists, engineers, and politicians. But what justifies the disavowal of moral culpability by those in the best position to reveal new and deadly aspects of nature? Is there any moral duty on the part of individual scientists to simply say ‘no’ to investigating those things whose only or best uses are to cause harms?

The Scientific Firewall

It has long been a staple of ethical debate regarding scientific research that science is open and free, and only engineers need worry about the applications to which they put scientific discoveries. The canard goes: science is ethically neutral, and there is, in fact, an ethical duty to investigate all natural phenomena, so therefore, no scientist need stifle his or her own research. The next step of this argument is to assert that while science ought to be utterly unfettered and free, technologists, engineers, and politicians are both practically responsible and morally culpable for the uses to which any scientific discovery is put. This argument works best with 'dual-use' scientific subjects, like bioweapons, nuclear fission, and fusion.² The tremendous possible peaceful uses of thermonuclear technologies argue well for most scientific investigation regarding the underlying sciences. But is this true for all sciences; do they all have 'dual-use' features that insulate scientists from moral culpability when doing basic research?

Consider the recent, real-world example of smallpox. By 1977, a worldwide concerted effort led to the successful eradication of smallpox in the natural world. Its only host is humans, and in the years since its successful eradication, no naturally occurring infection has been documented. This was one of the most successful and heralded scientific and technical enterprises ever. The smallpox virus was virtually extinct, except for some notable stockpiles. The two nations that spearheaded the eradication, coordinated by the World Health Organization (WHO), were the United States and the Soviet Union. Each maintained frozen stockpiles of smallpox samples following the eradication, ostensibly for the purposes of doing scientific research. Then things got out of hand. While WHO and others debated whether the remaining stockpiles ought to be destroyed, some scientists chimed in against the plan. They argued that stockpiles ought to be kept so that further research on smallpox could be done. Some even argued against the eradication of a virus species on moral grounds. The decision to destroy the stockpiles was delayed, and then the stockpiles began to be exploited. There is evidence, including the statements of former Russian president Boris Yeltsin, that during the Cold War, the Soviet Union defied the Biological Weapons treaty and weaponized smallpox, producing it in bulk, and making it more deadly by 'heating' it up, essentially making it less vulnerable to existing vaccinations and anti-viral drugs by exposing it to evolve more robust strains.³ In the process, stores of smallpox apparently

left at least one of the two designated repositories, and now the genie is likely once again out of the bottle.

Once again, in 1999, the world's two custodians of the only known stockpiles of smallpox were on the brink of deciding to destroy the stockpiles (inasmuch as they were believed still to solely exist in the hands of Russian and US scientists) when again some scientists chimed in with a chorus of objections. There were also scientists, some of whom had worked on the original eradication, who argued for the final destruction of smallpox everywhere. In the United States, President Bill Clinton was convinced by those who favored preserving the stockpiles, and the window has now finally closed. The Centers for Disease Control and others working with the US Department of Defense engaged in some controversial studies with smallpox in 2000 and 2001, and successfully created an animal model of the disease, a scientific feat that had hitherto been deemed impossible. This research has since been criticized as being over-hyped, as the animal models that resulted required extravagant forms of exposure before contracting smallpox, making them less-than-ideal subjects for experimentation. The research has further been criticized as being unlikely to lead to any useful discoveries, given that smallpox has been eradicated from the environment and only poses a threat from the current custodians of the virus: the United States and Russia, who could have eliminated it once and for all, but didn't.⁴

In a similar vein, and related to the decision to revitalize US smallpox research, some Australian scientists made quite a stir when they decided to see what would happen if they did some genetic engineering on the mousepox virus. By tinkering with the virus, inserting a mouse gene into the virus itself, they discovered they could defeat any immunity acquired by mice who had been vaccinated, and created a lethal, vaccine-proof mousepox virus with some simple genetic engineering. When US military researchers got wind of this experiment, reported both at a poster-session at a conference and in a journal article, the repercussions for potential mischief with the smallpox virus were obvious.⁵ There are obvious ethical questions that arise with both the Australian mousepox experiment and the US/Russian decision not to destroy the last vestiges of the smallpox virus when the opportunity existed. In each case, to differing extents, one might ask whether the risks of the research justified the potential benefits. Weighing the scientific justification against the potential risks of the research seems inadequate, however, to convey the ethical quandary posed by this and similar research. It is a

quandary posed by research and development of other technologies, notably in the twentieth century, and that was partly responsible for the development of modern principles of research ethics. The question one might reasonably ask is: Do scientists owe a duty to humanity beyond the relentless, unfettered search for natural laws? The verdict, at least in the realm of bioethics, has been established to be affirmative ... there are general ethical duties that outweigh research itself, and that temper behaviors at least when they directly affect human subjects.⁶

