

Foresight Guidelines for Responsible Nanotechnology Development

Draft Version 6: April, 2006

Neil Jacobstein

© 1999-2006 [Foresight Institute](#) and [IMM](#). All rights reserved. Do not reproduce.

Contents

Nanotechnology will alter our relationship with molecules and matter as profoundly as the computer changed our relationship with bits and information. Research on productive nanosystems will eventually develop programmable, molecular-scale systems that make other useful nanostructured materials and devices. These systems will enable a new manufacturing base that can produce both small and large objects precisely and inexpensively. The Foresight Guidelines are designed to address the potential positive and negative consequences of this new technology base in an open and scientifically accurate matter. The objective is to provide a basis for informed policy decisions by citizens and governments, and guidelines for the responsible development of productive nanotechnology by practitioners and industry.

The Guidelines are presented in the active format of self-assessment scorecards for nanotechnology practitioners, industry organizations, and regulatory agencies. Industry organizations for example can assess and score their own degree of compliance with the Guidelines, in much the same way they do with quality programs. This allows the dialog about nanotechnology safety to move from loose recommendations to self assessment of compliance with an operational set of nanotechnology development guidelines. Precise scoring is not necessary at this point, but the process of regular self assessment is critical. As the dialog progresses, more precise scoring guidelines are likely to evolve.

Version 6 includes consideration of near and long term forms of nanotechnology, and tradeoffs in balancing various risks with the spectrum of nanotechnology benefits addressed by the Foresight Challenges and longer term applications of the technology. This version utilizes a discussion of different types of replicator designs instead of relying only on the general term self-replication, which has many connotations. It distinguishes between specialized manufacturing machinery that utilizes designs with no autonomous replicators and assembler systems that may be general purpose and designed without embedded safety controls. This version of the Guidelines discusses the need for a mix of practitioner, industry, NGO, and government cooperation in enforcing controls. It also addresses some potential health, environmental, and military consequences of the technology and the implications of potential means of circumventing embedded controls. There are several new and reworded guidelines that address thoughtful critiques of the Guidelines that have appeared in books and articles over the past year. As always, feedback is welcome for the next version.

Introduction

All technologies are double-edged swords. Fundamentally new technological capabilities and benefits are accompanied by new risks, and new responsibilities for managing risks appropriately. The acceptance of these responsibilities is not optional. Dealing with the benefits and risks of nanotechnology openly and proactively, neither amplifying nor downplaying them, will be critical to the positive development of the field.

The first version of the Foresight Guidelines was developed during and after a workshop on Molecular Nanotechnology (MNT) Research Policy Guidelines sponsored by the Foresight Institute and the Institute for Molecular Manufacturing (IMM). The workshop was conducted over the February 19-21, 1999, weekend in Monterey, California. Participants included: James Bennett, Greg Burch, K. Eric Drexler, Neil Jacobstein, Tanya Jones, Ralph Merkle, Mark Miller, Ed Niehaus, Pat Parker, Christine Peterson, Glenn Reynolds, and Philippe Van Nederveelde. The Guidelines have been revised many times in the intervening years. The resulting Foresight Guidelines ("the Guidelines") include assumptions, principles, and some specific recommendations intended to provide a basis for the responsible development of nanotechnology.

Continued research and education are needed to create a shared understanding and sufficient knowledge base on the entire set of nanotechnology development and risk management issues that must be addressed. While discussion of guidelines can begin today, the scientific and technical community will continue to evolve its understanding of the issues. The Guidelines have already changed over time to reflect that dynamic understanding and specific feedback from a wider community (see [Background section](#)).

The term nanotechnology refers to several distinct classes of technology, each with its own set of capabilities, potential applications, and risks. The specific terms used for these technologies vary over time; however, it is important to be clear about the fundamental distinctions between them.

Nanoscale Science and Engineering

The nanoscale science and engineering conducted today has been defined as technology with a size range less than 100 nanometers or billionths of a meter. Using that definition only, most of chemistry would qualify. The practitioners in the field would add that nanoscale science and engineering researches and exploits the unique properties of materials in the less than 100 nanometer size range. For example, depending on its specific configuration, carbon in that size range may exhibit extraordinary tensile strength greater than diamond, or act as an electrical conductor, insulator, or semiconductor. This branch of nanotechnology exists today with many research programs throughout the world, and many companies commercializing their applications. These companies may sell purified nanotubes made of tubular lattices of carbon, soccer ball like polyhedra made of 60 atoms of carbon called Buckyballs, dendritic polymers called dendrimers, or other particles in the less than 100 nanometer size range. The applications of these technologies are numerous and significant. They enable fundamentally new types of pharmaceuticals, electronic memory and semiconductor devices, sensors, renewable energy capture and storage systems, water purification devices, super strong fabrics and materials, security and military components, as well as antipollution devices. These applications are already beginning to emerge, and will gather momentum over the coming years.

