Nanoethics (2008) 2:43–60
DOI 10.1007/s11569-008-0028-9

ORIGINAL PAPER

Risk Management Principles for Nanotechnology

Gary E. Marchant · Douglas J. Sylvester · Kenneth W. Abbott

Received: 8 January 2008 / Accepted: 8 January 2008 / Published online: 19 February 2008

Abstract Risk management of nanotechnology is challenged by the enormous uncertainties about the risks, benefits, properties, and future direction of nanotechnology applications. Because of these uncertainties, traditional risk management principles such as acceptable risk, cost–benefit analysis, and feasibility are unworkable, as is the newest risk management principle, the precautionary principle. Yet, simply waiting for these uncertainties to be resolved before undertaking risk management efforts would not be prudent, in part because of the growing public concerns about nanotechnology driven by risk perception heuristics such as affect and availability. A more reflexive, incremental, and cooperative risk management approach is required, which not only will help manage emerging risks from nanotechnology applications, but will also create a new risk management model for managing future emerging technologies.

Keywords Nanotechnology · Regulation · Risk management · Risk perception · Heuristics · Voluntary programs

G. E. Marchant (✉) · D. J. Sylvester · K. W. Abbott
Center for the Study of Law, Science and Technology,
Sandra Day O'Connor College of Law,
Arizona State University,
P.O. Box 877906, Tempe, AZ 85287-7906, USA
e-mail: gary.marchant@asu.edu

Introduction

Nanotechnology presents both an unprecedented challenge and an unparalleled opportunity for risk management. On the one hand, nanotechnology does not "fit" traditional risk management models, thereby impeding effective actions to manage nanotechnology risks using those existing approaches. On the other hand, nanotechnology will force risk managers to devise innovative new risk management approaches that may be applicable to other emerging technologies in the future.

As nanotechnology has emerged from the laboratory into industrial manufacture and commercial distribution, the potential for human and environmental exposure, and hence, risk, have become an increasing reality and priority. For purposes of this paper, we focus on the health, safety and environmental risks of nanotechnology, rather than more socio-economic or future risks such as privacy, terrorism, and economic displacement. As discussed below, the difficulties in identifying, never mind quantifying, the health, safety, and environmental risks of nanotechnology are a major impediment to applying traditional risk management approaches to nanotechnology. Risk management of nanotechnology is further challenged by the broad range of technologies and products encompassed within the term "nanotechnology," both in terms of current products and applications and even more in terms of future generations of products [26, 55]. The rapid pace of development of nanotechnologies, the difficulty in defining nanotechnology, and the substantial potential offsetting health and environmental benefits of some nanotechnology applications further complicate risk management of nanotechnology. Finally, risk management of nanotechnology must take into account public perceptions about the risks and benefits of nanotechnology and the growing public demands for regulatory oversight.

This paper analyzes the applicability of traditional risk management principles and new approaches based on the precautionary principle to nanotechnology, and finds these available approaches to be inadequate and unworkable. Nanotechnology will therefore require and force the development of new risk management models, an example of which we suggest here.

Existing Risk Management Principles

Both traditional risk management principles such as acceptable risk, cost–benefit analysis, and the feasibility principle, along with more recent innovations such as the precautionary principle, are inadequate to meet the risk management challenges presented by nanotechnology.

Traditional Risk Management Principles

The three most common traditional models for risk management of hazardous agents are (a) acceptable risk, (b) cost-benefit analysis, and (c) feasibility (or best available technology). *Acceptable risk* approaches rely on risk assessment to describe the risks of an agent, and then seek to reduce risks to levels that are socially acceptable. Current understanding of nanotechnology risks is too uncertain to permit meaningful risk assessment, and is likely to remain so for some time [11, 37, 43, 44, 48, 56]. There are no accepted test methods or validated data that can be used to prepare scientifically credible quantitative estimates of risk of specific nanotechnology applications at this time [59, 73].

Some initial animal studies have indicated the potential for toxicity in at least some nanomaterials, but these studies are very preliminary involving very high exposures that do not permit human risk assessment [47, 73]. Moreover, the initial studies give early indications of the likely complexity of nanotechnology risk assessment. For example, different forms of single-walled nanotubes present strikingly different risks depending on the manufacturing process and facility [37, 77]. The toxicity of nanomaterials appears to be determined by a complex set of characteristics, including size, surface area, chemical composition, coating, shape, and route of exposure [47, 48, 73]. Given this complexity, extrapolation of toxicological properties from other materials, including other nanomaterials, is currently unreliable, requiring risks to be determined on a case-by-case basis [19, 22, 59], a daunting prospect given the hundreds of nanotechnology products currently on the market and the thousands more to come.

Not only are risk-based approaches infeasible from a scientific perspective, they are also legally suspect, as regulators generally lack the risk information they need to make the threshold findings required to take regulatory action under, for example, most US environmental statutes [19]. The nanotechnology risk assessment dilemma is thus aptly summarized by Kristen Kulinowski, Executive Director of the Center for Biological and Environmental Nanotechnology: “We are in this awkward middle territory where we have just enough information to think there is an issue, but not enough information to really inform policymakers about what to do about it” (quoted in 46).

Another complication is the rapid pace of nanotechnology development, which is rapidly outpacing the development of risk assessment for these technologies [56]. For example, even though hundreds of nanotechnology products are already on the market, some involving significant exposure to workers and consumers, the US Environmental Protection Agency’s recently issued White Paper on nanotechnology provides a timeline for oversight, in which it will not be until the year 2011 or 2012 that the agency has sufficient risk knowledge to develop a systematic approach for managing the risks of nanotechnologies ([77], p. 112). Of course, by that time, virtually every citizen will have been exposed to nanomaterials, and new generations of nanotechnology products will be entering the market, creating new risk uncertainties. As David Rajeski of the Woodrow Wilson Center warns, “[i]f you think that any existing regulatory framework can keep pace with this rate of change, think again” ([55], p. 45). This is not to say that the development of risk assessment approaches for nanotechnology is not needed, but only that we cannot rely on risk-based approaches to provide the primary risk management solution in at least the short-term for such a rapidly emerging technology.

Finally, acceptable risk approaches generally suffer from a structural disadvantage: by only considering risks and their acceptability, they disregard other important factors such as the benefits of the technology creating the risks and the costs of reducing risks. As discussed below, these factors are likely critical for socially optimal decisions about nanotechnology.

A second traditional risk management model is *cost–benefit analysis* or balancing, in which the costs and benefits of proposed risk management options are balanced. Unlike the acceptable risk model, the cost–benefit model has the advantage of considering both the benefits and risks of nanotechnology, which is important given that nanotechnology is likely to present both risks and benefits for public health and the environment [19]. Nonetheless, the cost–benefit model is ill-equipped for managing nanotechnology at this time, given the immense uncertainties about its risks and benefits.

The enormous number and diversity of potential nanotechnology applications also make this approach unfeasible—a global cost–benefit balancing for nanotechnology as a whole would mask the significant cost–benefit variance that likely exists between different applications. Alternatively, performing separate cost–benefit balances for each specific nanotechnology application would likely overwhelm available risk management resources given the large number of potential applications. While some qualitative weighing of the risks and benefits of nanotechnology may be a useful exercise for purposes of thinking about how those risks should be managed, cost–benefit analysis does not provide a workable risk management approach for nanotechnology at this time.

The third and final traditional risk management principle is the *feasibility* or *best available technology* approach. This approach, which requires reduction of risks to the lowest level technologically or economically feasible, has the advantage of not requiring information about risks or benefits. Indeed, the feasibility approach has achieved considerable popularity among policymakers in recent years because it allows circumvention of controversies over risk analysis and jumps straight to reducing risks to the extent possible [5, 14, 61, 78]. Given the enormous uncertainties about nanotechnology risks, this approach has some appeal. The strength of the feasibility approach is also its key weakness, however, because, while ignoring risk information avoids controversy, it also avoids addressing what is truly important, which is risk [45, 66]. The feasibility approach thus may over-regulate or under-regulate risks depending on whether the best available technology is necessary or sufficient to reduce unacceptable risks. This problem could be particularly problematic for an emerging technology such as nanotechnology.

Moreover, it is not at all clear how the best available technology approach could or would work for nanotechnology. Certainly, the application of work practices and other precautions that seek to limit exposure to nanomaterials appears sensible. But we may not know enough to go further and require technology controls on production processes, including pollution control of emissions, when we know so little about the nature, level, controllability, and risks of released nanoparticles. It is also difficult to apply a best technology approach to the growing number of nanotechnology consumer products given that the public's use and disposal of such products is difficult to control. Finally, a best available technology approach deters companies from developing better control technologies, since doing so will only result in more stringent (and hence costly) regulations [2]; this is the wrong incentive for a rapidly emerging technology such as nanotechnology [57].