The Bioethics Example

It took some terribly visible ethical lapses by Nazi physicians to begin the discussion of codes of ethics governing research on human subjects. The Nuremberg Code was instituted because of the Nuremberg trials, and revelations about the use of concentration camp prisoners for experiments, devoid of pain-relief measures, any semblance of consent (much less informed consent), or any shred of human dignity. The Nuremberg Code served as the founding basis for the evolution of bioethical principles throughout the rest of the twentieth century. Principles such as the right of subjects to provide informed consent before being experimented on, and of being treated with dignity rather than used as mere means to ends, derive from well-known and generally accepted philosophical traditions, but were ignored historically by researchers even outside of Nazi Germany. Well-known examples such as the Tuskegee syphilis study, the Milgram study (both in the United States), and others prompted the development and institutionalization of bioethical principles in both professional codes and federal and state laws.⁷

Simply put, before the 'common rule', the Belmont Report, and similar codes in European nations specifically applied through laws and regulation the principles enumerated in the Nuremberg Code, human subjects continued to suffer in the name of science. We might speculate as to why scientists would fail to apply commonly held ethical principles, such as truth-telling, seeking consent, and preventing foreseeable harms, but motivations are ultimately not the issue. The fact is, it took creating institutions intended to oversee research on human subjects in order to finally begin to systematically prevent such abuses. It is very likely that many of the scientists over the ages who have misused human subjects in the course of their experiments never intended to cause undue harms, or justified any harms by the potential for greater rewards from their study. A prime example is the completely un-consented

to use of a child by Edward Jenner, the inventor of the smallpox vaccine. Jenner's work involved deliberately trying to infect a child, without adequate controls, animal studies, or consent. Fortunately, his hypothesis turned out to be accurate, and the use of cowpox to vaccinate the child prevented his death. Jenner's work saved countless millions, but his ethics was clearly wanting. Such a study today would not have been possible given that animal trials would be useless without a proper animal model for smallpox.⁸

It is likely that Jenner and other scientists similarly situated never meant specific harm, or at least that they justified the potential for harm to a particular subject by the potential for life-saving new treatments for many. What cases like this illustrate, however, is the fact that science has at times proceeded as though ethical concerns were an afterthought, or even completely tangential to the scientific enterprise. Even after the Nuremberg trials, scientists fell into the trap of weighing more heavily the value of potential benefits to be gained from research over individual duties of upholding dignity, providing informed consent, justice, and beneficence. Now, Institutional Review Boards and Ethics Committees provide oversight where human subjects are used in research, and help to give guidance to scientists who might make similar errors. But not all research involves direct use of human subjects. Rather, some research has only potential, future impact on populations, ecosystems, or even humanity as whole. No regulatory body requires vetting of that sort of research.

The example of the development of bioethical principles and institutions intended to apply them suggests that sometimes scientists do not self-regulate when it comes to ethical behaviors. It suggests nothing about motivations, however. It seems likely that ethical lapses are generally caused by lack of introspection, rather than maleficence. This lack of introspection may be exacerbated by the prevailing attitude among philosophers of science and scientists themselves, namely, scientific investigations ought to proceed without limit, and only politicians, technologists, and engineers are to blame for the unethical *applications* of scientific discoveries through technologies. But as is clear in the example of bioethics, and numerous documented examples of ethical lapses by researchers conducting experiments on human subjects, sometimes bad things are done in the name of science itself, well before the application phase of a new technology.

There is a vast and growing literature addressing the ethics of so-called 'dual-use' scientific research, often in the context of the smallpox example,

and other bio-security or bio-weapons agents and research. The frame of this discussion has largely included the perpetuation of the notion that 'legitimate' research often has illegitimate uses or consequences. Some have argued from this context that scientists must take upon themselves certain ethical duties and responsibilities, while others have maintained the standard argument that moral culpability lies with those *applying* the research, not those doing it.⁹ Looking at the development of bioethics as a field, and considering its institutions and principles, one might ask whether the Belmont Report needs some updating. With a little tweak, we might well fashion a set of principles just as elegant and concise as those enumerated in the Belmont Report, aimed not just at scientists doing research on human subjects, but rather at those whose research impacts *humanity* as a whole. Let's consider this possibility, see how a modified set of Belmont Principles might be applied more generally to all scientific research, and then take up the question of how institutions might then be created that could implement these principles. The discussion is framed with examples, both real-world and hypothetical, and considers also the extent to which some scientific research might never be considered 'dual-use'.