The risks associated with passive compounds in the less than 100 nanometer size range concern their ability to be inhaled, absorbed through the skin, or to pass through biological compartment barriers such as the blood brain barrier. They thus pose a range of potential health and environmental risks that are associated with their potential toxicity or mutagenicity in their interactions with biological systems. While the range of effects vary, most of the risks may be addressed by advanced industrial hygiene and environmental health practices and techniques that seek to characterize the specific risks, exposure patterns, and control methods and enforce them through a combination of practitioner education, industry self-regulation, monitoring and government regulation. This is an important emerging field in the environmental and health sciences, since most of the existing legislation on environmental, safety, and health risks may cover particulates, but do not take the change in physical and biological properties at the nanoscale into account. It is reasonable to assume that passive nanoscale particle risks, although potentially serious if not addressed, will be characterized and addressed systematically under new versions or extensions to existing occupational, industrial hygiene, environmental, and medical regulations.

Productive Nanosystems

Productive nanosystems are currently a research oriented class of nanotechnology that will produce programmable, molecular-scale systems that make other useful nanostructured materials and devices. These systems may take over a decade to develop and mature. They will be qualitatively different from nanomaterials, particularly regarding regulatory issues. These systems may be used as infrastructure for manufacturing, specifically the ability to build molecularly precise, inexpensive, three-dimensional products of arbitrary size. The most straightforward infrastructure for manufacturing will be built with special purpose molecular fixtures and components that are analogous to macroscale factory components that produce devices that are inherently incapable of replication. These special purpose manufacturing systems will eventually be able to manufacture very large structures by scaling specific components and sub systems.

The simplest, most efficient, and safest approach to productive nanosystems is to make specialized nanoscale tools and put them together in factories big enough to make what is needed. People use simple tools to make more complex tools, from blacksmiths' tools to automated machinery. The convergent assembly architecture developed by Ralph Merkle (1997 *Nanotechnology* **8** 18-22), describes how small parts can be put together to form larger parts, starting with nanoscale blocks and progressing up the hierarchy to macroscopic systems. The machines in this would work like the conveyor belts and assembly robots in a factory, doing similar jobs. If you pulled one of these machines out of the system, it would pose no risk, and be as inert as a light bulb pulled from its socket.

The eventual applications of these special purpose manufacturing systems include the ability to build almost any mechanical device cheaply, and in large quantity. This is why productive nanotechnology manufacturing capabilities will eventually do for our relationship to molecules and matter what the computer did for our relationship with bits and information. The computer enabled an ever expanding number of people to access billions of dollars worth of information. Productive nanotechnology will enable an ever expanding number of people to enjoy significant material wealth, based on carbon feedstock, which currently is in overabundant supply. It will

also enable the technical infrastructure to address effectively many of our most pressing transportation, environmental, medical and global warming issues.

The primary risks of manufacturing enabled by productive nanosystems concern what is manufactured, not the manufacturing infrastructure itself. Special purpose manufacturing systems can be designed to be safe and reliable. They could be made to build a wide range of devices cheaply, in place, and on-demand. These products could include components for large scale buildings, computing, mass transit systems, energy storage, and spacecraft. On the other hand, they could also include tiny new security devices, and large quantities of inexpensive and super strong conventional weapon systems.

Replicators: Autonomous and Non Autonomous

The concern about advanced forms of nanotechnology, after feasibility issues have been addressed, tends to center on the possibility of the technology getting out of control. Popular science fiction describes molecular robots that are autonomous self-replicating machines or autonomous replicators that evolve beyond human control, or can't stop reproducing. For the purpose of our Guidelines, an autonomous replicator is a specific kind of device that both (1) contains a set of materials processing and fabrication mechanisms sufficient to perform the operations necessary build devices like itself and (2) contains a set of design instructions and instruction-interpreting mechanisms sufficient to direct the operations necessary to build a device like itself. All other machines, lacking these exceptional properties, are not autonomous replicators. Productive nanotechnology enabled manufacturing of small or large products does not require any autonomous replicators, either in development or in application. It is also important to distinguish between the special purpose ability to manufacture many copies of a specific product, and the ability of the manufacturing infrastructure to replicate itself.

Nature employs its own biological form of molecular manufacturing to produce organisms. Productive nanosystems can be completely non-biological. They can be designed to function in a narrow and controlled range of physical conditions. Autonomous self-replicating assemblers are not necessary to achieve significant manufacturing capabilities. As Drexler and Phoenix indicated in their Safe Exponential Manufacturing paper (2004 *Nanotechnology* **15** 869-872), developing manufacturing systems that use general purpose replicators able to extract their own energy sources is unnecessary.

