Social and Ethical Implications of Artificial Cells

Mark Bedau* & Mark Triant
Department of Philosophy, Reed College

Rough Draft 050208

Abstract: Artificial cells are microscopic self-organizing and self-replicating autonomous entities built from simple organic and inorganic substances. A number of research efforts are currently aimed at creating artificial cells within the next generation. The ability to create artificial cells would have many social and economic benefits but it would also raise significant social risks. This paper reviews the social and ethical implications of artificial cells. We first respond to the objections that creating artificial cells would be wrong because it is unnatural, it commoditizes life, it fosters reductionism, or it is playing God. Then we raise skepticism about the ultility of decision theory for resolving whether to create artificial cells. Finally, we consider two principles for acting in the face of uncertain risks—the Doomsday Principle and the Precautionary Principle—and find them wanting. We end by proposing a new method—which we dub the "Cautious Courage" Principle—for deciding whether and how to develop artificial cells. Our conclusions generalize to analogous debates concerning related new technologies, such as genetic engineering and nanotechnology.

Keywords: artificial cell, genetic engineering, nanotechnology, Precautionary Principle, Frankencell, decision theory.

Introduction

A striking biotechnology research program has been quietly making incremental progress for the past generation, but it will soon become public knowledge. One sign of this is a recent article in the widely distributed Sunday supplement *Parade Magazine*, in which one could read the following prediction:

> Tiny robots may crawl through your arteries, cutting away atherosclerotic plaque; powerful drugs will be delivered to individual cancer cells, leaving other cells undamaged; teeth will be self-cleaning. Cosmetically,

---

* To whom correspondence should be directed: Reed College, 3203 SE Woodstock Blvd., Portland OR 97202, USA; (503) 517-7337; bedau@reed.edu

> you will change your hair color with an injection of nanomachines that circulate through your body, moving melanocytes in hair follicles…

This may sound incredible and it is certainly science fiction today, but scientists working in the field believe that within the next decade or so the basic technology underlying this prediction will exist. That technology could be called "artificial cells."

Artificial cells are microscopic self-organizing and self-replicating autonomous entities created artificially from simple organic and inorganic substances. The *Parade* article quoted above was written by Michael Crichton as an effort to explain the science underlying his most recent international best-seller, *Prey* (Crichton 2002). Crichton's book imagines the disastrous consequences of artificial cell commercialization gone awry (humans are the prey of swarms of artificial cells). Although one can question many scientific presuppositions behind Crichton's imagined artificial cells (Dyson 2003), the underlying research is proceeding apace, and the potential risks and benefits to society are enormous. And the soon-to-be released movie based on *Prey* will make the wider public much more aware of these issues. So it is appropriate to ask whether we as a society are ready for this future.

Society was certainly unprepared for Dolly, the Scottish sheep cloned from an adult udder cell, when her picture was first splashed across the front pages of newspapers around the world. President Clinton immediately halted all federally funded cloning research in the United States (Brannigan 2001), and polls revealed that ninety percent of the public favored a ban on human cloning (Singer 1998). This paper aims to start a critical and informed public discussion about the implications of artificial cells, so that the announcement of the first artificial cells will not provoke similar knee-jerk reactions. Although the probable eventual social ramifications of artificial cells are quite significant, there is little risk that we will confront them this year or next, so we have time for thoughtful reflection and informed discussion. But the clock is ticking.

We aim here to review the main issues that the public discussion of artificial cells should address. Our perspective on this social and ethical landscape has been influenced by the public controversies about genetic engineering and nanotechnology. Indeed, our conclusions about artificial cells could equally be applied to these related technologies, but those parallels will be largely silent here. After explaining the trends in contemporary artificial cell research and outlining their risks and benefits, we will discuss the main strategies for deciding whether and under what conditions to create them. One set of considerations focuses on intrinsic features of artificial cells. These include the suggestions that creating artificial cells is unnatural, that it commoditizes life, that it fosters a reductionistic perspective, that it is playing God, and that we should use religious texts as sources of authority. We find all these considerations unconvincing.

The alternative approaches focus on the consequences of creating artificial cells. Utilitarianism and decision theory promise scientifically objective and pragmatic methods for deciding what course to chart. Although we agree that consequences are of primary importance, we are skeptical whether utilitarianism and decision theory can provide much practical help because the consequences of creating artificial cells are so uncertain. What we need is some principle for choosing the best course of action in the face of this uncertainty. In this kind of setting some people advocate following what we call the Doomsday Principle, but we find this principle to be incoherent. Given the possibility of significant risk and scientific uncertainty about its nature and scope, an increasing number of people are turning for guidance to the Precautionary Principle, but we also find this principle unattractive. We conclude that the best guide for deciding whether and how to develop artificial cells given their uncertain consequences is what we call the "Cautious Courage" Principle.