In sum, none of the three traditional models of risk management are capable of effectively managing the risks from nanotechnology at this time given the tremendous uncertainties that exist for this emerging technology. This realization has prompted some to advocate application of the most recent entry in the risk management toolbox, the precautionary principle.

The Precautionary Principle

Numerous public interest groups and scholars have called for the precautionary principle to be applied to nanotechnology [17, 21, 27, 38, 79]. The precautionary principle, which has emerged in recent years as an alternative approach to risk management, is often summarized by the phrase "better safe than sorry." The precautionary principle recognizes that health and environmental decisions often must be made in the face of pervasive uncertainty, and therefore calls on decision makers to err on the side of safety by delaying new technologies until their safety can be adequately ensured. This requirement is often framed in terms of shifting the burden of proof to the proponent of a technology to demonstrate its safety. Given the massive uncertainty about nanotechnology risks, this technology might appear to be an ideal candidate for application of the precautionary principle. Yet in fact nanotechnology vividly demonstrates the limitations of the principle as a decision-making tool; the precautionary principle too is not a workable risk management model for nanotechnology.

The first problem with the precautionary principle is that it is too poorly defined to serve as a decision-making rule. While lawmakers and proponents frequently cite to "the" precautionary principle, there is no standard text for the principle, and the dozens of formulations that have been suggested differ in important respects [52]. Moreover, no version of the precautionary principle answers the critical questions that need to be considered in moving forward with regulatory decisions, such as what level or type of evidence (if any) of harm is sufficient to trigger the principle, what quantum and types of data must a manufacturer produce to satisfy the principle, what level of risk is acceptable, and how should the benefits of a technology be weighed against its risks (if at all; [40]). Without any criteria or guidelines to resolve these questions, the precautionary principle is prone to arbitrary and capricious decision-making, if not outright mischief. Examples of such unreasonableness include the invocation of the precautionary principle to ban corn flakes enhanced with essential vitamins in The Netherlands, prohibit caffeinated energy drinks in France, prohibit cranberry juice beverages in Denmark because they contain vitamin C, and reject food aid containing some genetically modified corn in the famine-affected nation of Zaire [41].

The precautionary principle suffers from another flaw in that it is biased toward the status quo, impeding new technologies even if they may ultimately prove beneficial for the environment or public health [12, 25]. It is quite possible that a freeze or moratorium on nanotechnology per the precautionary principle would do more harm than good to human health and the environment, never mind the many other benefits of nanotechnology that would be forgone. Nanotechnology offers many promising health and environmental benefits, including more effective and safer cancer treatments, improved medical diagnostics, remediation of hazardous wastes, cleaner energy sources, and improved control of pollution emissions [24, 77]. Given these potential benefits, precaution and emphasis on protecting health and the environment might actually weigh in favor of promoting rather than restricting nanotechnology. The precautionary principle fails to provide guidance on which direction to pursue [69, 70]. Finally, there is evidence that application of the precautionary principle increases, rather than addresses, the public's concerns and anxiety about a technology [82], and thus cannot be defended on the instrumental ground of enhancing public assurance.

Some have suggested that the precautionary principle makes the most sense for protecting against "catastrophic risks" that could irreversibly destroy major parts of the human population or the earth's ecosystem [72]. This argument has been applied to nanotechnology. For example, two weeks after the 9/11 terrorist attacks on the United States, the *New York Times* published an article debating whether, in light of the horrific application of a generally peaceful technology demonstrated by the 9/11 attacks, humanity might be better off in the long run to forgo a powerful new technology such as nanotechnology that could be used for enormous good or evil [34]. Commentators such as Bill Joy have speculated that nanotechnology could be used to develop swarms of self-replicating nanobots that could destroy the planet, often referred to as the "grey goo" scenario [29]. Applying the precautionary principle to this possibility, the argument goes that no amount of potential benefits from nanotechnology would justify assuming a risk (no matter how small) of such a catastrophic consequence.

There are several problems with this nano-catastrophism argument. First, virtually all serious analyses of the grey goo scenario—including a recent analysis by Eric Drexler, who first posed the problem [13]—have concluded that it is extremely implausible if not impossible [51]. Second, catastrophic scenarios can be envisioned for virtually any technology, but we would be paralyzed into inaction if we avoided any technology that could be associated with such a scenario, no matter how implausible. The first environmental release of a genetically modified organism, the so-called ice minus bacterium, was alleged to create the risk of destabilizing global climate by spreading throughout the upper atmosphere and disrupting the normal cloud seeding processes [30]. The new particle accelerator built for the Brookhaven National Laboratory in 2000 allegedly could have produced a shower of quarks that might have turned the entire earth into some new type of matter [53]. Any given international traveler could conceivably have inoculated himself with the smallpox virus and entered the country with the intention of starting a devastating pandemic, a catastrophic risk that could only be prevented by banning all international travel. While preventing catastrophic risks should be a top priority, it would not be practical to allow the mere possibility of some remote catastrophic risk to be a sufficient rationale for banning a promising technology [53, 72].

Finally, while it may be possible to imagine potential catastrophic risks from nanotechnology, it is also possible to envision potential future applications of nanotechnology that could save us from other catastrophic risks. For example, medical applications of nanotechnology may give us the tools to stave off the next pandemic virus that could be mutating in some distant corner of the world right now. Or nanotechnology may be used to protect us from some asteroid hurtling through space on a path that will intercept the earth sometime in the future. While these scenarios may be remote, so too are the catastrophic scenarios for nanotechnology. Truth be told, no one knows for sure which of these implausible scenarios are the most implausible. Banning nanotechnology based on the precautionary principle could just as easily prevent as create a future catastrophic risk to humankind. Thus, while the precautionary principle provides a useful general philosophy, and while some application of precaution is certainly appropriate to guide the development of nanotechnology, the precautionary principle itself fails to provide a workable risk management approach.

In sum, neither the traditional risk management principles nor the new precautionary principle provides an acceptable approach for regulating nanotechnology. For past technologies, that would normally have meant that risk regulation would be postponed until further evidence of a real problem had emerged [19, 37, 83], perhaps with the interim use of some secondary risk management tools such as risk communication, liability, self-regulation, and insurance [6, 26].

For the same reasons, it is perhaps not surprising that, despite the enormous attention given to nanotechnology in recent years, no national government has yet enacted any traditional nano-specific regulation [8, 9].

Public Perceptions of Nanotechnology Risks

As the previous section demonstrated, anticipatory regulation of nanotechnology pursuant to current models appears inappropriate and ill-advised. It is equally inadvisable, however, for regulators to sit back and do nothing. As we have previously noted, legal regulation of nanotechnology is inevitable [42]. The ethical, environmental, and economic issues that nanotechnology raises, whether probable or fantastical, will necessitate some regulatory response, if only because of political pressure. The public is already making up its mind about nanotechnology's risks and benefits, and they are doing so on the basis of very little information [3, 31, 50, 62].

The public's willingness to form premature opinions about nanotechnology's benefits and risks may jeopardize the development of the technology. As one commentator put it, "[o]ne of the greatest challenges facing nanotechnology is avoiding a backlash from the public that slows or even halts the progress of research and development." ([62], p. 335). This experiential concern, derived from examples of public backlashes against other technological developments including genetically modified organisms, nuclear power, and recombinant DNA technology, merges with another area of research on how the public forms opinions about risks in the face of uncertainty.

Decades of research into risk perception have discovered a series of cognitive and emotional responses that influence how individuals perceive risks. Unlike prior models of human cognition as calculating, rational, and logical, this research has revealed that individuals employ "heuristics," mental short-cuts based on experience and emotion, to assess potential risks and benefits. Although these heuristics work well in many cases, they are especially vulnerable in cases of uncertainty, where they are prone to systematic and predictable errors.

Adding to the concern that individuals may make skewed determinations of risk based on heuristics is the fact that such decisions are not made in a vacuum. Social interactions, for example, play an important role in the formation and reinforcement of heuristic reasoning [31, 71]. As a result, reactive media attention and the concerted actions of public interest groups can directly affect how individuals initially perceive risks—often resulting in the cementing of opinion on a given technology's risks and benefits, making that opinion exceedingly difficult to change and often resulting in extreme positions [31, 64].