Respect, Beneficence, and Justice

These three essential principles of biomedical ethics frame all reviews of proposed biomedical research involving human subjects. Although based on long-debated principles of ethics in general, and owing much to standard notions of both utilitarian and deontological ethics theories, the Belmont Principles are thoroughly pragmatic, and derived clearly from the most prominent ethical lapses that stoked the report's authorship. They include:

1. respect for persons,
2. beneficence, and
3. justice.

The principle of respect for persons is akin to the Kantian notion that people may not use each other as means to ends, but must treat each other with equal dignity – as ends in themselves. In the Tuskegee study and other notable lapses of scientific ethics, human subjects were used as means to ends, just as other scientific tools might be, without regard for equal dignity of the subject. The principle of beneficence simply requires that the research on human subjects be conducted with good intentions. It ought to be pursued

not merely for the sake of scientific curiosity, but rather to cure some ill, to correct some harm, or otherwise benefit humanity. Finally, the principle of justice requires that populations or individuals who are vulnerable (like children or historically mistreated minorities) must be specifically protected in the course of research.

Debate about the merits and application of these principles continues, but they have also become institutionalized in the form of guiding principles used by governmentally created review bodies that now oversee all research on human subjects in most of the developed world. Despite the philosophical status of debates regarding the Belmont Principles, they are in effect already enacted, accepted as part of the background of all research on human subjects, and devised as a hurdle that must be overcome when proposing new experiments involving human subjects. While ethical lapses still occur, as we have seen with such widely disseminated stories as that of Jesse Gelsinger,¹⁰ we can now gauge the conduct (or misconduct) of researchers involved in these lapses, and educate researchers about how to avoid them in the future. In other contexts, both laws and moral codes do not prevent every harm, but provide contexts for judgments when harms occur. Laws and moral codes still serve to educate, and when agreed upon generally, frame our moral debates over particular acts, intentions, and consequences.

These principles are not unique to the realm of bioethics. They frame many of our everyday acts and intentions, and serve as the basis for both moral and social education and regulation in our everyday lives. Despite their expression and adoption in the realm of 'bioethics', what prevents their application to other fields of investigation? Perhaps it is because the sciences outside biomedicine have had fewer public and noteworthy instances where research has caused visible harms. The deaths or injuries of human subjects used in research typically cause public outrage and provoke action. Research which has no such immediate consequence is unlikely to get that sort of notice. But does this mean that the Belmont Principles are not more generally applicable? If we believe that these principles have no application outside of biomedical research involving human subjects, then we must justify some moral horizon for intentions and acts of scientists. In other words, we would have to justify ignoring the potential for misuse or harms to those not immediately within the control of the researcher, even where those harms might well outweigh or outnumber harms posed to potential human subjects.

To put this into context, let's consider a fictional technology at the center of Kurt Vonnegut's well-known breakthrough novel, *Cat's Cradle*. In this novel, a fictional scientist named Felix Hoenikker discovers a permutation of water that is solid at room temperature. He hides his discovery before he dies, but the secret remains in the possession of his children. Ice-nine possesses the ability to turn any body of water solid given that a single molecule of it will 'teach' all other molecules next to it to become ice-nine, creating a chain reaction that freezes any body of water with which it makes contact. It does the same to any water in an organism's body if ingested. The fictional ice-nine is clearly a terrifying scientific discovery. Vonnegut based the character Hoenikker on the Nobel Prize-winning Irving Langmuir, whom Vonnegut knew through his brother's association with Langmuir at General Electric. Of Langmuir, Vonnegut said: '[he] was absolutely indifferent to the uses that might be made of the truths he dug out of the rock and handed out to whomever was around. But any truth he found was beautiful in its own right, and he didn't give a damn who got it next'.¹¹

In the book, ice-nine inevitably gets released into the environment, essentially bringing about the end of the world, all life on it, and all but a handful of people who manage to survive. The research on and discovery of ice-nine would never have fallen under the purview of bioethical principles as enunciated in the Belmont Report. While ice-nine is fictional, smallpox is not and it poses many of the same questions, real-world threats, and ethical challenges as that posed by Vonnegut's book. Is the beauty and truth of science justification enough to investigate even the deadliest potential technologies or discoveries? Are there ethical principles that bind individual scientists in the absence of regulatory institutions? Can the Belmont Principles be extended to scientists doing research only *indirectly* involving human subjects, where the potential effects of an avenue of study impact *humanity as a whole*?