What are artificial cells?

Artificial cells would be microscopic aggregations of simple organic and inorganic molecules that construct and replicate themselves autonomously. Although artificial, for all intents and purposes they would be alive, for they would spontaneously regenerate and repair themselves and they would adapt and evolve in an open-ended fashion through natural selection.

There are two main motivations behind artificial cell research. One is pure science. If one could make artificial cells from scratch, especially using materials or methods that are not employed by natural forms of life, one would have dramatic proof that one fully grasps the essential molecular foundations of living systems. But artificial cells also have a practical appeal. Natural cells are much more complicated than anything yet produced by man, and many people believe that the next watershed in intelligent machines depends on bridging the gap between non-living and living matter (e.g., Brooks 2001). So, making artificial cells that organize and sustain themselves and evolve to their environment would open the door to creating technologies with the impressive capacities of living systems, and the promise of harnessing those capacities for social and economic gain is quite attractive. The next section develops these promises along with their associated risks.

Nobody has yet created an artificial cell, but research aimed at this goal is actively under way. Two main approaches are being pursued. One is the "top-down" strategy of modifying existing life forms. J. Craig Venter and Hamilton Smith recently publicized their intention to create a partly man-made artificial cell, with $3,000,000 in support from the US Department of Energy (Gillis 2002), and the Department of Energy has since increased its support by an order of magnitude (Smith, personal communication). Venter and Smith are simplifying the genome of the simplest existing cell with the smallest genome: *Mycoplasma genitalium*. Sequencing showed that 580 kb genome of *M. genitalium* contained only 480 protein-encoding genes and 37 genes for RNA species, for a total of 517 genes (Fraser et al. 1995). Random shotgun gene knockout experiments subsequently determined that approximately 300 of those genes were essential for *M. genetalium* to survive and reproduce in laboratory conditions (Hutchison et al. 1999). Venter and Smith plan to use existing DNA synthesis technology to construct an entirely artificial chromosome that contains the 300 genes needed by *M. genitalium*. They will then remove the genetic material from an existing *M. genitalium* bacterium and insert their artificial chromosome. If the bacterium cytoplasm can be coaxed to express that synthetic DNA, it will grow and reproduce and thus start a lineage of bacteria that has never existed anywhere before (Gillis 2002). One perfected, it will be possible to repeat the process but add further genes that perform various useful functions, such as generating hydrogen for fuel or breaking down carbon dioxide that is polluting the atmosphere.

The top-down approach to making artificial cells is a logical extension of our experience with genetic engineering over the past thirty years. It has the virtue that it can simply adopt as a black box all the biological wisdom produced by million of years of evolution and currently embodied in biochemistry. It has the corresponding disadvantage that its insights will be constrained by the same evolutionary contingencies. If there are essentially different solutions to the fundamental problems that simple cells must face, they will not be discovered by building on *Mycoplasma*'s minimal genetic requirements. An alternative "bottom-up" strategy sacrifices the head-start provided by existing life form and attempts to create artificial cells *de novo*, entirely from non-living materials. The goal is to create a structure of molecules that is simple enough to form by self-assembly but complex enough to reproduce itself and evolve without using any products of pre-existing life forms such as protein enzymes. The advantage of this approach is that, freed from the contingent constraints within existing life forms, it can explore a much broader canvas of biochemistries and thus can eventually deliver a more fundamental understanding of the molecular mechanisms required for life.

Most bottom-up artificial cell research incorporates key elements of existing cellular biology. The three primary elements are the formation of enclosed membranes from amphipathic lipids, the replication of information-carrying molecules like DNA or RNA by a templating process, and the harvesting of chemical energy to construct cellular structures from small molecules that are transported across the membrane from the environment into the cell (Szostak 2001, Pohorille and Deamer 2002). Nucleic acids like DNA and RNA will replicate under appropriate laboratory conditions, but some kind of container is needed to keep reactants physically proximate, and a population of such containers is required to allow competing reaction systems to evolve by natural selection. So an artificial cell will consist of three integrated systems: an informational chemistry (genes) and an energy-harvesting chemistry (metabolism) both enclosed within a self-assembling container (cell wall).