The danger that nascent public opinion on the risks and benefits of nanotechnology will be inappropriately informed by unbalanced media attention, reinforced by social interactions, and cemented into extreme positions of fear and dread is certainly real. Worse yet, there is growing evidence that many individuals are predisposed to fear nanotechnology, as demonstrated by the overreaction to the initial (but false) reports that people made ill by the German cleaning product "Magic Nano" were the first demonstrated injuries from a nanotechnology product (it was subsequently revealed that, despite its name, Magic Nano was not a nanotechnology product) ([83], p. 704; [75]).

The decades-long research of behavioral scientists has led to one unmistakable conclusion—human beings seldom assess uncertain risks based solely on information, probabilities, or logical assessment. When confronted with questions such as, "will rapidly advancing nanotechnology revolutionize health-care, the nature of computers, and the structures of materials, or will it lead to as yet uncontemplated new forms of pollution and cancer," ([39], p. 119) people do not wait until they have information to make an assessment. Instead, people rely on heuristics to make a quick, intuitive, and at times emotional assessment about the likelihood a given risk will occur. In place of rational assessment, behavioral science posits that "people rely on a limited number of heuristic principles which reduce the complex tasks of assessing probabilities and predicting values to simpler judgmental operations" ([76], p. 3). In this way, heuristics "play a role in aiding individuals in ascertaining the relative risks posed by future events without resorting to more accurate, but time consuming, statistical analysis" [74]. Unfortunately, "heuristics serve people well in many circumstances, but they also create vulnerabilities to the predations of advertisers, political spin doctors, trial attorneys, and ordinary con artists" ([54], p. 1165).

The research into risk perception has identified numerous common heuristics that people employ to assess risks and benefits. For example, in many cases people react to possible risks by focusing solely on the type of harm rather than the probability of harm (probability neglect) [68]. In other situations, when confronted with actual harms, individuals seek human causes and downplay the possibilities that chance and nature are responsible (the mythical benevolence of nature; [69, 70]. Finally, other researchers have concluded that individuals cannot simultaneously perceive that individual technologies have both benefits and risks, and as a result cognitively privilege benefits to the exclusion of risks or vice versa (cognitive dissonance avoidance; [39]). These are but a few of the cognitive and emotional processes that researchers have identified as affecting risk perceptions [52]. The key to understanding these heuristics and biases is to see that individuals do not arrive at them analytically: they are knee-jerk, unreflective, intuitive, reactive, and experiential.

The most important of these heuristics for nanotechnology is "Affect." First introduced by Paul Slovic, the Affect heuristic "refers to people's tendency to rapidly and automatically have positive or negative feeling when confronted with a certain word, concept, or other stimulus" ([39], p. 161). In other words, individuals have a predisposition, most likely unconscious, towards various stimuli. When confronted with such a stimulus, individuals react to it affectively [28, 31, 52, 63, 65, 67, 70, 76].

An interesting effect of the Affect heuristic is the tendency of individuals to negatively correlate a technology's perceived risk with its benefit. In other words, numerous studies have shown that where individuals believe a technology has high benefits, they automatically believe its risks are low. Conversely, where risk is perceived to be high, the benefits are correspondingly seen as low. This negative correlation has been shown to affect both lay and expert opinions, and is robust even in the face of countervailing evidence [63]. Based on this finding, researchers hypothesize that individuals possess an emotive and ingrained response to various stimuli. In short, how individuals *feel* about a particular stimulus directs how they perceive its dangers or benefits [63].

The question, of course, is how do people arrive at these affective orientations? In many cases, individuals do not assess technologies as separate risks (or benefits), but instead adopt a world-view that automatically, or "affectively," views technology as risky [63]. A recent study by Kahan et al. [31] investigated the role that Affect played in determining assessments of nanotechnology risks. According to Kahan, "the visceral, emotional responses of our subjects, pro or con, determined how beneficial or dangerous they thought nanotechnology was likely to be…" ([31], p. 3).

Of some comfort to those interested in nanotechnology is the finding that "these instantaneous judgments were not static. Individuals exposed to information on the risks and benefits of nanotechnology formed different views from individuals not so informed" ([31], p. 3). Less heartening, however, was the finding that "the *ways* in which information influenced our subjects—whether it inclined them to see nanotechnology as more risky or more beneficial—was highly conditional on the values they held" ([31], p. 3). Thus, individuals who viewed themselves (or, better put, were affectively inclined) as hierarchical or egalitarian, individualistic or communitarian, interpreted information about nanotechnology risks and benefits to *conform* to their affective dispositions. The fact that Affect so deeply guided individual views on the risks or benefits of nanotechnology, even in the face of evidence contrary to their initial beliefs, led Kahan and others to conclude:

> These results paint a picture … of at least one possible future for nanotechnology. It is one in which citizens rapidly take affect-driven positions, which harden as they conform what they learn thereafter to their more basic cultural attitudes toward technology and risk. The result is likely to be a state of political polarization over the desirability of nanotechnology that very much resembles the one that now exists for other controversial environmental issues, including nuclear power and global warming. Or at least that is how things are likely to play out absent the development of strategies that neutralize the tendency of persons to assimilate information in a manner that confirms their emotional and cultural predispositions ([31], pp. 3–4).

The bleak picture painted by this study is made worse when we consider a second common heuristic, "Availability." The Availability heuristic is among the most widely studied and has achieved the most attention in legal circles. Although it is employed in numerous aspects of cognition, its chief role is in assessing the *likelihood* of risks. [76]. According to researchers, individuals who can easily recall a memory specific to a given harm are predisposed to overestimating the probability of its recurrence, compared to other more likely harms to which no memory is attached. In other words, "the availability heuristic captures the mental process by which people assume that events more easily recalled are more likely to recur" [74]. A classic example of Availability is that individuals who have seen or read about a house burning down are more likely to believe their own house will burn down than: (a) they were prior to witnessing the event, and (b) others who have not shared a similar experience.

The nature of the initial experience also determines the influence of the Availability heuristic on the assessment of recurrence probabilities. Thus, "[t]he impact of seeing a house burning… is probably greater than the impact of reading about a fire in the local paper" [74]. The recency of an experience also affects the level of risk individuals assign to a given event.

Numerous factors affect individual recall and Availability, including suggestion, memorization, recency, and the amount of information recalled. If we are to take the existence of the Availability heuristic seriously, we must assume that the subjective perception of risk for a given event or policy may be linked to the respondent's experiences, with some experiences having a more direct effect than others; "availability may be endogenous to individual predispositions" ([69, 70], p. 759).

An important point for understanding these heuristics is that each is sensitive, perhaps keyed into, the imagery and meaning associated with a given risk; and imagery and meaning are deeply affected by social interactions. With Affect, the imagery and perception of a risk may be endogenous, but can be influenced and exacerbated by social interactions [63, 69, 70]. With Availability, the imagery and meaning attached to a risk are by definition found in public messages and experiences. The effect that external imagery and social reinforcement can have on individual perceptions of risk has led many to conclude that the media, public interest groups, industry, and government have real power to influence the public's perceptions about the nature and probability of risk associated with a given technology.

Many have noted, and it is indeed intuitive, that as media attention and the actions of interested parties continually emphasize certain risks, especially when they tend to stigmatize the relevant technology, it is possible for such images to overwhelm objective and balanced information about true probabilities and harms. Thus, there are real incentives for those who oppose implementation of a new technology (for whatever reason) to highlight images that evoke dread, fear, disgust, and similar emotions (cf. "Frankenfood"). Where political actors are able to persistently project such images, technologies may be stigmatized as inherently dangerous [33]. Once so stigmatized, "vivid images and concrete pictures of disaster can 'crowd out' other kinds of thoughts, including the crucial thought that the probability of a disaster is really small" [69].

The real importance of stigmatization is that it has effects far beyond the individuals who may be convinced of the dangers of a given technology or application. As just noted, individuals form perceptions of risk based in large measure on social interactions with peer groups [63]. Thus, "representative anecdotes and gripping examples can move rapidly from one person to another. Once several people start to take an example as probative, many people may come to be influenced by their opinion, giving rise to cascade effects…[a] problem [that] might well be aggravated by certain media and new technologies" ([70], p. 759). Cascade effects are most closely associated with Availability, but can also have impacts on Affect by reinforcing affective dispositions and associating technologies with specific risks that individuals are affectively predisposed to fear or loathe ([33, 81], p. 27, [23, 35, 54, 63], p. 221).