It is important to note that molecular nanotechnology not specifically developed for manufacturing could be implemented as *non autonomous replicating systems* that have many layers of security controls and designed-in physical limitations. This class of system could potentially be used under controlled circumstances for nanomedicine, environmental monitoring, and specialized security applications. There are good reasons to believe that when designed and operated by responsible organizations with the appropriate quality control, these non autonomous systems could be made arbitrarily safe to operate. However, a determined and sophisticated group of terrorists or "non state entities" could potentially, with considerable difficulty, specifically engineer systems to become *autonomous replicators* able to proliferate in the natural environment, either as a nuisance, a specifically targeted weapon, or in the worst case, a weapon of mass destruction. Both conventional nanoscale technology and manufacturing enabled by productive nanosystems can be implemented by responsible parties quite effectively without these risks, but as with other technologies, the risk of abuse must be considered seriously. Thus,

in addition to the need for professional ethics and multiple layers of embedded industrial controls, there will also be a need for thoughtful regulation, monitoring, and potentially the development of “immune responses” to external threats.

Embedded Controls

There are many additional dimensions of safety controls that can be engineered into designs for replicators. Robert Freitas and Ralph Merkle (2004, *Kinematic Self-Replicating Machines*, Landes Bioscience, p. 152) have described a multidimensional space of 137 replicator design properties, characterized by structure, function, inputs, and outputs, and the percentage of its own components that a system can fabricate, extract, transport, inspect, warehouse, repair, control, or energize. Each of these design properties represents a potential control point for identification, embedded safety features, and deactivation. The authors support an earlier version of these Guidelines in their book, and agree with its recommendations against developing inherently unsafe replicator designs which permit: 1) surviving mutation or undergoing evolution, and 2) competing with or using biology as a raw materials resource.

Replicator designs may be made arbitrarily safe by employing layers of redundant embedded controls. For example, a broadcast architecture which transmits encrypted manufacturing instructions to a machine without an on-board instruction set. Another example is a vitamin architecture which designs in dependence on an exotic fuel or substance not available in the natural environment. While no technology is absolutely safe and free from the risk of abuse, we will refer to these systems as inherently safe replicator designs. Specifically, they are inherently safe compared to systems that are built without redundant layers of built-in safety controls. It is important to note that molecular manufacturing systems could be made considerably safer and more vulnerable to deactivation than natural bacteria or viruses which evolve rapidly, use biological materials as food, are readily available today, and are easier to hijack for nefarious purposes. Understanding the underlying properties of different forms of non autonomous replicating systems will make it much easier to engineer them with embedded safety and security features, as well as detect and control any attempts to abuse them.

Inherently safe designs may be needed for non autonomous replicators that can monitor the environment and respond rapidly to attacks or natural disasters. The development of nanotechnology is part of a general technological trend towards miniaturization that is proceeding globally at an accelerating pace. Bad actors who might try to abuse the technology could potentially be thwarted if replicator countermeasures or an “immune system” is deployed in advance. These systems would be analogous to those developed by nature for when organisms come under attack by viruses and bacteria, and anti-virus programs developed by humans to protect computing systems. Like our own very successful but imperfect immune system, this type of system could be compromised by sophisticated technologists and programmed to attack itself, as the body does in autoimmune diseases. This does not mean immune defenses are useless. In fact, they have been selectively developed during millions of years of evolution. It does mean that we will need to build redundant layers of security into the design of these systems, and keep improving them rapidly. This may become a “predator-prey” cycle as competitors respond to the latest defense or attack mechanisms. More research and resources are required on the large space of design for safety in non autonomous replicator systems, design for security and rapid response in immune systems, and guidelines concerning their development and use.

Balancing Benefits and Risks

Policy related discussions of nanotechnology require consideration of the economic and environmental benefits of the different types of nanotechnology, as well as the potential problems. Some of the potential benefits of advanced forms of the technology include new lifesaving systems to address massive poverty and hunger, prevent and repair damage from diseases that are not responsive to today's antibiotics or anti-viral drugs, and provide rapid material response to natural disasters such as tsunamis, hurricanes, earthquakes, and volcanoes. Poverty, disease, and natural disasters kill thousands, in some cases millions annually, and the potential to ameliorate their effects significantly should not be relinquished lightly, particularly by those least affected.

The Foresight Nanotechnology Challenges address critical needs that could be met by developing a range of near and long term nanotechnology solutions. They include: 1) meeting global energy needs through more efficient generation, storage and distribution, 2) providing abundant clean water through improved water purification and filtration, 3) increasing health and longevity of human life through medical diagnostics, drug delivery and customized therapy, 4) maximizing the productivity of agriculture through precision farming, targeted pest management and the creation of high yield crops, 5) making powerful information technology available everywhere through reduced cost and higher performance of memory, networks, processors and components, and 6) enabling the development of space resources through improved fuels, as well as smart materials and environments. Meeting any one of these challenges with nanotechnology would have a transformative effect on our quality of life. Thus, the need for new controls should not prevent the responsible development of the field.