Extending the Moral Horizon

Most arguments concerning the morality of certain types of basic research focus on issues of dual-use and unintended consequences. These arguments concentrate on the distinction between 'legitimate' scientific investigation versus unethical uses of the research. As discussed above, this presupposes that scientific research is always in a different moral position than the application of that research. What justifies this assumption? Do scientists enjoy a unique position occupied by no other fields or professions? Let's

examine some reasons why they might before considering whether scientific inquiry itself, prior to its application through a particular technology, may ever confer moral culpability.

Some might contend that scientific inquiry alone can never confer moral culpability because inquiry is personal, a matter of conscience, and subject to no restrictions at all. Limiting inquiry in one realm might lead us on a slippery slope of censorship, thought-police, and other clearly unsavory interference with free thought. We don't want regulators to prevent scientists from doing legitimate research, looking over shoulders to police scientific investigations, and preventing the acquisition of knowledge about nature. Indeed, governmental interference with scientists' research has provoked the wrath of both scientists and the public, especially when done in the name of particular ideologies.¹² Let's take it as a given that this sort of regulation is tricky at best, Orwellian at worst. But just because we don't want government or regulators overseeing the actions of an individual, doesn't mean that that person's actions, or even intentions, are free from moral scrutiny. We often and comfortably make moral judgments about conduct and intentions that have no *direct* effect on others, even when we don't tolerate or desire any institutional regulation. Yet there are still strong arguments supporting the notion of unfettered scientific inquiry, devoid both of governmental and self-regulation.

Science doesn't kill people; people with *technologies* kill people. Of course, this is a perversion of the US National Rifle Association motto: 'guns don't kill people; people kill people'. There is some sense to this. Artifacts like crossbows, rifles, and nuclear weapons cannot be morally culpable, only people can be. In the name of greater personal freedom, of both conscience and property, governments ought not to restrict inquiry into, or ownership of, dangerous items. The law, codes of ethics, and public and private morality are well-equipped to deal with unethical uses of artifacts, so the principle of maximal freedom requires that we allow not only inquiry into, but also possession of knives, rifles, and nuclear weapons (at least for certain nations). But this is not quite the case in practice, and we do tolerate restrictions on owning certain artifacts. Thus, in the United States, even the most ardent gun aficionado cannot legally own a fully automatic weapon, to say nothing of a tactical nuclear bomb. Moreover, we do not restrict thinking about, and inquiring to an extent, laws of nature generally. Indeed, many of us would consider it immoral to restrict such thought or inquiry as an intrusion into

matters of personal conscience. But does this necessarily imply that while the freedom of personal conscience enables us to think about and inquire into all the universe's natural laws, taking any and all further steps must escape moral judgment?

Take, for instance, the problem of child pornography. Do we hold a pedophile morally guiltless just because, while he might have amassed a collection of pedophilic literature or cartoons, he or she never actually abused a child or contributed to the abuse of a child? Notions of free speech and conscience might protect that behavior, but we are willing to judge certain further positive acts relating to pedophilia as not only morally blameworthy, but worthy also of legal culpability. Intentions matter, even when intention alone is not enough to spur public regulation. Stated intentions matter more, and even when they do not rise to the level of legal culpability, they may spark appropriate moral indignation or outrage. And finally, positive acts based on intention matter even more, and can provoke appropriate moral indignation, outrage, or public recrimination. The pedophile who possesses photos, even while he or she might not have directly contributed to a harm, has indirectly done so and our moral outrage and legal repercussions grow accordingly.

The case of the pedophile might make us reconsider the notion that, while all thought and conscience should be totally free of external restriction, both are nonetheless immune to *moral judgment*. Similar cases may be made about individuals in both their personal and professional capacities who hold intentions, and even take positive actions without direct consequences or harms, yet that invoke some moral culpability. Do we hold the businessperson who thinks about the social or personal consequences of his or her actions in the same regard as one who does not, even where no real difference accrues to customers, colleagues, or society?