Nanotechnology is especially susceptible to cascade effects and affective hardening of positions. The experiences of GMOs and other "controversial" technologies have led political actors to consciously highlight dreaded harms and to persistently publicize anecdotes that reinforce the availability of such harms in connection with new technologies like nanotechnology. As a result, some now believe that "functional discourse… is largely absent from technology debates, and the climate necessary for productive discourse is poisoned" ([39], p. 117). Indeed, in the case of nanotechnology, some already believe that "the only messages… currently reaching the public are negative ones portrayed in movies and television…" ([62], p. 335). Media depictions of nanotechnology tend to emphasize fantastical risks, including widespread environmental degradation, increased cancer, and even the destruction of the human race [62]. In the face of such overwhelming negative publicity, it is quite likely that nanotechnology will be subject to availability cascades, strongly anti-nano affective attitudes, and overestimates of the probability of specific risks.

This bleak picture is not, however, the only possible future for nanotechnology. Although we have focused here on the ways in which heuristic processes may lead to overestimation of risks, it is equally possible that individuals may focus on benefits. Availability cascades and affective attitudes may also produce perceptions of the benefits of the technology as overly probabilistic and overly desirable. If available imagery and affective reasoning (and messages sent to appeal to such reasoning) strongly support the technology, these may offset risk-based reasoning or even overwhelm it.

Of course, no one yet knows whether nanotechnology's risks do outweigh its benefits. What is clear, however, is that fear-based, or risk-focused, attention to nanotechnology has recently begun to outweigh attention to its potential benefits [62]. In order to keep open the possibility of a reasoned discourse on the risks and benefits of nanotechnology, therefore, some action must be taken to reassure the public that nanotechnology risks are being actively managed, while avoiding the pitfalls of anticipatory regulation under pre-existing models. As Cass Sunstein has noted, "Government should take action that reassures people, even if such actions are not justified on technical grounds" [64, 70].

The question is, what action? In the next section, we outline a plan for a gradual, flexible, and evolutionary approach to nanotechnology regulation that we believe would be an important first step in restoring balance to the discourse on nanotechnology risks and benefits, while reassuring the public that steps are being taken to identify and control those risks. Our approach is consistent with a developing consensus among regulatory scholars that a non-traditional "soft law" approach that is incremental, reflexive, and cooperative will be needed to manage the risks of nanotechnology, at least in the near future [7–9, 18, 20, 36, 37, 49, 57, 60, 80, 83; but see 27].

New Risk Management Model

Given the pervasive uncertainty and dynamism of current nanotechnology developments, we suggest a flexible, evolutionary approach to risk "regulation," especially in the immediate, near and medium terms. The approach we suggest is flexible in two senses. Substantively, it draws on multiple approaches to addressing risk, not only the accepted models of risk regulation discussed above, but also more general approaches to the appropriate handling of risky technologies, such as the notion of product stewardship and the professional ethics of researchers. Procedurally, as these examples suggest, our approach favors subsidiarity and decentralization: the participation of a wide range of private and public stakeholders, including those currently researching and commercializing nanotechnology, in developing and applying risk management norms. Broad stakeholder participation should help us gain a better understanding of the actual risks and benefits of particular nanotechnology products and processes, communicate that understanding to the public, and enable multiple approaches to managing risk. Over time, the experience and learning these approaches produce should allow societies to gradually develop appropriate and cost-effective systems of regulation.

While no regulatory approach can overcome all the heuristics that distort individuals' assessments of risks, this approach does address some of the major issues. Our approach emphasizes transparency and dissemination of information to the public by all participating actors; in addition, it emphasizes active participation of stakeholders, including public interest groups, in developing and applying risk management norms. The result would be to create and disseminate accurate images, and even actual experiences, of the benefits, risks, and probabilities associated with nanotechnology. Transparency and participation would be ongoing, providing regular updating of availability perceptions. In addition, this approach would demonstrate to concerned onlookers that multiple actors, from academic researchers to manufacturers to government agencies, are actively identifying and addressing potential risks through a range of techniques. Finally, if properly managed, this approach should generate an increased level of public trust in those responsible for the development of nanotechnology.

As noted above, the current risk status of nanotechnology is dominated by uncertainty. Many potential environmental, health, and safety risks are uncertain in terms of severity, threshold exposure levels, variations among even closely related products and processes, and the like; other risks are "unknown unknowns." A lack of definitional certainty as to what exactly constitutes nanotechnology exacerbates this uncertainty and hampers efforts to identify and address the risks of particular applications, technologies, and research processes.

Despite this pervasive uncertainty, the growing political pressure to regulate and the need to control social responses mean that some form of regulatory response will likely be required in the near future. Since nanotechnology R&D, manufacturing and other activities are proceeding apace, it is important to begin immediately to develop ways to deal with the potential risks, whether fully actualized or merely perceived.

Theoretical Background

Although traditional models of technology risk regulation are inappropriate for the reasons discussed above, lessons from the broader study of regulation are highly relevant. One of the most influential recent works in that literature, Ayres & Braithwaite's *Responsive Regulation* [4], provides the theoretical inspiration for our approach.

Responsive Regulation was written at a time of widespread debate about the merits of "deregulation." Ayres & Braithwaite, however, argued that the choice between "regulation" (understood as traditional command and control) and "deregulation" is a false one: Regulation always involves a symbiotic relationship between public and private actions, and the interactions between these two realms can be managed, responding precisely to varying conditions and behavior across industries and even firms, to obtain better regulatory outcomes. Their emphasis on substantive flexibility and subsidiarity or decentralization mirrors our own views on appropriate risk regulation for nanotechnology.

Ayres and Braithwaite employed the construct of a "pyramid" to illustrate their flexible approach to regulation. Most narrowly, this pyramid depicts the spectrum of possible *sanctions*—from persuasion and warnings, at the base, up through civil, licensure, and criminal penalties, at the peak—available to regulatory agencies. A broader version of the pyramid captures the range of regulatory *strategies* available at the national level—from self-regulation at the base, through supervised or enforced self-regulation and other forms of public–private interaction in the middle, to standard forms of command-and-control regulation, still with a range of possible penalties as depicted on the original pyramid, at the peak (Fig. 1).

With this tool kit at hand, regulators can play a tit-for-tat strategy: they allow firms to self-regulate so long as the firms reciprocate with responsible action; if instead some firms act opportunistically, regulators respond to the defectors with appropriate penalties and more stringent regulation. The threat of regulatory intervention both deters non-compliance by potential defectors and encourages all firms to develop an attitude of social responsibility. If this strategy is skillfully deployed, the majority of regulatory activity will occur near the base of the pyramid, at the lowest levels of governmental intrusion (e.g., persuasion or self-regulation), with more intrusive actions taken only when softer measures prove unworkable.

Recent scholarship has attempted to apply the regulatory pyramid to nanotechnology. Bowman and Hodge [9] use a pyramid model to argue for a complex regulatory system for nanotechnology (Fig. 2). Their pyramid is hexagonal, with six sides that correspond to families of issues including occupational health and safety, environmental protection, product safety, privacy and civil liberties, intellectual property, and international law. In addressing each set of issues, the pyramid suggests that regulators should deploy a range

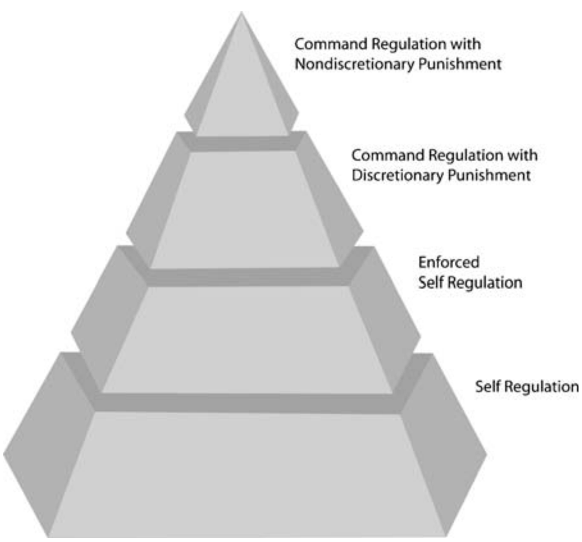

Fig. 1 Ayres & Braithwaite regulatory pyramid (adapted with permission from Ayres and Braithwaite [4]

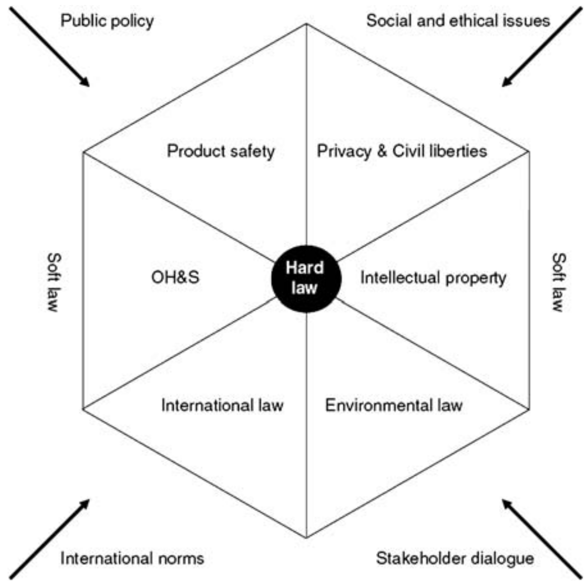

Fig. 2 A conceptual model for the regulatory frontiers of nanotechnology. Source: Bowman and Hodge [7]

of regulatory options, from "soft law" to "hard law," as called for by Ayres & Braithwaite [4].