There has been considerable confusion about the sources of risk in the development and future deployment of nanotechnology. Most of the risk assessment of current nanoscale technology concerns compatibility of nanomaterials with humans and the environment. Relevant ecological and public health principles must be utilized in the development of any technology, particularly one as fundamental and broad as nanotechnology. Nanoparticles may be absorbed or inhaled if the proper industrial hygiene precautions are not utilized. They may also pass readily through body compartments such as the blood brain barrier. Manufactured diamondoid products may not break down easily in the natural environment. Consumers may not initially have means readily available to recycle them. Thus, total "product lifecycle" considerations about health and environmental effects should be taken into consideration as industry develops new manufacturing techniques based on productive nanosystems.

Since some controls on the different classes of nanotechnology will eventually be put into place, it makes sense for them to be as well informed as possible. Rather than have external controls imposed upon an R&D community that is not addressing potential risks openly, the developing nanotechnology R&D community should adopt appropriate self-imposed controls proactively. They can also participate in policy discussions on external controls that may be formulated in light of current knowledge and the evolving state of the art. The quantity and quality of these additional controls will depend to some extent on the success of voluntary controls.

Successful Precedent

The NIH Guidelines on Recombinant DNA technology are an example of self-regulation taken by the biotechnology community over 30 years ago. While the kind of artificial molecular machines of primary interest for nanotechnology are expected to be very different from the kind of biological systems covered by the NIH Guidelines (just as a 747 is very different from a sparrow, even though both fly), the NIH Guidelines illustrate that advance preparations are possible and can be effective. Those guidelines were so well accepted that the privately funded research community has continued to submit research protocols for juried review, in spite of the fact that it was optional for them to do so. In addition, although the NIH Guidelines have been progressively relaxed since they were first released, they did achieve their intended effect.

Reducing Risks, Improving Opportunities

Industry and government should have the maximum opportunity to develop and commercialize a manufacturing industry based on productive nanosystems designed for safety and reliability. In addition, nanotechnology should be developed in ways that make it possible to distribute the substantial benefits of the technology to the majority of humanity currently desperate to achieve material wealth at any environmental or security cost. Manufacturing based on productive nanosystems will eventually be capable of producing widespread material abundance with significantly less environmental impact than technologies in common use today. Providing technical abundance alone cannot make a country wealthy and secure. This also requires education, and social arrangements that include rule of law, and other features of civil society. However, technological abundance can alleviate many of the conflicts that stem primarily from rivalry over resources. Reducing this specific cause of conflict could make the world considerably more secure than it is today. In addition, the release from bare economic subsistence could enable billions of people to take advantage of the emerging global classroom enabled by the World Wide Web. This education effect could compound the positive security benefits of nanotechnology.

Given the accelerating world wide research on various types of nanotechnology it is highly probable that advanced capabilities will eventually be developed, and thus, it is important for society to consider proactively the range of controls that should be in place to minimize risks and maximize economic benefits. Unlike the proliferation of nuclear technology, which can be partially controlled by limiting access to fissionable materials, once the technology matures, it could utilize carbon, which is readily available. Further, there is no guarantee that limiting research and development on these systems in open and democratic countries would effectively slow their development and deployment elsewhere. China, for example, has an active nanotechnology research program, and it is training large numbers of qualified scientists and engineers in the relevant disciplines for developing advanced forms of the technology. The Guidelines take the position that the safest and most responsible path is to enact reasonable technology-specific controls at the practitioner, industry, and governmental levels, and simultaneously develop the monitoring and counter measures to control potentially rogue or offensive use of the technology.

Regulation and International Treaties

Effective means of restricting the misuse of molecular nanotechnology in the international arena will need to be developed. If weaponized versions of MNT are developed, they may not fall under existing arms-control treaties. Adding productive nanosystems designed for manufacturing

to the list of technologies covered in Chemical, Biological and Nuclear Weapons treaties could be inappropriate because MNT is not a weapon, but a productive technology with broad applications. It is more similar to chemical technology than to chemical weapons, and more similar to biotechnology than to biological weapons.

Adding particular weapons related applications of MNT to the list of technologies covered in Chemical, Biological and Nuclear Weapons treaties may be appropriate in certain cases. It should be remembered, however, that the capabilities of productive nanosystems will be extensions of general manufacturing technology. The military applications of MNT will include the manufacture of high performance aerospace vehicles and precision munitions at low cost. The high value and dual use of MNT for civilian and defense purposes will require making distinctions between the enabling technology and its specific applications, balancing health and economic benefits against security concerns. Since nanotechnology research is now global, the security challenges will be present, with or without the ability to capture the wide variety of benefits.