The spontaneous growth and replication of lipid bilayer vesicles has already been demonstrated in the laboratory (Walde et al. 1994), as has the synthesis of information-carrying molecules inside lipid vesicles (Pohorille and Deamer 2002, Oberholzer et al. 1995). Furthermore, self-replicating RNA molecules can spontaneously become encapsulated in self-replicating lipid vesicles. Much effort is directed at finding an appropriate RNA replicase, i.e., an RNA molecule that functions both as a repository for genetic information and as an enzyme directing replication of that very genetic material. Finding a molecule that performed both critical functions would vastly simply the artificial cell's biochemistry. A promising initial step toward this goal includes using *in vitro* selection to find ribozymes—RNA enzymes for breaking down RNA—that act as primitive polymerases—enzymes for building up RNA (Ekland and Bartel 1996, Bartel and Unrau 1999). One key remaining challenge is to find an RNA replicase that is efficient enough to accurately replicate itself, and another is for this function to work quickly inside a self-replicating vesicle. A third challenge is to couple the functioning of the genetic, metabolic, and container chemistries so that the entire system evolves as a unit. Rasmussen et al. (2003) have recently proposed an especially simple artificial cell in which these three chemical systems are explicitly coupled in a novel simple fashion. In this design PNA[1] chemistry (Nielsen et al. 1991) replaces RNA chemistry and lipid micelles replace lipid vesicles. As with all other artificial cell projects, Rasmussen's new design has not yet been experimentally realized.

Just as the top-down approach to creating artificial cells can be considered the logical extension of genetic engineering, the bottom-up approach can be considered a branch of nanotechnology. Most treatments of nanotechnology consider tiny versions of familiar machines with molecular-scale gears and levers. Artificial cells would be different because their chemical mechanisms would more reflect the molecular biology world. But traditional discussions of nanotechnology mention self-assembling, self-replicating, and even evolving molecular machines, and artificial cells would fit that bill. The potential benefits of both genetic engineering and nanotechnology have received a lot of hype, and their dangers have also generated much controversy. Artificial cells would raise similar risks and benefits. To this topic we turn next.

Potential risks and benefits

The ability to make artificial cells is a scientific milestone of immense proportions. Achieving it would mark a profound understanding of the biochemical systems that embody life, and it would also provide a fast track toward a series of fundamental scientific insights. In addition to these scientific benefits, many technological, economic, and social benefits would follow because artificial cells would be a threshold technology that opens the door to new kinds of applications. Pohorille and Deamer (2002) note many pharmacological and medical diagnostic functions that artificial cells could perform. One application is drug-delivery vehicles that activate a drug in response to an external signal

[1] PNA is peptide nucleic acid, an analog of DNA in which the backbone is a pseudopeptide rather than a sugar.

produced by target tissues. Another function is microencapsulation of proteins, such as artificial red blood cells that would contain enhanced hemoglobin or special enzymes. A third application is multi-function biosensors with activity that can be sustained over a long period of time. After reviewing these examples, Pohorille and Deamer (2002, 128) conclude:

> Artificial cells designed for specific applications offer unprecedented opportunities for biotechnology because they allow us to combine the properties of biological systems such as nanoscale efficiency, self-organization and adaptability for therapeutic and diagnostic applications. ... it will become possible to construct communities of artificial cells that can self-organize to perform different tasks and even evolve in response to changes in the environment.

It is easy to expand the list of artificial cell applications with possibilities ranging from molecular chemical factories and the metabolism of environmental toxins to defenses against bioterrorism and a cure for heart disease (artificial cells that flow through our bloodstream ingesting atherosclerotic plaque).

Artificial cells are the tip of an iceberg of opportunities provided by living technology. Living systems have a remarkable range of distinctive useful properties, including autonomous activity, sensitivity to the environment and robustness in the face of environmental change, and automatic adaptation and ongoing creativity. There is increasing need felt for technology that has these features, i.e., for what could be called "living" technology. Conventional engineering is hitting a complexity barrier because it produces devices that are non-adaptive, brittle and costly to redesign. There is growing recognition that the creation of truly intelligent and adaptive physical artifacts depends on bridging the gap between nonliving and living matter (Brooks 2001). The only physical entities that now exhibit self-repair, open-ended learning and spontaneous adaptability to unpredictably changing environments are forms of life, so it is plausible to conclude that the future of intelligent, autonomous, automatically adaptive systems depends on the creation of living technology. And artificial cells would be the first concrete step down this path.

Artificial cells also raise significant social and ethical worries. Ethical issues related to creation of artificial forms of life have a long history, dating back at least to the artificial production of urea, the first man-made organic compound. Concerns about nanostructures proliferating in natural environments were expressed in the nanotechnology community a decade ago (Merkle 1992), and a recent cautionary piece by Bill Joy in *Wired* about the combination of nanotechnology with genetic engineering (Joy 2000) sparked extensive commentary on the web. Similar public concerns have surfaced over the minimal cell research of Venter and Smith, which the popular press has dubbed "Frankencell.". This public outcry prompted Venter to halt research while an independent panel of ethicists and religious leaders reviewed the ethics of synthesizing artificial cells (Cho et al. 1999). When Venter and Smith announced the resumption of their artificial cell project (Gillis 2002), they attracted quick commentary on editorial pages (e.g., Mooney 2002). And as it happened, within a week of the Venter/Smith announcement Michael Crichton published *Prey*. Events like these are increasingly bringing the social and ethical implications of artificial cells to the attention of the general public.