While the pyramid model provides the theoretical inspiration for our approach, both the Ayres/Braithwaite and Bowman/Hodge models are largely static: they envision a fully developed regulatory system that can effectively manage a particular set of risks. Both models, moreover, are designed for advanced nations with highly developed legal systems, in which legislatures and agencies can create, communicate, and utilize a range of regulatory options. Finally, both at least implicitly assume high levels of information and understanding on the part of regulators. As such, both models may better depict a potential future regulatory end state than an immediate approach capable of addressing the risks of a dynamic technology in an environment of pervasive uncertainty, as is the case with nanotechnology.

Incremental Regulation

The concepts behind these models, however, can be deployed in a different way to address the uncertain risk situation of nanotechnology today. Rather than viewing persuasion, soft law, self-regulation, command-and-control regulation and the like as components or "layers" within a static regulatory system, we extend the regulatory pyramid through time, viewing these approaches as sequential.

In the near term, we would begin with softer and more decentralized measures, including self-regulation. We would emphasize those measures that will produce the greatest information, coupled with mechanisms for learning from them. As in the original pyramid model, such soft regulatory strategies are the most flexible and the least costly and intrusive. To be sure, however, in a sequential approach the initial stages will not be subject to the ongoing agency oversight and threat of tit-for-tat regulatory intervention that characterize a fully developed regulatory system; the risk of opportunism undoubtedly exists. That risk is the principal basis for criticizing incremental or soft law approaches. For example, in proposing a set of "Principles for the Oversight of Nanotechnologies and Nanomaterials," a coalition of civil society organizations recently argued: "Voluntary approaches are wholly inadequate to oversee nanotechnology. Voluntary programs lack incentives for 'bad actors'… to participate, thus leaving out the entities most in need of regulation" [27].

Yet there are other incentives for responsible action by those involved in nanotechnology. These include the need to overcome public fears and avoid reputational costs, stigmatization, and backlash; the risk of costly litigation and liability; and the desire to forestall inappropriate mandatory regulation. To a considerable extent, these social forces fill the role of the regulator's tit-for-tat strategy. Some actors, such as scientific researchers, are influenced by professional norms of responsibility, quite apart from legal incentives. Civil society organizations can pressure those working with nanotechnology to adopt responsible approaches, and can participate in multi-stakeholder programs; soft regulation need not be limited to self-regulation.

In addition, as we argue below, even in the near term we would encourage regulators and other public officials to promote responsible private actions and steer them in desirable directions (for example, by encouraging transparency and participation, two other principles proposed by the coalition); fund research on the risks of nanotechnology (perhaps focusing on broader and longer-term risks that private actors have weaker incentives to consider); and take other supportive actions short of mandatory regulation.

Finally, the call for immediate regulation overlooks two countervailing considerations. The first is the social cost of regulation that is inappropriately designed; the risk of design error is quite high, since immediate regulation would almost inevitably be based on one of the existing risk management approaches we have criticized as inadequate. The second is the difficulty of achieving political consensus and action on specific forms of regulation—even in one country, let alone on a harmonized basis across countries. Soft approaches may lack the potency of mandatory law, but they are often much easier to put in place.

Over time, as society learns about the actual risks and benefits of nanotechnology from its early experiences with state-supported voluntary measures, a regulatory end state similar to the Aryes/Braithwaite and Bowman/Hodge pyramids can be gradually built up. All stages of this approach should be managed with a high degree of transparency and stakeholder participation, to provide accurate images of risks, benefits and probabilities, build trust, and set the stage for an effective regulatory system.

Figure 3 provides a rough graphical depiction of our incremental approach. We suggest only a few broad regulatory categories in Fig. 3, but these appear to correlate well with current initiatives seeking to address the risks of nanotechnology. Of particular importance is the "Immediate" section of our gradual pyramid. If it is true, as suggested above, that nanotechnology may be sliding towards stigmatization, producing more complete information about current research, manufacturing and other applications, and about the actual risks and benefits of nanotechnology, is essential to avoid stigmatization and help form appropriate affective reactions to this emerging technology.

Information Gathering and Dissemination

In the immediate future, the greatest need for "regulation" is to increase the accumulation and dissemination of knowledge about the current state of nanotechnology research, development, and application, and about the risks and benefits of specific products and processes. Information gathering and dissemination are essential to properly assess the risks of nanotechnology and to provide accurate and trustworthy messages to the public. Such information may not overcome all negative affective reactions, but accurate and trustworthy messages should at least avoid any widespread affective antagonism or availability cascades.

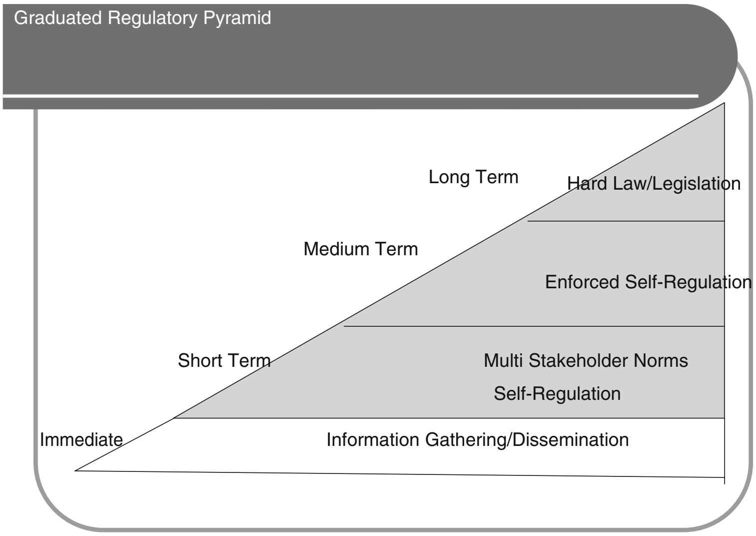

Fig. 3 Building up the regulatory pyramid over time

Numerous nanotechnology stakeholders support this early emphasis on information. For example, the American Chemistry Council (ACC) has adopted a formal *Position on Nanotechnology*. It calls for increased public funding for research into methods to assess the impact of nanotechnology on environment, health, and safety. The ACC also calls for global coordination of regulatory, research and standard-setting activities, including an assessment of the adequacy of existing statutes and regulatory programs that might apply to nanotechnology. The ACC has been joined by an unlikely ally in this effort: Environmental Defense (ED), a non-governmental environmental organization. In a 2005 *Joint Statement of Principles* [15], Environmental Defense and ACC note the urgency to "identify and better understand nanotechnology's potential risks up front." To this end, the two groups call for "an international effort to standardize testing protocols, hazard and exposure assessment approaches, and nomenclature and terminology," as well as a significant increase in public funding for safety research and an assessment of existing regulations. More broadly, the two groups call for a broad "multi-stakeholder dialogue" including all interested parties. In the view of the ACC and ED, such an "open and transparent process" is the best guarantee that the potential risks of nanomaterials will be identified and minimized. We would add that such a process is also the best guarantee of public understanding and trust.