Overly restrictive treaties or regulatory regimes applied to core MNT technologies could lead to the unintended consequence that only the rule-following nations would be at a competitive disadvantage technologically, economically, and militarily. While most nations, companies, and individuals are likely to adhere to reasonable safety restrictions, guidelines that are viewed as too restrictive will simply be ignored, paradoxically increasing risk. In addition, not all guidelines and laws will be followed, and enforcement is rarely perfect. Non-state actors could become quite significant, particularly when the relevant knowledge and raw materials are available globally. They may well not be signatories to any agreement. While a 100% effective ban could, in theory, avoid the potential risks of certain forms of molecular nanotechnology, a 99.99% effective ban could result in development and deployment by the 0.01% that evaded and ignored the ban. For example, the international Biological Weapons Treaty was being violated on a massive scale even before the ink was dry.

On the other hand, international cooperation on restricting the proliferation and use of atomic weapons has been partially effective in limiting their development and use over the past several decades. In that case, however, the raw materials were not widely available. There are reasonable arguments on both sides of the treaty question. It is wise to avoid an unnecessary nanotechnology arms race, particularly when manufacturing enabled by productive nanosystems could be used to greatly reduce competition for material resources and mutually improve quality of life for the competitors. Treaties may become easier to verify, at least when the verifiers are more advanced than the verified. The “trust, but verify” security concept will increase in importance in a potentially dangerous, but increasingly transparent world. In fact, nanotechnology based sensors are likely to increase multifold the options for transparency.

The Guideline participants as a group have not endorsed any specific means to address MNT security concerns through treaty arrangements. However, as nanotechnology capabilities increase, governmental, NGO, industry cooperative arrangements, and self-funding enforcement mechanisms will need to be implemented. Self policing NGOs and industry groups will play an important role in this, but the public is likely to want some oversight provided by elected representatives. This may reflect distinctions between addressing potential voluntary and localized risks vs. involuntary and more widespread risks, as well as differences in the distribution of benefits. In addition, monitoring defense applications will have to be handled

primarily by government agencies. There will be no easy answers in this arena, only a set of dynamic tradeoffs that will have to be incorporated into flexible organizations and policies that can survive the vagaries of the real world of politics, economics, and accelerating technology.

The international community of nations, industry, and nongovernmental organizations will need to develop effective means of enforcing nanotechnology guidelines and regulations, and responding to misuse. Such means should not restrict the development of peaceful applications of the technology or defensive measures by responsible members of the community. Today there are obvious difficulties in achieving consensus on the definition of “responsible members” and enforcing international agreements on chemical and biological weapons is known to be problematic. However, given the importance of making progress in this area, further research, collaboration, and innovation is encouraged.

Education and Enforcement

The safe development and use of nanotechnology depends, in part, on the good judgment and ethical behavior of the researchers carrying out this work. This is an imperfect, but important first line of defense. The more this is recognized as critically important, the more effective the vast majority of researchers are likely to be in actively preventing unsafe designs or uses of nanotechnology, and in insuring that manufacturing systems have built-in safeguards. The natural and responsible path for the development of productive nanotechnology based manufacturing makes use of no autonomous replicators. However, defense against potential rogue elements who might seek to abuse replicators is a problem not unlike the challenge of controlling the developers of viruses on the Internet. In both cases, a combination of moral and technical education, active industry and government cooperation, inherently safe system designs, legal frameworks, and R&D on secure immune systems for defense may be the best solutions available.

Nanotechnology policy will have to balance risks with benefits, and distinguish between different classes of risks. Molecular manufacturing and nanotechnology are not one technology, but rather a spectrum of technologies, with radically different risk profiles. A substantial R&D program is needed to clarify the nature, magnitude and likelihood of the potential risks, as well as the options available for dealing with them effectively. For example, toxicology analyses relating to nanomaterials have already been identified as an early priority. Nanomaterial safety is a matter that is distinct from the key risks of productive nanotechnology based manufacturing, but both require good industrial hygiene practices.

There are significant risks associated with failing to address the increasingly costly economic, political, environmental, energy, and security problems that the development of productive nanosystems could help resolve. Likewise, there are real costs to restrictive policies that limit nanotechnology innovation by responsible actors and allow rogue entities to move ahead. The Guidelines were not intended to cover every risk or potential abuse of the technology. People may abuse automobile technology, and society has responded by making cars safer to operate, holding drivers accountable for their actions through laws that are enforced, and requiring drivers to pay for automobile insurance. Likewise, industry and governments are held responsible for their use of technologies that have widespread impact.

The Guidelines are intended to cover most of the risks associated with normal development and use of the technology, and to mitigate, as much as possible, the risks associated with potential abuse of the technology. However, most guidelines and embedded control regimes can be circumvented by a sophisticated and determined adversary. Defending against this kind of threat may require an active monitoring and immune system for detection and deactivation.