One of the most wide-spread worries about artificial cells is their potential threat to human health and the environment. Bill Joy's *Wired* article did not mention artificial cells by name but he worried about essentially the same thing: molecular machines that had the ability to reproduce themselves and evolve uncontrollably. Referring to the dangers of genetic engineering and Eric Drexler's (1986) warnings about the dangers of self-reproducing nanotechnology, Joy concludes that "[t]his is the first moment in the history of our planet when any species, by its own voluntary actions, has become a danger to itself—as well as to vast numbers of others," and he describes one key problem thus:

> As Drexler explained: "Plants" with "leaves" no more efficient than today's solar cells could out-compete real plants, crowding the biosphere with an inedible foliage. Tough omnivorous "bacteria" could out-compete real bacteria: They could spread like blowing pollen, replicate swiftly, and reduce the biosphere to dust in a matter of days. Dangerous replicators could easily be too tough, small, and rapidly spreading to stop—at least if we make no preparation. We have trouble enough controlling viruses and fruit flies.

To the health and environmental risks of artificial cells Joy adds the threat of new and vastly more lethal forms of bioterrorism.

These dangers posed by artificial cells stem from two key features. First, since they would be self-replicating, any danger that they pose has the potential to be magnified on a vast scale as the artificial cells proliferate and spread around the globe. Second, because they would be evolving, their properties could change in ways that we never anticipated. For example, they could evolve new ways of competing with existing life forms and new ways to evade our eradication methods. This potential for open-ended evolution makes the long-term consequences of creating them extremely unpredictable. Much of the positive potential of artificial cells stems from their ability to replicate and evolve, and the very same power raises the specter of life run amok.

One can envision strategies for coping with these dangers, of course. One is simply to strictly contain artificial cells in confined areas and not let them escape into the environment. This is a familiar way of addressing dangerous pathogens found in nature (e.g., the Ebola virus). Another method is to take advantage of the fact that they are artificially created and build in mechanisms that cripple or control them. A common proposal is to make them dependant on a form of energy or raw material that can be blocked or that is normally unavailable in the environment, so that they would survive in the wild only if and when we allow. Another common suggestion is to make artificial cells have a strictly limited life span, so that they die before they could do any harm. It might even be possible to engineer them so that they remain alive only upon receiving regular external signals or they die when an externally triggered on/off switch is tripped.[2] A further form of crippling would be to block their ability to evolve. It is thought that the mixing of genetic material brought about by sexual reproduction is the primary driving force of evolution, so one could hamper artificial cell evolution by preventing sexual reproduction. Merkle (1992) has also proposed encrypting artificial genomes in such a way that any mutation would render all the genetic information irretrievable. One final suggestion is to put a unique identifier (a genetic "bar code") inside each artificial cell, so that we can track down the source of any artificial cell that does damage and seek redress from the responsible parties.

Such measures would not placate the concerns about artificial cells, though, for the safeguards are fallible and costly. No containment method is perfect, and more effective containment is more expensive. Another cost is that containment significantly hampers research, thus impeding our knowledge of how artificial cells work and what beneficial uses they might have. And many of the potential benefits of artificial cells involve them inhabiting our environment or even living inside our bodies, and all such applications would be off the table if they were to be isolated inside strict containment devices. Furthermore, methods for crippling or controlling artificial cells could well be ineffective. When humans have introduced species into foreign environments, it often proves difficult to control their subsequent spread. More to the point, viruses and other pathogens are notorious for evolving ways to circumvent our methods of controlling or eradicating them, and artificial cells would experience significant selection pressure to evade our efforts to cripple or control them. Another kind of social cost of crippling artificial cells is that this would defeat many potential benefits of living technology. The appeal of living technology includes taking advantage of life's robustness and its flexible capacity to adapt to environmental contingencies, and crippling life would sacrifice this capacity.

Creating artificial cells would dramatically alter our world, and the potential upside and downside are both quite large. Artificial cells could enable many impressive benefits for human health, the environment, and defense, and they would dramatically accelerate basic science. But they could also significantly threaten human health and the environment and enable new forms of bioterrorism. In addition, creating artificial cells could change public perception about life and its mechanistic foundations in a way that would

---

[2] For example, there is evidence that magnetic fields can be used to turn genes on and off (Stikeman 2002).

undermine some entrenched cultural institutions and belief systems. Given these significant consequences, how should we decide how to proceed down this path? To this issue we turn next.