ED has also joined with other partners from industry to seek appropriate regulatory responses to nanotechnology. Perhaps most notably, ED partnered with DuPont to develop a *Nanorisk Framework* [16]. This project aims to develop a practical framework to understand and manage nanotechnology risks. It seeks a responsible approach to development, production, use, and disposal of nanoscale materials across the entire nanotechnology product lifecycle. Like our regulatory model, the project includes several phases. The initial phase includes several steps that are consistent with our approach: (1) identifying the risks of the nanomaterials dealt with by DuPont and identifying the tests appropriate for particular products at various stages of development; (2) developing techniques of risk management, with a focus on safe procedures for handling nanomaterials at different stages of the lifecycle; (3) developing transparency mechanisms, techniques for informing internal and external stakeholders, including the general public, about risk identification and risk management decisions; and finally (4) establishing systems to track the implementation and determine the efficacy of risk management techniques, with appropriate feedback, evaluation, and adjustment. This program seems well designed to produce and disseminate accurate information, develop unconventional forms of "regulation" in the form of safe handling protocols, and build trust among stakeholders and the general public.

Self-Regulation

In the short term, our approach emphasizes self-regulation, coupled with the development of multi-stakeholder norms and policies. Once again, this approach not only recognizes the nascent nature of nanotechnology, but also corresponds to the actions of relevant stakeholders.

The ACC again offers a relevant example. In its *Position on Nanotechnology*, the organization makes clear that, even as other regulatory options are studied, member firms are to apply to their nanotechnology activities the "product stewardship" principles incorporated in the ACC Global Chemicals Management Policy and the Responsible Care program developed by the International Council of Chemical Associations (ICCA) in response to the 1984 Bhopal incident. To be sure, "product stewardship" is a highly amorphous concept, but it incorporates important elements, such as commitments to a product lifecycle approach and to continuous improvement, and represents a familiar self-regulatory approach that business can readily implement. The ICCA has declared that the product stewardship concept will increasingly pervade self-regulatory programs such as Responsible Care, and has pledged to establish a strengthened and unified product stewardship program to evaluate and manage chemical risks and benefits, including those relating to nanotechnology. Meanwhile, the Responsible Care program—which was very weak in its initial form—has been strengthened by the involvement of outside stakeholders through third-party audits of compliance, an approach that should help strengthen public trust in self-regulation within competitive industries. The ACC–ED *Joint Principles* likewise call for "appropriate protective measures" to be implemented as an aspect of responsible technology development while information is gathered and standards created.

The ED–Dupont agreement observes that it would be harmful to business as well as society to commercialize a new technology like nanotechnology before appropriately identifying and managing its risks; potential costs to business might include costly litigation, liability and cleanup requirements. The agreement specifically declares that it is in the best interests of industry, the public, and the environment for business to proactively develop, in advance of government regulation, a framework for responsible nanotechnology management. In addition to directly enhancing safe product development and public acceptance, these actors view a self-regulatory framework as providing a workable model for reasonable government policy at a later time, consistent with our gradual approach to building up a complete regulatory system. In their agreement, DuPont and ED commit to developing both broad self-regulatory principles (e.g., that new nanomaterials will be tested before being marketed) and specific guidelines on implementing those principles. Finally, the partners agree to demonstrate the framework they develop by applying it to an actual product or process; revise and refine the framework following the pilot project; apply the framework across DuPont's activities involving nanotechnology; and disseminate the framework to other firms, industry associations, and government agencies, promoting it as a regulatory model.

Multi-Stakeholder Norms

Ayres & Braithwaite [4] correctly observe that self-regulation alone is unlikely to be satisfactory. Experience with other industries and technologies where firms face strong competitive incentives to minimize self-regulatory expenditures makes clear the need to monitor compliance and exert pressure for compliance on firms, especially in the absence of mandatory regulation; failures of self-regulation can be highly damaging to public confidence, especially if they cause visible harms that influence public attitudes through the action of the availability heuristic. Ayres & Braithwaite conclude that the best way to ensure industry compliance is to bring groups representing the public interest into the regulatory process.

In our incremental regulatory approach, each stage can, and probably should, involve the participation not only of firms, researchers, and other targets of regulation (who may also be engaged in self-regulation), but also of appropriate advocates for the public interest and other stakeholders. Ayers & Braithwaite envisioned a complex formal process for incorporating these groups in public regulatory procedures; especially in the early stages of the gradual regulatory process we envision, however, participation will have to be more informal, as in several recent developments.

The ACC–ED and ED–DuPont agreements already reflect a decentralized multi-stakeholder approach. Moreover, both explicitly endorse broader multi-stakeholder participation. The ED–DuPont agreement, in particular, commits the parties to "engage a wide range of … stakeholders at various stages throughout the project to draw on their expertise and solicit input…. The project will ... include interim checkpoints for the Parties to … share interim results with other stakeholders, and solicit input from other stakeholders." In addition, the goal of the ED–DuPont project is to develop a framework for (self-)regulation "that will be accepted, endorsed and adopted by a wide range of stakeholders, including other companies, other public interest groups, academia and government agencies."

The best and most prominent example of a multi-stakeholder approach to voluntary regulation is the Foresight Guidelines for Responsible Nanotechnology Development (Foresight Institute 2006). The Guidelines were sponsored by the Foresight Institute, which was organized explicitly to provide a public forum for discussion of the risks and benefits of nanotechnology and to "pave the way" for its societal acceptance. Institute members include scientists, engineers, business people, investors, ethicists, policy makers and laypersons as well as firms; thus the organization represents a broad spectrum of stakeholders, interests, and opinions. It has been at the forefront of public discussions of nanotechnology risks and benefits.

The Foresight Guidelines (FGL) focus on "productive nanosystems." Currently, these systems form "a research oriented class of nanotechnology that will produce programmable, molecular-scale systems that make other useful nanostructured materials and devices." The Institute sees such systems as qualitatively different from nanomaterials, especially in their regulatory implications.

The FGL were initially developed at an expert workshop held in 1999; they have been revised multiple times through subsequent workshops, web-based community discussion, and other modalities. The version we discuss here is Draft Version 6, released in April 2006.

The FGL make an extraordinarily strong argument for the value of "soft law," especially as applied by researchers and firms themselves. They define soft law broadly, to include ethical behavior, good judgment, "professional guidelines and practices" based in science and knowledge of environmental and ethical issues, "cultural norms" of good practice that pervade scientific research, and professional ethics. According to the FGL, these norms are at least as effective as "hard law" in preventing unsafe practices and promoting action against them. Even in dealing with "rogue" actors who might abuse NT, the FGL suggest that much of the regulatory action can remain at the bottom or in the early stages of the regulatory pyramid. Two examples of soft regulation discussed in the FGL are moral and technical education, and the promotion of safe system designs that make abuse more difficult (in essence, embedding regulation within technology). The FGL include separate guidelines for different groups of actors working with nanotechnology; each guideline is cast as a self-assessment "scorecard," a notably soft form of implementation consistent with the FGL's reliance on ethics and professional norms.

To illustrate the types of norms contained in the FGL, consider the guideline for nanotechnology professionals (basically those involved in R&D). This guideline indicates that researchers should: adopt professional guidelines and ethical practices; engage in proactive stewardship by considering the possible negative consequences of any products subject to research and planning to prevent those consequences or minimize their harmful effects; conceive and develop products using total lifecycle analysis; quickly address any problems that arise; and practice inherently safe system design, avoiding the use of autonomous replicators. Similarly, the guideline for industry calls for proactive self-regulation, tailored to the specific risk profile of individual products and processes; this requires rigorous, balanced analyses of risks and benefits. This guideline also urges the use of inherently safe system design.

Other multi-stakeholder approaches to regulation are being pursued within standard-setting bodies. For example, the ISO established a Technical Committee on nanotechnologies in November 2005. In addition to the basic work of harmonizing terminology, nomenclature, and measurement, the Committee has established a working group, chaired by the USA, to consider "regulatory" standards pertaining to health, safety, and the environment.

Moving Up the Pyramid

In the medium term, as information about nanotechnology risks and benefits is gathered and disseminated and as society learns from diverse experiments in voluntary self- and multi-stakeholder regulation, it may be necessary to move toward greater governmental involvement, though still short of full-fledged command-and-control regulation, corresponding to the middle levels of the regulatory pyramid. This area of regulatory techniques is highly diverse, and it would be fruitless to speculate on the forms of government involvement that might be established. We briefly discuss one current example of government activity that captures the flexibility and public-private interchange that characterize this area of the pyramid.

Late in 2005, the US Environmental Protection Agency (EPA) launched a "collaborative" project to develop a Nanoscale Materials Stewardship Program. The EPA noted that it has statutory authority under the Toxic Substances Control Act to regulate many nanomaterials as "chemical substances," and that it would continue to implement that authority. Yet it sees the Stewardship Program as a valuable complement to its more traditional regulatory actions. The EPA invited broad stakeholder participation, and intends both to develop and to implement the program collaboratively and with public input.