I argue that when considering the ethics of scientists, we must not only look at regulations, laws, and codes used to review or punish their actions, but also consider intentions and motivations with an eye toward education. Moral training of scientists, as with other professionals, presupposes not only that we wish to keep them from breaking laws or running afoul of professional codes of conduct, but also that we wish to help develop moral insight that can guide behaviors.¹³ Take an example from another profession whose members affect peoples' lives daily, with sometime dire consequences. Even where an attorney, for instance, injures no one by his lies, the fact of the lie alone ought to concern us. Both in their individual

and professional capacities, people who lie are not to be trusted and deserve our moral judgment. In professions like engineering, science, medicine, and law, moreover, the consequences of actions taken with ill-intentions matter much more to clients, colleagues, patients, and society as a whole simply because the potential for harms is so great.

We should take account of a broader moral horizon when considering the ethics and morality of scientists, just as we do with other professionals, and ask whether and when intentions and positive actions on those intentions trigger an individual duty to refrain, and subsequent moral judgment by others, even where the law or regulations ought not to be invoked. All of which brings us back to the scenarios presented at the outset, and the problem of ice-nine as described by Kurt Vonnegut. Do the principles of beneficence, dignity, or justice provoke any ethical duties for scientists while inquiring into natural laws? Are these duties, if any, different in so-called 'dual-use' cases than for instance in the case of nightmarish scenarios like ice-nine?

Smallpox, Ice-nine, and Nanowares

Almost anything can be considered 'dual-use' if we want to be technical. A nuclear bomb could be used to level a city, or to create a canal. Smallpox research could be used to develop new cures, new therapies, antiviral medications, or new biological weapons. Even ice-nine could serve a dual use, providing ice to skate upon in the middle of summer, or destroying the entire ecosystem. For that matter, the most seemingly benign inventions could, given sufficient creativity, be put to questionable or immoral uses. Printing presses can produce great works of art, or hateful screeds. The dual-use debate, then, may be a bit of a red herring. We are concerned with the ethical implications of certain types of scientific research, and the capacity for a certain discovery or technology's dual-use is not what matters. Instead, we should ask under what conditions a scientist ought to refrain from investigating some aspect of nature, and under what conditions he or she ought to disseminate certain knowledge, regardless of whether the science in question has both a beneficial and harmful use. Let's reconsider the issue of smallpox and its near-eradication.

The global public-health initiative to eradicate smallpox was nearly successful. Its final success was abandoned, and now, despite the fact that smallpox does not exist 'in nature' it still exists, and poses a real threat to humanity. That need not have been the case. Because smallpox has no other

vectors for its survival apart from human hosts, when it was finally eradicated from all human hosts its extinction could have been guaranteed. But for the fact that the United States and Soviet Union insisted on maintaining stockpiles of the virus, we would not need now to worry about the use by rogue states or terrorists of stolen quantities of smallpox. But for the efforts of the United States and the Soviet Union to 'study' the use of smallpox in bio-warfare, and the mass production and subsequent loss of control over the remaining stockpiles of smallpox virus under Soviet science, smallpox would be but a distant memory of nature's capacity for horror and destruction. Scientists cannot be held immune from the moral implications of having preserved stockpiles of the virus. Some made arguments based upon the value to science posed by continuing study of the nearly extinct virus. Their arguments won the day, even if there may have been ulterior motives on the part of the two sponsoring governments maintaining the last known samples. Do any ethical principles mitigate against either the active encouragement of or complicity in the decision to retain the last remaining smallpox samples?

Let's consider first the Belmont Principles. As it turns out, one of the big obstacles to conducting any legitimate science with smallpox is that it has no animal vectors. The Nuremberg Code, and subsequent codes of biomedical ethics, requires that research on human subject be preceded by animal research. To do useful, beneficial research using smallpox would require a human subject, and no researcher could ethically employ human subjects in smallpox research. Not only would the danger be too great, but without first doing animal research, no human subject research could be approved. In the last-ditch effort to save the smallpox stockpiles in the United States in 1999, researchers proposed that a useful application of the smallpox samples was in attempting to find an animal model for the disease. Toward this end, researchers exposed monkeys to massive doses of smallpox until they finally sickened the subjects. Nearly every monkey exposed died without developing human-like symptoms of the disease. But a couple monkeys developed responses similar to human smallpox. This was written of as a triumph in smallpox research, and for some has justified the maintenance of the smallpox stockpile. Finally, a potentially useful animal model of smallpox infection may exist which now justifies maintaining the stockpiles so that further research can be done. And all of this is further justified by the very real potential that smallpox, while no longer a natural threat, is a threatened potential agent of bio-terrorism.¹⁴ In this context, what are the implications

of the decision to preserve smallpox, considering the principles of *respect for persons*, *beneficence*, and *justice*? Does an extended moral horizon alter our view?