The intrinsic value of life

Arguments about whether it is right or wrong to develop a new technology can take one of two forms (Reiss and Straughan 1996, ch. 3; Comstock 2000, chs. 5-6). *Extrinsic* arguments are driven by the technology's consequences. A technology's consequences often depend on how it is implemented, so extrinsic arguments do not usually produce blanket evaluation of a technology. Presumably, any decision about creating a new technology should weigh its consequences, perhaps along with other considerations. Evaluating extrinsic approaches to decisions about artificial cells is the subject of the two sections following this one. *Intrinsic* argument for or against a new technology are driven by the nature of the technology itself, yielding conclusions pertinent to any implementation of it. The advances in biochemical pharmacology of the early 20th century and more recent developments in genetic engineering and cloning have been criticized on intrinsic grounds, for example. Such criticisms include the familiar injunctions against playing God, tampering with forces beyond our control, or violating nature's sanctity, and the prospect of creating artificial cells raises many of same kinds of intrinsic concerns. In this section we will address intrinsic arguments about whether there is something intrinsically objectionable about creating artificial life forms.

Reactions to the prospect of synthesizing new forms of life range from fascination to skepticism and even horror. Everyone should agree that the first artificial cell will herald a scientific and cultural event of great significance, one that will force us to reconsider our place in the cosmos. But what some would hail as a technological milestone, others would decry as a new height of scientific hubris. The "Frankencell" tag attached to Venter's minimal genome project reveals the uneasiness generated by this prospect. So it is natural to ask whether taking this big step would be crossing some forbidden line. In this section, we will examine four kinds of intrinsic objections to the creation of artificial cells, all of which frequently arise in debates over genetic engineering and cloning. These arguments all stem from the notion that life has a certain privileged status and should in some respect remain off-limits from human intervention and manipulation.

One objection against creating artificial cells is simply that doing so would be unnatural and, hence, unethical. The force of such arguments depends on what is meant by "unnatural" and why the unnatural is wrong. At one extreme, one could view all human activity and its products as natural since we are part of the natural world. But then creating artificial cells would be natural, and this objection would have no force. At the other extreme, one could consider all human activity and its products as unnatural, defining the unnatural as what is independent of human influence. But then the objection would deem all human activities to be unethical, which is absurd. So the objection will have any force only if "natural" is interpreted in such a way that we can engage in both natural and unnatural acts and the unnatural acts are intuitively wrong. But what could that sense of "natural" be? One might consider it "unnatural" to intervene in the workings of other life forms. But then the unnatural is not in general wrong; far from it. For example, it is surely not immoral to hybridize vegetable species or to engage in animal husbandry. And the stricture against interfering in life forms is not something that arises especially for humans, for it is not wrong to vaccinate one's children. So there is no evident sense of "unnatural" in which artificial cells are unnatural and the unnatural is intrinsically wrong.

Another objection is that to create artificial life forms would lead to commoditizing life, which is immoral.[3] Underlying this objection is the notion that living things have a certain sanctity or otherwise demand our respect, and that creating them undermines this respect. The commoditization of life is seen as analogous to the commoditization of persons, a practice most of us would find appalling. By producing living artifacts, one might argue, we would come to regard life forms as one among our products and so valuable only insofar as they are useful to us. This argument is easy to sympathize with, but is implausible when followed to its conclusion. Life is after all one of our most abundant commodities. Produce, livestock, vaccines, and pets are all examples of life forms that are bought and sold every day. Anyone who objects to the commoditization of an artificial single-celled organism should also object to the commoditization of a tomato. Furthermore, creating, buying, and selling life forms does not prevent one from respecting those life forms. Family farmers, for example, are often among those with the greatest reverence for life.

The commoditization argument reflects a commonly held sentiment that life is special somehow, that it is wrong to treat it as we treat the rest of the material world. It can be argued that while it is not inherently wrong to commoditize living things, it is still wrong to create life from nonliving matter because doing so would foster a reductionistic attitude toward life, which undermines the sense of awe, reverence, and respect we owe it[4]. This objection doesn't exactly require that biological reductionism be false, but merely that it be bad for us to view life reductionistically. Of course, it seems somewhat absurd to admit the truth of some form of biological reductionism while advocating an antireductionist worldview on moral grounds. If living things are really irreducible to purely physical systems (at least in some minimal sense), then creating life from nonliving chemicals would presumably be impossible, so the argument is moot. On the same coin, if living things are reducible to physical

---

[3] For discussions of this argument as applied to other forms of biotechnology, see Kass (2002), ch. 6, and Comstock (2000), pp. 196-198.