Some potential components relate to the immediate need for better information about risks and benefits: EPA suggests that it may use the program to compile existing information from researchers and industry, and to encourage the development of testing procedures to produce additional information. More notably, EPA suggests that it may identify and encourage use of a basic set of risk management practices for R&D and commercial applications. To identify such practices, EPA will almost certainly rely on the existing experiments in product stewardship discussed here, such as the ED–DuPont agreement. Once EPA identifies and encourages use of those practices, they will no longer constitute pure self-regulation, but will remain soft law, complementing and to a considerable extent substituting for traditional hard regulation.

Building Up a Pyramid

The EPA Nanomaterials Stewardship Program reflects precisely the approach we espouse here: begin with information gathering and assessment, encourage experiments with self-regulation and multi-stakeholder norms, move gradually to greater governmental involvement to standardize, scale up and supervise voluntary programs, perform all these steps with high levels of transparency and participation, and over time build up to a regulatory end state that retains the best of these voluntary mechanisms at the base of the pyramid, along with formal regulation at the peak of the pyramid, as required.

For a valuable perspective on this approach, consider the actions the FGL suggest for regulators. Many of these guidelines speak to formal, mandatory regulation at the peak of the pyramid: for example, they call for regulators to be granted specific responsibilities and authorities, and for governments to designate a single regulatory entity to coordinate nanotechnology activities across agencies. More strikingly, however, the FGL devote considerable attention to voluntary actions at the base or in the early stages of the pyramid, and to relatively subtle interactions between regulators and the regulated, in the middle levels of the pyramid. For example, the FGL suggest that governments should:

a. Rely not only on "regulations," but also on "consensus standards promulgated by researchers, industry, or government." Whatever their source, regulatory norms should provide clear and specific guidelines and require the use of inherently safe systems.
b. Provide incentives for collaboration among firms, public interest groups and government on mechanisms for continuous improvement and the application of best practices in the handling of nanotechnology. This is a clear endorsement of the product stewardship approach.
c. Provide disincentives for those that fail to follow reasonable principles and guidelines. For example, such actors might be disadvantaged with regard to access to funding, designs, advanced nanotechnology capabilities cooperative market relationships, or collaborative relationships with public interest groups.
d. Enlist public interest groups and other actors in the international community to help prevent deliberate misuse through external verification.

Clearly, then, the FGL envision a regulatory pyramid, with the enforcement of existing and new laws through civil and criminal liability forming only the peak, called into action when lower-level measures prove to be insufficient. In our view, a similar approach applied over time would serve well both this technology and the interests of the public.

Implementing the Pyramid

A final question, of course, is how our incremental model—which in its early stages relies on decentralized measures taken by a range of private and public actors, including self-regulation and the development of multi-stakeholder norms and policies—can be adopted and implemented by the relevant decision makers. To begin, the preceding discussion has identified numerous spontaneous self-regulatory initiatives and multi-stakeholder programs that seek to identify and address the risks of nanotechnology. Civil society organizations continue to pressure researchers and business firms to act responsibly. Governments support research on nanotechnology impacts and risks, although most observers agree that more could be done. In short, our near-term approach is already being implemented, at least to a considerable degree.

Pure bottom-up initiatives are not, however, the only ways our approach can be implemented. In the near term, public authorities can signal concern that the risks of nanotechnology be adequately identified and addressed; such actions reinforce civil society pressure and remind researchers and industry that the state can intervene with mandatory regulation and a tit-for-tat strategy if necessary. Public authorities can also promote and support private initiatives and steer them in desirable directions, through means such as persuasion, financial incentives, publicity, and the implicit threat of regulation. Indeed, a few of our examples already reflect conscious action by government agencies or public officials; the EPA stewardship program is but one. Over time, of course, the role of public regulation will gradually increase.

In some areas more formal state involvement may be desirable even in the near term. For example, we have previously noted the need for transnational coordination of nanotechnology regulation [1] and urged a multilateral response [42]. One approach might be a "framework convention," adopted by states like any treaty. A framework convention could incorporate the kinds of near- and mid-term approaches we propose here: committing states to support research on the risks of nanotechnology, promote self-regulatory and multi-stakeholder initiatives, encourage and practice transparency and participation, exchange information and experiences, and consult on the form and transnational effects of any proposal for mandatory regulation [1]. Yet the "official" nature of a framework convention would be an important element in achieving cooperation among diverse and competitive states.

In general, however, formal authorization for an incremental approach will not be necessary. The approach may arise through spontaneous actions by researchers, industry and stakeholders, civil society pressure, government support and steering, or even formal government action. It is the incremental approach itself that is important, as we believe it is best suited to address the concerns identified in this paper, and ultimately to build an appropriate and effective regulatory system for the dynamic field of nanotechnology.

Conclusion

Nanotechnology presents enormous challenges to risk management, and existing risk models (including the new precautionary principle) will not be up to the challenge. A more incremental, multi-actor, and multi-component oversight model is needed for nanotechnology. The successful development of such a new risk management approach would not only facilitate the responsible development of nanotechnology, but will create a new precedent that could be used for other emerging technologies of the future. For example, looking to the recent past, a model similar to that proposed here might have helped smooth the introduction of genetically modified foods. As we look to other technology revolutions looming in the future, including emerging developments in telecommunication technologies, surveillance technologies, genetic enhancement, cognitive sciences, and many others, the need to develop new, better models for risk management (starting with nanotechnology) becomes all the more urgent.

References

1. Abbot KW, Marchard GE, Sylvester DJ (2006). A framework convention for nanotechnology? Environmental Law Reporter 36:10931–10942