If we consider that the subjects of the smallpox investigations (conducted in part to justify continuing to maintain smallpox stockpiles) include not just the monkeys that were infected and ultimately died, but also humanity as a whole, did this experiment satisfy the Belmont Principles? It would arguably meet these principles *if* smallpox remained a natural threat. The dignity of individual humans was not infringed. No individual was treated as less than autonomous or deserving of equal dignity. Moreover, if smallpox were still a natural threat, then presumably all experiments would be conducted with the goal of treatment or, as was the case before 1979, eradication. And finally, the principle of justice is satisfied as long as no vulnerable populations were treated less than equally in the course of the experiment. But if we consider the implications of the experiments in the context of a disease that could historically have been eradicated completely, then we can be more critical of the intentions of the scientists and their decisions to take part in the research.

Let's imagine, since smallpox had been eradicated from the natural environment and only posed a threat from intentionally maintained stockpiles held by humans, that smallpox and ice-nine pose nearly identical risks, and are similar technologies. Ice-nine, like smallpox, posed no natural risk in Vonnegut's book, but only posed a risk as a human-devised technology. The dual-use argument that might justify experimenting with ice-nine breaks down in light of its artificial nature. Moreover, the potentially catastrophic results of an accident involving ice-nine (namely, the total destruction of the biosphere) argue in favor of a positive duty not to investigate it beyond mere surmise or theory. Neither beneficence nor justice warrant investigating ice-nine. We might argue that beneficence argues in favor of investigating smallpox because we worry about terrorist uses of it and need to devise treatments. All of which is recursively self-satisfying, because we would not have had to worry about this had scientists done the right thing to begin with, and supported its ultimate destruction. In the world of *Cat's Cradle*, we could similarly argue in favor of ethically pursuing ice-nine research *only* in a post-ice-nine-apocalypse environment.

An argument that is often used to justify these sorts of scientific inquiries is that 'someone will devise the technologies, and employ them

harmfully – eventually. Thus, we should investigate these things first (because we have good intentions)’. Of course, this reasoning justifies investigating any and all science and technologies, no matter how potentially destructive or threatening to humanity or the environment. But it presupposes (a) that the investigators doing the work have good intentions, (b) that the technology or discovery would eventually be carried out by others, and (c) that once discovered or applied, it can be contained. Let’s assume that, in fact, ice-nine, or smallpox for that matter, will come into the hands of individuals or groups with less-than-good intentions. Will discovering it or investigating it now do any good? In the case of ice-nine, clearly the answer is no. In the case of smallpox, beneficence would argue for the research if for some reason we believed that smallpox could not be contained with existing technologies. If, for instance, we believe that the Australian ‘mousepox trick’ could be applied to smallpox, then devising ways to treat genetically altered mousepox might be an act of beneficence. But without an animal model for similarly altered smallpox, we’d need first to try the ‘trick’ on smallpox. Again, we have a recursive, self-fulfilling justification to pursue any and all research, including on any devastating, horrific, or deadly technology one can think of. Moreover, there’s always reason to question whether one’s own motivations will always be pure, or that a technology will always remain in one’s control or contained.

The ‘Eventual’ Fallacy

The ‘eventual’ fallacy justifies any investigation, and scientific inquiry, no matter the potential consequences. It fails if we broaden the moral horizon offered by the Belmont Principles to include *humanity as a whole* when we are considering sciences and technologies posing no natural threat. Implicit in bioethical principles is some utilitarian calculus. Science proceeds not in a vacuum, but as a socially devised institution. It is conducted by professionals, with funding from mostly public sources, and with relative freedom under the auspices of mostly academic environments. As a largely public institution, and as the beneficiaries of the public trust and wealth, scientists must consider the consequences of their inquiries. They are not lone, mad scientists, hunched over Frankenstein apparatus in private attics. Nor is worrying about the possible existence of Dr. Frankenstein sufficient to warrant all inquiries. Rather, science must be free to inquire into any and all of nature’s mysteries, but scientists must also be aware of

being beholden to a world at large, bound by concerns about consequences of their research, and ultimately dependent upon public support for their institutions.

The ‘eventual’ argument makes sense when the risks posed by investigating a deadly thing are outweighed by the likelihood of that deadly thing being discovered and used by others combined with the potential of a scientific investigation developing a plausible protection of the public at large. So, roughly

R = risk,

L = likelihood of independent discovery and use, and

P = potential benefit from scientific investigation *now*.