[4] See and Dobson (1995) and Kass (2002), ch. 10 for discussions of this objection in other contexts.

systems, it is hard to see why fostering reductionistic beliefs would be unethical. It is by no means obvious that life *per se* is the type of thing that demands the sense of awe and respect this objection is premised upon, but even if we grant that life deserves our reverence, there is no reason to assume that this is incompatible with biological reductionism. Many who study the workings of life in a reductionistic framework come away from the experience with a sense of wonder and an enhanced appreciation and respect for their object of study. Life is no less amazing by virtue of being an elaborate chemical process. In fact, only after we began studying life in naturalistic terms have we come to appreciate how staggeringly complex it really is.

Inevitably, the proponents and eventual creators of artificial cells will have to face up to the accusation that what they are doing amounts to *playing God*[5]. Animals have been cloning themselves for millions of years, and viruses splice DNA all the time, but to create life from non-life is to accomplish what has only occurred once before in the history of the planet. Though accusations of hubris may come as no surprise, it is unclear what "playing God" really amounts to, and why it is considered immoral. The playing-God argument can be hashed out in one of two ways: it could be an observation that by creating life from scratch, we are opening up a very large can of worms that we simply aren't prepared to deal with, or it could amount to the claim that, for whatever reason, creating life from scratch represents a line that humans are simply not meant to cross. The former construal, because it is couched in terms of the potentially disastrous consequences of scientific progress, poses the playing-God accusation as an extrinsic objection to the creation of artificial cells, and will be treated in a later section.

Regarding the latter construal, we might ask ourselves, so what? Implicit in accusing a person of playing God is the understanding that doing so is wrong, but setting aside the *consequences* of playing God, we are left with little to justify this assumption. The term "playing God" was popularized in the early 20th century by Christian Scientists in reaction to the variety of advances in medical science taking place at the time. With the help of new surgical techniques, vaccines, antibiotics, and other pharmaceuticals, the human lifespan began to extend, and many fatal or otherwise untreatable ailments could now be easily and reliably cured. Christian Scientists made it their policy to opt themselves and their families out of medical treatment because they believed it wrong to "play God" – healing the ill was God's business, not ours. Yet if a person living today were to deny her ailing child medical attention on the grounds that doing so is playing God, we would be appalled. If saving a life through medicine is playing God, then it is not only morally permissible to play God; it is sometimes morally required of us. And if treating the ill doesn't count as playing God, how

---

[5] For discussions of the playing-God objection as it has entered into the genetic engineering controversy, see Comstock (2000), pp.184-185 and Reiss & Straughan (1996), pp. 79-80, 121.

do we know what does? Dogmatic pronouncements hardly amount to justified arguments, and if theological authorities disagree on the matter, who if any should be believed?

Questions surrounding the playing God argument are related to the more general question of what role religious authority should be allowed to play in influencing scientific policy decisions. Though religious doctrine will surely be invoked in future discussions of artificial cell science, it is questionable whether it could be helpful as a guide to policymakers. In the first place, religious nonpartisanism is generally considered a keystone of any modern democratic state. But more importantly, religious dogma remains fundamentally unequipped to provide answers to the kinds of issues this paper addresses. One Christian might consider the creation of artificial cells an unconscionable act of hubris, while another might see it as just another extension of the capacities bestowed on us by God, and there would be no way to resolve this disagreement, because nowhere in any sacred text is there any mention of creating artificial life forms, "playing God", or violating the natural order. One's stance on the issue, if influenced by spiritual beliefs, remains a matter of individual interpretation. To let any one spiritual perspective guide policy decisions would serve only to divide and alienate those involved in determining policy.

All of the intrinsic objections raised against the creation of artificial cells we have addressed in this section have proved after some scrutiny to be vague, overly simplistic, or ill-conceived. These kinds of arguments are raised with surprising frequency in debates over the ethics of new movements in medicine and biotechnology, and so demand attention in the context of the artificial cell controversy as well. Insofar as artificial cell technology poses ethical concerns, we have found most commonly raised intrinsic concerns to be inadequate or irrelevant as arguments to artificial cells. We have also argued that religious doctrine, though by no means irrelevant to one's ethical conduct, should not have a place in determining public policy decisions of the kind addressed in this paper. This leaves us to appeal to the imagined and expected *consequences* of artificial cell technology in determining how we should proceed.