2. Ackerman BA, Stewart RB (1985) Reforming environmental law. Stanford Law Rev 37:1333–1365
3. Austin C. Study shows Americans encouraged by prospects of nanotechnology (July 14, 2004). http://www.ncsu.edu/news/press_releases/04_07/211.htm. Cited May 4, 2007
4. Ayres I, Braithwaite J (1992) Responsive regulation: transcending the deregulation debate. New York, Oxford University Press
5. Babich A (2003) Too much science in environmental law. Columbia J Environ Law 28:119–184
6. Baram MS (1984) Alternatives to regulation: managing risks to health, safety, and the environment. Lexington Books, Lexington, MA
7. Breggin LK, Carothers L (2006) Governing uncertainty: the nanotechnology environmental, health, and safety challenge. Columbia J Environ Law 31:285–329
8. Bowman DM, Hodge GA (2007) A small matter of regulation: an international review of nanotechnology regulations. Columbia Science and Technology Law Review, 8:1–36
9. Bowman DM, Hodge GA (2006) Nanotechnology: mapping the wild regulatory frontier. Futures 38:1060–1073
10. Braithwaite V, Levi M (eds) (1998) Trust and governance. New York: Russell Sage
11. Colvin VL (2003) The potential environmental impact of engineered nanomaterials. Nat Biotechnol 21:1166–1170
12. Cross FB (1996) Paradoxical perils of the precautionary principle. Wash Lee Law Rev 53:851–925
13. Drexler E (1986) Engines of creation. Anchor, New York
14. Driesen DD (2005) Distributing the costs of environmental, health, and safety protection: the feasibility principle, cost-benefit analysis, and regulatory reform. Environ Aff 32:1–95
15. Environmental Defense, American Chemistry Council (2005) Joint statement of principles. Available at www.environmentaldefense.org/documents/4857_ACC-ED_nanotech.pdf
16. Environmental Defense, DuPont (2007) Nano risk framework. Available at www.nanoriskframework.com
17. ETC Group (2003) The big down: Atomtech—technologies converging at the nano-scale. http://www.etcgroup.org/upload/publication/171/01/thebigdown.pdf. Cited Sept. 29, 2007
18. Fiedler F, Reynolds G (1994) Legal problems of nanotechnology: an overview. S Cal Interdisc LJ 3:593–629
19. Florini K, Walsh S, Balbus JM, Denison R (2006) Nanotechnology: getting it right the first time. Nanotechnol Law Bus 3:39–53
20. Forrest D (1989) Regulating nanotechnology development, Foresight Nanotech Institute, March 23, 1989. http://www.foresight.org/nano/Forrest1989.html. Cited September 29, 2007
21. Friends of the Earth (Australia) (2007). Who's afraid of the precautionary principle? http://nano.foe.org.au/node/186. Cited Sept. 29, 2007
22. Greenwood M (2007) Thinking big about things small: creating an effective oversight system for nanotechnology. Woodrow Wilson International Center for Scholars, Washington, DC
23. Gregory R, Flynn J, Slovic P (1995) Technological stigma. American Scientist 83:220–223
24. Gwinn MR, Vallyathan V (2006) Nanoparticles: health effects—pros and cons. Environ Health Perspect 114:1818–1825
25. Holm S, Harris J (1999) Precautionary principle stifles discovery (letter). Nature 400:398
26. International Risk Governance Council (IRGC) (2006). White paper on nanotechnology risk governance (Geneva, IRGC). http://www.irgc.org/irgc/_b/contentFiles/IRGC_white_paper_2_PDF_final_version.pdf. Cited September 29, 2007
27. International Center for Technology Assessment (ICTA) et al. (2007) Principles for the oversight of nanotechnologies and nanomaterials. http://www.icta.org/doc/Principles%20for%20the%20Oversight%20of%20Nanotechnologies%20and%20Nanomaterials_final.pdf. Cited Sept. 29, 2007
28. Johnson BB (1993) Advancing understanding of knowledge's role in lay risk perception. Risk Issues Health Saf 4:189–212
29. Joy B (2000) Why the future doesn't need us. Wired, August 4, 2000. http://www.wired.com/wired/archive/8.04/joy_pr.html. Cited Sept. 29, 2007
30. Jukes TH (1986) Frost resistance and Pseudomonas. Nature 319:617
31. Kahan D, Slovic P, Braman D, Gastil J, Cohen GL (2007) Affect, values, and nanotechnology risk perceptions: an experimental investigation. http://papers.ssrn.com/sol3/papers.cfm?abstract_id=968652. Cited May 3, 2007
32. Kahneman D, Slovic P, Tversky A (eds) (1985) Judgement under uncertainty: heuristics and biases. Cambridge University Press, Cambridge
33. Kasperson JX et al (2003) The social amplification of risk: fifteen years of research and theory. In: Pidgeon N, Kasperson RE, Slovic P (eds) The social amplification of risk. Cambridge University Press, Cambridge
34. Kolata G (2001) Scientists debate what to do when findings aid an enemy. New York Times Sept. 25, 2001:D1
35. Kuran T, Sunstein CR (1999) Availability cascades and risk regulation. Stanford Law Rev 51:683–768
36. Kuzma J (2007) Moving forward responsibly: oversight for the nanotechnology–biology interface. J Nanopart Res 9:165–182
37. Lin AC (2007) Size matters: regulating nanotechnology. Harvard Environ Law Rev 31:349
38. Lin-Easton PC (2001) It's time for environmentalists to think small—real small: a call for the involvement of environmental lawyers in developing precautionary policies for molecular nanotechnology. Georget Int Environ Law Rev 14:106–134
39. Mandel G (2005) Technology wars: the failure of democratic discourse. Mich Telecommun Law Rev 11:117–190
40. Marchant GE (2003) From general policy to legal rule: the aspirations and limitations of the precautionary principle. Environ Health Perspect 111:1799–1803
41. Marchant GE, Mossman KL (2004) Arbitrary and capricious: the precautionary principle in the European Union courts. AEI, Washington, DC
42. Marchant GE, Sylvester D (2006) Transnational models for regulation of nanotechnology. J Law Med Ethics 34:714–725
43. Maynard AD (2006a) Nanotechnology: a research strategy for addressing risk. Woodrow Wilson International Center for Scholars, Washington, DC
44. Maynard AD (2006b) Safe handling of nanotechnology. Nature 444:267–269
45. McCubbin PR (2005) The risk in technology-based standards. Duke Environ Law Policy Forum 16:1–56

46. Morrissey SR (2007) Understanding nanotechnology. Chemical & Engineering News April 16, 2007:35–37
47. Nel A, Xia T, Maedler L, Li N (2006) Toxic potential of materials at the nanolevel. Science 311:622–627
48. Oberdörster G et al (2005) Principles for characterizing the potential human health effects from exposure to nanomaterials: elements of a screening strategy. Particle & Fibre Toxicology 2:8
49. Paddock L (2006) Keeping pace with nanotechnology: a proposal for a new approach to environmental accountability. Environ Law Report 36:10943–10952
50. Peter D. Hart Research Associates, Inc. (2006) Report findings. http://www.nanotechproject.org/file_download/98. Cited May 4, 2007
51. Phoenix C, Drexler E (2004) Safe exponential manufacturing. Nanotechnology 15:869–872
52. Pidgeon N, Kasperson RE, Slovic P (eds) (2003) The social amplification of risk. Cambridge University Press, Cambridge
53. Posner R (2004) Catastrophe. Oxford, University Press
54. Rachlinski J (2003) The uncertain psychological case for paternalism. Northwest Univ Law Rev 97:1165–1225
55. Rajeski D (2004) The next small thing. The Environmental Forum, March/April, pp. 42–49
56. Renn O, Roco MC (2006) Nanotechnology and the need for risk governance. J Nanopart Res 8:153–191
57. Reynolds GH (2003) Nanotechnology and regulatory policy: three futures. Harv J Law Technol 17:179–209
58. Sandin P (1999) Dimensions of the precautionary principle. Hum Ecol Risk Assess 5:889–907
59. Scientific Committee on Emerging and Newly Identified Health Risks (SCENIHR), European Commission (2007) Opinion on the appropriateness of the risk assessment methodology in accordance with the technical guidance documents for new and existing substances for assessing the risks of nanomaterials. http://ec.europa.eu/health/ph_risk/committees/04_scenihr/docs/scenihr_o_004c.pdf. Cited Sept. 29, 2007
60. Segal S (2004) Environmental regulation of nanotechnology: avoiding big mistakes for small machines. Nanotechnol Law Bus 1:290–303
61. Shapiro SA, McGarity TO (1991) Not so paradoxical: the rationale for technology-based regulation. Duke Law J 1991:729–752
62. Sheetz T et al (2005) Nanotechnology: awareness and societal concerns. Technol Soc 27:329–345
63. Slovic P (2000) The perception of risk. London, Earthscan
64. Slovic P (2004) What's fear got to do with it? It's affect we need to worry about. Miss Law Rev 69:971–990
65. Slovic P, Peters E, Finucane ML, MacGregor DG (2005) Affect, risk, and decisionmaking. Health Psychol 24:S35–40
66. Sunstein CR (1991) Administrative substance. Duke Law J 1991:607–646
67. Sunstein CR (2002a) The laws of fear. Harvard Law Rev 115:1119–1168
68. Sunstein CR (2002b) Probability neglect: emotions, worst-cases, and law. Yale Law J 112:61–108
69. Sunstein CR (2003a) Beyond the precautionary principle. Univ PA Law Rev 151:1003–1058
70. Sunstein CR (2003b) Hazardous heuristics. Univ Chicago Law Rev 70:751–781
71. Sunstein CR (2004) Precautions against what? The availability heuristic and cross-cultural risk perception. Alabama Law Rev 57:75–101
72. Sunstein CR (2006) Irreversible and catastrophic. Cornell Law Rev 91:841–897
73. Sweet L, Strohm B (2006) Nanotechnology—life-cycle risk management. Hum Ecol Risk Assess 12:528–551
74. Sylvester D, Lohr S (2005) Counting on confidentiality: legal and statistical approaches to federal privacy law after the USA Patriot Act. Wisconsin Law Rev 2005:1036–1138
75. Thayer AM (2006) Chance of a lifetime. Chem Eng News 1:10–18
76. Tversky A, Kahneman D (1982) Judgment under uncertainty: heuristics and biases. In Kahneman et al. (eds) Judgment under uncertainty: heuristics and biases, pp 3–19
77. U.S. Environmental Protection Agency (EPA) (2007) Science Policy Council, Nanotechnology White Paper, EPA 100/B-07/001. http://es.epa.gov/ncer/nano/publications/whitepaper12022005.pdf. Cited Sept. 29, 2007
78. Wagner WE (2000) The triumph of technology-based standards. Univ Ill Law Rev 2000:83–113
79. Weckert J, Moor J (2006) The precautionary principle in nanotechnology. Int J Appl Philos 20:191–204
80. Wejnert J (2004) Regulatory mechanisms for molecular nanotechnology. Jurimetrics J 44:323–350
81. Wexler L (2006) Limiting the precautionary principle: weapons regulation in the face of scientific uncertainty. U. C. Davis Law Rev 39:459–527
82. Wiedemann PM, Schutz H (2005) The precautionary principle and risk perception: experimental studies in the EMF area. Environ Health Perspect 113:402–405
83. Wilson RF (2006) Nanotechnology: the challenge of regulating known unknowns. J Law Med Ethics 34:704–713