If $L + P > R$, then a scientist can make a moral case for pursuing an investigation into something posing a large, general risk. Otherwise, there is simply no moral justification for further inquiry. Taking ice-nine as an example

$R = 100$ (near-likelihood of worldwide catastrophe if released into the environment),

$L = 99$ (being generous, on a long enough timeline, this number grows to 100), and

$P = 0$ (there’s no ‘cure’ for ice-nine).

Or taking smallpox research (now, as opposed to before the eradication)

$R = 90$ (smallpox could escape and cause enormous human devastation),

$L = 70$ (there’s a chance that Russian stockpiles have already made their way into others’ hands), and

$P = 10$ (we already have smallpox vaccines that work well, but maybe we can develop vaccines for other strains or genetically modified versions).

Note that the value of P changes as the likelihood of potential independent discovery changes because of the temporal *caveat*. Thus, it is right to inquire into the state of scientific knowledge elsewhere. However, this imposes an additional burden not to increase the value of L by *disseminating* knowledge that leads others to dangerous knowledge where there is no supervening imperative due to potential benefits from the knowledge.

The ‘eventual’ argument changes over time, and depending upon actual conditions in the world. If, for instance, we know that some rogue state or terrorist group has been experimenting with smallpox, then the calculus

changes. It changes even more if we can identify the nature of those experiments and thus target scientific inquiry to a specific threat. But a generalized threat posed by the potential of someone acquiring knowledge or technology somewhere at some time means that this calculus requires scientific caution. For sufficiently deadly inquiries or applications, scientists should perceive a duty to refrain at least from disseminating certain types of knowledge, if not necessarily from theoretical inquiry alone (unless that inquiry may reasonably be justified by the above calculus). The 'eventual' fallacy is committed by simple recourse to the truism that over an infinite time-span, all natural truths will be discovered, and all potential uses and misuses of technology will be developed, so present research on any science or application of technology is morally justified.

Implications for Institutions

Unlike the Belmont Principles, which could be used to guide the development of regulatory institutions, the expanded ethical horizon I have argued for above requires individual responsibility on the part of scientists. The calculus proposed must be employed by scientists before they ever get to the point of disseminating their ideas. It is a personal, moral responsibility that must be cultivated. Nonetheless, encouraging the development and adoption of these principles, and adopting the notion of a broad horizon of scientific responsibility (encompassing not just individual human subjects, but also responsibility toward *humanity in general*), can best be encouraged through new institutions. Legal and regulatory bodies ought to devise these institutions both within and among sovereigns. Professional organizations as well ought to embrace and adopt ethical training of their members, understanding that scientists are citizens of broader groups whose funding and support they require. Education in principles, not just of scientific integrity, but also social responsibility, ought to be developed and embraced. Currently, scientific integrity and ethics are taught only in the briefest and most superficial manners, and are not generally necessary for any scientist not doing research on human subjects. But in light of the potential for sciences and their technologies to be used for harm, and given the scale of some of these potential harms, more general education in science and morality should be required. This is especially true where the potential impact of a particular science is great, as with nanotechnology, genetic engineering, and similar technologies.¹⁵

As discussed above, scientists are generally beholden to public funding, at least to some extent, and just as many governments now require some minimum training in the core bioethical principles of Belmont and its offspring, so too could grant funding in technologies like nanotech and genomics depend on some minimum education in ethics.

Besides education, the principles and proposed calculus of risks, harms, and benefits could be used in *post hoc* analyses to determine culpability where scientists release dangerous sciences or technologies which actually cause harms. Just as medical doctors were culpable in the Nuremberg trials, so too might future scientists be morally and legally culpable for the apocalyptic (or even slightly less so) repercussions of their negligence or recklessness. Of course, *mens rea* must be considered, but merely citing the ‘eventual’ fallacy will not suffice to defend all scientific inquiry and its resultant dissemination, either through publication or technology. Scientists must not only have a sense that they are morally culpable for the uses of their discoveries where they understand the risk–harm–likelihood calculus, but they must also be liable to be held culpable where harms result from their acts, and where they possess a culpable *mens rea*.