Evaluating the consequences

The question of whether to *create* artificial cells is just one of many decisions our scientists and policymakers face. To what extent should artificial cells be developed for commercial, industrial, or military applications? How strictly should they be regulated? Before any of these decisions can be made, we must decide on how to make them. We have already presented a variety of reasons for pursuing artificial cell research and development, as well as reasons to doubt the safety of artificial cell technology. Especially this early in the decision-making process, there is no obvious way of knowing whether the speculated benefits are worth pursuing in light of the speculated risks. Ruling out any intrinsic ethical qualms against creating artificial cells, the choices we make will be for the most part a matter of how we weigh the risks and benefits of artificial cell science, and what strategies and principles we employ in making these decisions in the face of uncertain consequences. These issues will be discussed in the remainder of the paper.

The utilitarian calculus is one obvious tool we might employ in assessing a course of action in light of its consequences: possible actions are measured according to the overall "utility" they would produce (e.g., the aggregate happiness or well-being that would result), where the course of action we *should* pursue is the one that would produce the greatest utility. There are of course problems with any utilitarian calculus, one being the question of how to quantify a given outcome. Is it worse for a person to die or for a corporation to be forced to lay off ten thousand employees? At what point do benefits to public health outweigh ecological risks? The answers to these questions only become less clear when realistically complex situations are taken into account, but this is not a problem peculiar to utilitarianism. Normal acts of deliberation often require us to evaluate and compare the possible outcomes of our actions, no matter how different the objects of comparison. Though one may balk at the notion of assigning a monetary equivalent to the value of a human life, insurance companies face this task every day. Whether or not the value of money is comparable to the value of a human life in any objective sense, certain acts of deliberation require that their values be compared insofar as their values enter into deciding between mutually exclusive courses of action.

This problem of comparing apples to oranges is not the only obstacle faced by the utilitarian approach. As is, this framework is insensitive to the distribution of risks and benefits resulting from a course of action. One outcome is considered better than another only if it possesses the greater *aggregate* utility. So for instance, if a test batch of artificial cells were released into rural Asia, infecting and killing ten thousand people, but was later put to use by doctors to save eleven thousand Americans from fatal heart disease over the course of the next ten years, this simple utilitarian calculus would advocate the release of the test batch. Similarly, suppose a commercial artificial cell manufacturer makes a fortune off one of its products, but in doing so causes widespread damage to the environment. Depending on the details of the situation and the values assigned to factors like corporate wealth and environmental health, the utilitarian calculus may deem this course of action morally advisable, even though only a select few would benefit while local populations suffer, which seems utterly unjust. In deciding on a course of action, the distribution of harms and benefits must be taken into account as well as their quantity and magnitude.

Beyond the question of distribution, we might want other factors to influence how harms and benefits are weighted. For instance, many consider the ongoing process of scientific discovery to be crucial for society's continued benefit and enrichment, so that when in doubt about the consequences of a given research program, we should err on the side of free inquiry. Another plausible way one might bias one's assessment of risks and benefits would be to err on the side of the avoidance of harm. We normally feel obligated not to inflict harm on others, even if we think the overall human welfare would increase as a result. A surgeon should never deliberately kill one of her patients, even if doing so would mean supplying life-saving organ transplants to five other dying patients. So if, all else being equal, we should avoid doing harm, it is sensible to bias one's assessment of harms and benefits in favor of playing it safe.

Though it may be true that we have a special obligation against doing harm, a harm-weighted principle of risk assessment is unhelpful in deciding how to proceed with a program as long-term and uncertain as artificial cell science. As Stephen Stich (1978) observes, "The distinction between doing good and doing harm presupposes a notion of the normal or expected course of events" (201). In the surgeon's case, killing a patient to harvest his organs is clearly an instance of a harm inflicted, whereas allowing the other five patients to die due to a scarcity of spare organs is not. But when deciding the fate of artificial cell science, neither pursuing nor banning this research program could be described as inflicting a harm, because there is no way of knowing (or even of making an educated guess about) the kind of scenario to compare the outcome against. In the end artificial cells may prove to be too difficult for humanity to control, and their creation may result in disaster. But it could also be that by banning artificial cell technology we rob ourselves of the capability to withstand some other kind of catastrophe (such as crop failure due to global warming or an antibiotic-resistant plague). The ethical dilemma posed to us by artificial cells is not one that involves deciding between alternative *outcomes*; it requires that we choose between alternative *standards of conduct* where the outcome of any particular course of action is at best a matter of conjecture. So in the present case, the imperative against doing harm is inapplicable.

The ethical problem posed by artificial cells is fundamentally speculative, and cannot be solved by simply weighing good against bad and picking the choice that comes out on top. We can at best weigh *hypothetical* risks against *hypothetical* benefits and decide on the most prudent means of navigating this uncertainty.