Informed policy decisions could accelerate the safe development of peaceful and environmentally responsible uses of nanotechnology. This includes capturing the opportunity to develop powerful new approaches to medicine, and energy efficient, zero emission manufacturing and transportation technologies. Informed and thoughtful policies will also increase our security by enabling new types of immune systems to detect and respond to deliberate abuse.

The field of nanotechnology is very broad, like computing, and thus its eventual regulation spans human health and safety (NIH), environmental protection (EPA), and eventually weapon systems (DoD, DHS, CIA). It is important that appropriate distinctions be made between different classes of nanotechnology in the development of regulations, and that effective interagency coordination is done to promote and enforce consensus, and build on existing regulatory standards and monitoring where appropriate.

The self assessment scorecards are based on the notion that the people, organizations, and governments that work in the nanotechnology field should develop and utilize professional guidelines and practices. These guidelines and practices should be grounded in science and technology principles, and knowledge of the interacting environmental, security, ethical, and economic issues relevant to the development of the field. This is based on the notion that professional ethics, "soft law", and cultural norms regarding good practice are at least as effective as "hard law" in preventing unsafe practices, and in helping to ensure that unsafe practices are noticed and acted upon. The use of "soft laws" is a first line of defense, and is not meant to suggest that "hard laws" for safety and health are not useful, and at times appropriate.

Any regulation adopted by researchers, industry or government should provide specific, clear guidelines. Regulators should have specific and clear mandates, providing efficient and fair methods for identifying different classes of hazards and for carrying out inspection and enforcement. There is great value in seeking the least-restrictive necessary legal environment to ensure the safe and secure development of each specific type of nanotechnology. It is important to recognize in this context that some types of nanotechnology will eventually provide the best solutions available for remedying the existing environmental and public health damage resulting from our current, distinctly suboptimal, technology base.

Scorecard 1: Nanotechnology Professional Guidelines

Self Scoring: 0-5, 0 = no compliance, 5 = high compliance

Best Score in this section = 40

1. Nanotechnology developers adopt professional guidelines and ethical practices relevant to the responsible development of both near term and advanced nanotechnology.

2. Nanotechnologists attempt to consider proactively and systematically the environmental and health consequences of their specific technologies. Practitioners recognize that the scope and magnitude of potential problems are reduced to the extent that they consider the range of possible negative consequences, and plan to prevent them, or at least minimize their effects through embedded and redundant control systems.
3. Nanotechnology research and development is conducted with due regard for the principles of environmental science and standard practices of public health, with the understanding that significant changes in physical, chemical, and physiological properties may occur when macroscale materials are developed and utilized on the nanoscale.
4. Nanotechnology products are conceived and developed using total product lifecycle analysis.
5. Productive nanosystem based manufacturing makes use of inherently safe system designs requiring no autonomous replication.
6. When controversy exists concerning the theoretical feasibility or implementation timing of advanced nanotechnologies, such as specialized manufacturing components or scaling techniques, researchers address and clarify the issues rapidly, and attempt to resolve any controversy openly.
7. Any developers who consider the design or development of non autonomous replicators for specific R&D purposes should first explore the potential benefits and risks of alternative approaches actively, in a balanced and rigorous manner.
8. Any use of potentially autonomous replicators is avoided in manufacturing, and only utilized in R&D after institutional review and approval of extensive and redundant control systems.

Scorecard 2: Nanotechnology Industry Guidelines

Self Scoring: 0-5, 0 = no compliance, 5 = high compliance

Best Score in this section = 40

1. Industry self-regulation is practiced proactively, and tailored to the specific risk profile of the nanotechnology under development. Specifically, studies to assess the consequences of new nanotechnology processes, materials, and tools are scoped appropriately, advanced as rapidly as possible, and encompass both benefits and risks with rigor..
3. When molecular manufacturing systems are implemented, they use inherently safe system designs with no autonomous replicators.
4. Any molecular manufacturing device designs specifically limit unplanned distribution and provide traceability and audit trails.

5. Encrypted molecular manufacturing device instruction sets are utilized to discourage misuse.
2. Manufacturing systems are described and classified according to their specific characteristics, particularly with respect to autonomy and safety control systems.
6. Autonomous replicators for R&D are avoided through the use of inherently safe system designs that selectively utilize non autonomous system characteristics, and layers of redundant controls.
7. Replication systems used for R&D are designed to be incapable of autonomous replication in any natural environment. They have multiple system requirements (e.g., for externally supplied information, interventions, environmental conditions, materials, components, or exotic energy sources) that are available only where deliberately provided to enable operation of the machine.
8. Replicator R&D focused on detecting and responding to potential technology threats utilizes redundant embedded safety controls such as time limited operations, encrypted external controls to override internal operations, and anti-mutation protections. For example, the information that specifies their construction is stored and copied in encoded form, and the encoding is such that any error in copying randomizes and thus destroys the decoded information. These systems are continuously improved for security.