Just as governments take it upon themselves to fund and advance research and development, both out of scientific curiosity and as a way to grow economically, so should they adopt the responsibility to educate scientists to be better citizens. As taxpayers provide for investigations into nature’s truths, sometimes with no potential for economic benefit, they must also be considered as beneficiaries or targets of the fruits of scientific inquiry. An expansion of the Belmont Principles might include recognizing *we are all human subjects of certain inquiries*. Where discoveries possess the potential of great harms, environmental catastrophes, mass extinctions, or worse, the collapse of an entire biosphere (as with ice-nine, or ‘grey goo’), scientists must take it upon themselves to measure their aesthetic appreciation of truths in themselves with gravity of worst-case repercussions. Institutions and regulatory bodies must encourage this, and provide guidance in the form of practical moral education of all scientists, not just in medicine and bioethics, but for all fields of inquiry.

Teaching ethical principles to scientists need not stifle research. Nor does it imply that scientists must watch their thoughts. They need not restrict their thoughts, but they ought to guide them. Minds should be free to explore all possibilities, but the context for inquiry must always be considered to

encompass something broader than just the institutions of science. Where one realizes grave or catastrophic implications for a particular path of inquiry, one *does* owe a duty to those on whose behalf they are musing, and who would inevitably become the target of resulting catastrophic technologies. Just as any of us may privately muse about acts of horror or violence, we assume greater duties as we begin to discuss, plan, or execute those acts. The same must be true for scientists, as in any other public profession or private life. Respect, beneficence, and justice apply not only to human subjects of particular experiments, but also more generally to humanity as a whole. The result of all this should be better trust of scientists and their profession, and a greater realization on the part of scientists that their work proceeds through mutual trust and appreciation between scientists and the public. In the end, we all should benefit as scientists begin to realize their duties are personally held, and broadly applicable. When faced with the choice to inquire into something whose only or most likely application is harmful or deadly, scientists should have the moral strength, educational background, and public support to refuse in light of ethical principles generally accepted, well considered, and backed by strong public institutions.

Does the Future Need Us?

That all depends. The world is indifferent to our comings and goings, and species have perished before. Carl Sagan mused that perhaps the reason we had not yet been contacted by alien species was due to some law of development by which, when a species attains the capability to reach out physically beyond the stars, it also develops the technical means (via artifacts like nuclear weapons) to destroy itself. Having attained those technical means, he surmised, most species cannot help but drive themselves into another dark age, or become extinct. We would like very much to believe this won't happen to us, and that the better angels of our natures will be summoned in time to avoid our self-extinction. Whether we escape our extinction at the hands of runaway technology will very much depend upon the individual intentions of a great many actors. It cannot depend upon the actions of states, as states have not only been notoriously bad at regulation and containing dangerous technologies (which they often appropriate for their own wartime uses), but *cannot* technically contain the twin technologies of synthetic biology and nanotech, just because of the nature of these technologies.

The question is: Do we need the future? If we do, and if we wish to see the utopian rather than dystopian nanotech future unfold, and all its various promises met, then we need to make conscious efforts to act as though our present decisions matter, and the wrong ones may be disastrous. Curtailing certain research, or self-regulating and community policing, all have roles to play in ensuring that our technology does not harm ourselves and others, or irrevocably alter our environment for the worse. Recognition of general and specific duties to others, and across time, requires behaviors to be moderated despite a generally recognized and accepted duty to science itself. These duties impact both scientists and innovators, as the subject matter of the two overlap, and those with the closest contact with nascent technologies and sciences are best positioned to anticipate and avoid harms.

It is easy and tempting to brush off apocalyptic visions as the musings of apparent neo-Luddites and science fiction writers, but sci-fi dystopias sometimes play out as envisioned. *Fahrenheit 451* can be read as a warning of a future that is nearly here, and in both our popular culture and the political sphere, we are aided by even fictional scenarios in planning for real futures, or avoiding them. Crichton and Joy's nightmarish futures *should* be avoided, and we have clear duties to avoid them where possible. Utopian visions of nanotech similarly serve as beacons for where and how we can direct our intentions, and steer ourselves toward a particular destination.

We may also seek to avoid, as we build or modify institutions by which the future of the technology will be constrained or liberated, certain *social* dystopias. Stephenson's *Diamond Age*, as much as it portrays a technological utopia, also warns of the social dangers of monopoly, suggesting that central control of the technology will lead to new class divisions and both unnecessary and undesirable scarcities. Extending the current IP paradigm over nanotech is not only pragmatically unworkable, ethically unjustified, and economically inefficient, but also anathema to the (small 'l') liberal goals we promote in our politics. In *The Diamond Age*, it is overcome through technological and political revolution of a sort. We can start on the ground floor, ensuring that the technological means to ensure democratic technologies are built into our institutions as well as the technology. We can adopt the perspective that openness now will help us prevent tyranny later. History proves this true.