Scorecard 3 Government Policy Guidelines

Self Scoring: 0-5, 0 = no compliance, 5 = high compliance

Best Score in this section = 55

1. Regulatory studies and controls distinguish the wide variety of nanotechnologies, and recognize that their different risk profiles require different regulatory policies.
2. Regulations and consensus standards promulgated by researchers, industry, or government provide specific and clear guidelines, and encourage the use of inherently safe system designs for manufacturing and R&D.
3. The government has designated a division within a regulatory entity or a new agency with sufficient resources to ensure nanotechnology standards enforcement and coordination across agencies. It is building upon and when necessary augmenting existing regulatory structure and institutions (eg. for health and safety, environment, defense, and intelligence). Regulators have specific responsibilities and authorities for identifying different classes of hazards, providing development approvals when necessary, and for carrying out inspection and enforcement.

4. Economic incentives are provided for responsible innovation through discounts on liability insurance policies, access to royalties, or consortia membership for molecular manufacturing and development organizations that certify Guidelines compliance. Willingness to provide self-regulation and open access for third party inspection that safeguards proprietary technology is a condition to utilize advanced forms of molecular nanotechnology.
5. Access to special purpose productive nanosystems enabled manufacturing with inherently safe designs and no autonomous replication is unrestricted, unless the special purpose capabilities pose a specific risk.
6. Initiatives are in place to encourage the international community and non-governmental organizations to restrict the deliberate misuse of molecular nanotechnology by improving verification, monitoring, and detection techniques, and making the detection and enforcement of misconduct increasingly probable. Such means should not restrict the development of nanoscale materials, special purpose manufacturing systems, or non autonomous defensive measures utilizing inherently safe designs.
7. Accidental or willful misuse of nanotechnology is further constrained by legal liability and, where appropriate, subject to criminal investigation and prosecution. This should also pertain to those that enable and collaborate on the misuse of the technology.
8. Eventual distribution of advanced molecular nanotechnology capability is restricted, whenever possible, to responsible actors that have agreed to practice these Guidelines, and permit verification. No such restriction need apply to special-purpose, molecular machine systems with no autonomous replicators and inherently safe designs, or to the end products of molecular manufacturing that satisfy the Guidelines.
9. Governments, companies, and individuals who fail to follow reasonable principles and guidelines for development and dissemination of MNT are placed at a substantial competitive disadvantage with respect to access to companies, collaborative NGO organizations, R&D funding, plans, designs, software, hardware, and cooperative market relationships.
10. Incentives are in place to encourage industry, government, and NGO developers to collaborate on continuous improvement and use of best practices in nanotechnology and risk management, including the theory, mechanisms, and experimental designs for increasingly safe manufacturing, as well as effective monitoring and control systems.
11. Regulatory entities sponsor research on increasing the accuracy and fidelity of environmental and health models used for nanotechnology risk assessment and management, as well as the theory, mechanisms, and experimental designs for built-in safeguards and defensive nanodevice immune systems.

Background

The idea of guidelines for the safe development of nanotechnology has been discussed within the Foresight community for over a decade. It is inevitable that any guidelines put forth will be further discussed and perhaps substantively changed; but the dialog on specific proposals must

begin somewhere. This latest version of the Foresight Guidelines represents another step in an ongoing discussion.

In spite of the diversity of briefing materials and views represented at the initial Monterey workshop in February of 1999, the participants managed to discuss the technical and policy issues with both intensity and civility. While any one participant might have preferred more or less emphasis on a particular issue, the group was able to converge on a common set of draft guidelines for the development of nanotechnology.

The group agreed to review the Guidelines among themselves, discuss them in wider Foresight meetings during 1999, and then release them on the WWW for review by the larger community. The goal was to get the Guidelines endorsed and adopted by organizations sponsoring nanotechnology research and development projects, and to inspire effective self-regulation wherever necessary and possible. Explicit discussion of long term risks and potential regulation of industry made acceptance difficult.

Another goal of the Workshop members was to educate nanotechnology researchers about the potential benefits and risks of the technology. The long-term goal was to eventually produce a dialog and set of Guidelines that would be useful to policy makers, the public, and the MNT research and development community. We believe that this has happened.

The Foresight Guidelines are intended as a living document, subject to modification and revision. Early drafts have been reviewed and revised several times since the Monterey workshop, including during Foresight/IMM sponsored discussions led by Neil Jacobstein in May and November of 1999. They were also provided in the attachments to Ralph Merkle's June 1999 Congressional testimony on MNT, and referenced in Neil Jacobstein's presentation on Nanotechnology and Molecular Manufacturing: Opportunities and Risks at Stanford University's Colloquium for Doug Engelbart in January of 2000. The Workshop participants debated whether the Guidelines were sufficiently developed for widespread publication, when Bill Joy's article: "Why the Future Doesn't Need Us" was published in the April 2000 issue of *Wired Magazine*. This article raised public awareness of the potential dangers of self-replicating technologies, including nanotechnology.

Since that time, the Guidelines were reviewed critically by Robert Freitas, and revised by Ralph Merkle and Neil Jacobstein. Version 3.6 of the Guidelines was discussed in a May 2000 Foresight workshop session led by Neil Jacobstein. Bill Joy was invited to participate in this discussion. He made several constructive suggestions, including one that outlined a guideline on closing the economic incentives loop via an insurance policy requirement for developers. Jacobstein incorporated the feedback from this and subsequent discussions into [version 3.7](#) of the Guidelines, and they were then published for open review on the web.

Neil Jacobstein rewrote the Guidelines as a form of self assessment scorecards for version 4.0, based on the observation that this kind of self assessment is becoming a standard part of quality and six sigma programs in industry and government. He combined and added some new guidelines, including a guideline based on a paper by Eric Drexler and Chris Phoenix in the *Journal of Nanotechnology* on "Safe Exponential Manufacturing". This paper made the case for nanotechnology enabled manufacturing using a hierarchy of machine tools, without the need for general purpose self-replicating assemblers. Glenn Reynolds edited the draft and provided an

analysis by his law students on current treaties and the fact that weaponized MNT might not be covered by them. Eric Drexler also reviewed the draft and made additional editorial suggestions.

Jacobstein presented version 4 of the Guidelines at the Foresight Advanced Molecular Nanotechnology Conference in Washington DC in September of 2004, and discussed revised versions at a National Academy of Sciences workshop on the Feasibility of Molecular Manufacturing held in February of 2005, and an Aspen Institute Seminar on Future Perspectives in July 2005. He wrote version 5.0 with feedback from these meetings and specific comments on the Guidelines or the risks they address from Robert Freitas, Marc Gubrud, Eric Drexler, Ray Kurzweil, Bill Joy, Martine Rothblatt, Larry Millstein, Max Moore, Doug Mulhall, Christine Peterson, David Forrest, Chris Phoenix, and Glenn Reynolds. Robert Freitas and David Forrest provided detailed reviews of this draft. Jacobstein produced and presented version 5 of the Guidelines for the October 2005 Foresight Advanced Molecular Nanotechnology Conference in San Francisco. He produced a substantially revised version 6.0 based on feedback from John Bashinski, Christine Peterson, David Forrest and others. David Forrest referenced the Guidelines in his presentation to the Roundtable Discussion on Nanotechnology Regulation of the Senate Subcommittee on Environment and Public Works, April 6, 2006. The Guidelines represent a complex set of tradeoffs between competing concerns, and it must satisfy the needs of more than one special interest community. Thus, it is likely that few of its participants agree with all of it, though most would acknowledge its value. We acknowledge their value in improving the Guidelines.

Version 6 of the Guidelines will be available after a draft review period at the Foresight web URL: <http://www.foresight.org/guidelines/>. We encourage your ideas and constructive criticism about how to improve the Guidelines.

Eventually, the Guidelines need to become sufficiently specific that they can form the basis for a legally enforceable framework within which nanotechnology development can be safely pursued. Future versions of the Guidelines or legislation inspired by them might eventually be enforced via a variety of means, possibly including lab certifications, randomized open inspections, professional society guidelines and peer pressure, insurance requirements and policies, stiff legal and economic penalties for violations, and other sanctions. Enforcement will be inherently imperfect, but the deterrent effect of unpredictable inspection, combined with predictable and swift consequences for violations, may prove preferable to the available alternatives.

Care must be taken that future revisions of the Guidelines do not become so restrictive that they simply drive nanotechnology research and development underground. This could expose compliant countries to the increased risks associated with decreased technical, economic, and military capabilities. It would also sacrifice the many significant economic, environmental, and medical benefits of nanotechnology that counteract serious and certain risks that society now faces in industrialized countries, and particularly in the developing world.

Center for Responsible Nanotechnology founded (US) US President George W. Bush continues support for development of nanotechnology 2002 TIMELINE | 7 Event UK Advisory Group on Nanotechnology Applications issues report: "New Dimensions for Manufacturing: A UK Strategy for Nanotechnology" Who UK Advisory Group on Nanotechnology Application When June 2002 Publication of report From Genomes to Atoms: The Big Down ETC Group 2003 Publication of. Their paper Mind the Gap: Science and Ethics in Nanotechnology cautioned that there was a risk of derailing the potential benefits of nanotechnology, including for developing countries, if the study of the economic, legal and social aspects of nanotechnology