This still leaves room for something like a utilitarian calculus, however. Decision theory formulates principles for choosing among alternative courses of action with uncertain consequences, by appropriately weighing risks and benefits. Given a particular decision to make, one constructs a "decision tree" with a branch for each candidate decision, and then sub-branches for the possible outcomes of each decision (the set of outcomes must be mutually exclusive and exhaustive). A utility value is then assigned to each possible outcome (sub-branch). If the probabilities of each possible outcome is known or can be guessed, they are assigned to each sub-branch and the situation is called a decision "under risk." Decisions under risk are typically analyzed by calculating the expected value of each candidate decision (averaging the products of the probabilities of each possible outcome and their utilities), and then recommending the choice with the highest expected value. If some or all of the probabilities are unknown, then the situation is called a decision "under uncertainty." Various strategies for analyzing decisions under uncertainty have been proposed. For example, the risk averse strategy called "minimax" recommends choosing whatever leads to the best of the alternative worse case scenarios. The proper analysis of decisions under uncertainty is not without controversy, but plausible strategies can often be found for specific kinds of decision contexts (Resnick 1987). In every case, for decisions under both risk and uncertainty, decision theory relies on comparing the utilities of possible outcomes of the candidate choices.

Decisions under risk and under uncertainty should be contrasted with a third kind of decision—what we will term decisions in the dark—that are typically ignored by decision theory. Decisions in the dark arise when those facing a decision are substantially ignorant about the consequences of their candidate choices. This ignorance has two forms. One concerns the set of possible outcomes of the candidate choices; this prevents us from identifying the sub-branches of the decision tree. The other is ignorance about the utility of the possible outcomes; this prevents us from comparing the utilities of different sub-branches. In either case, decision theory gets no traction and has little if any advice to offer on the decision.

New and revolutionary technologies like genetic engineering and nanotechnology typically present us with decisions in the dark. The unprecedented nature of these innovations makes their future implications extremely difficult to forecast. The social and economic promise is so huge that many public and private entities have bet vast stakes on the bio-nano future, but at the same time their imagined risks are generating growing alarm (recall Bill Joy's worries, above). Even though we are substantially ignorant about their likely consequences, we face choices today about whether and how to support, develop, and regulate them. We have to make these decisions in the dark.

The same holds for decisions about artificial cells. We can and should speculate about the possible benefits and risks of artificial cell technology, but the fact remains that we now have substantial ignorance about their consequences. Statistical analyses of probabilities is consequently of little use. So, decisions about artificial cells are typically decisions in the dark. Thus, utilitarianism and decision theory and other algorithmic decision support methods have little if any practical value. Any decision-theoretic calculus we attempt will be limited by our current guesses about the shape of the space of consequences, and in all likelihood our picture of this shape will substantially change as we learn more. This does not mean that we cannot make wise decisions; rather, it means that deciding will require the exercise of good judgment. We cannot foist the responsibility for making wise choices onto some decision algorithm like utilitarianism or decision theory.

Deciding in the dark

Although the consequences of creating artificial cells will remain uncertain for some time, before than we will have to face decisions about whether to allow them to be created, and under what circumstances. And as the science and technology behind artificial cells progresses, the range of these decisions will only grow. We will confront decisions about whether to permit various lines of research in the laboratory, whether to allow various kinds of field trials, whether to permit development of various commercial applications, whether to assign liability for harms of these commercial products, whether to restrict access to research results that could be used for nefarious purposes, etc. Our uncertainty about the possible outcomes of these decisions does not relieve us from the responsibility of taking some course of action, and the stakes involved could be very high. So, how should we meet this responsibility to make decisions about artificial cells in the dark?

When contemplating a course of action that could lead to a catastrophe, many people conclude that it's not worth the risk and instinctively pull back. This form of reasoning illustrates what we call the *Doomsday Principle*, which is the idea that *we should not pursue a course of action if it might lead to a catastrophe*.[6] Something like this principle is employed by many people in the nanotechnology community. For example, Merkle (1992) thinks that the potential risks posed by nanomachines that replicate themselves and evolve in a natural setting is so great that they should not only not be constructed; they should not even be designed. He concludes that to achieve compliance with this goal will involve enculturating people to the idea that "[t]here are certain things that you just *do not do*" (Merkle 1992, 292, emphasis in original). This illustrates doomsday reasoning that absolutely forbids crossing a certain line because it might lead to a disaster.

A little reflection shows that the Doomsday Principle is implausible, because it would generate all sorts of implausible prohibitions. Almost any new technology could under some circumstances lead to a catastrophe, but we presumably don't want to ban development of technology in general. More dramatically, there is some risk that getting out of bed on any given morning *could* lead to a catastrophe. Maybe you will fall and disable yourself and thereby be prevented from discovering a cure for cancer; maybe the news that you stayed in bed will get out and spread and ultimately provoke similar behavior globally and cause the world's economy to collapse; maybe a switch triggering the world's nuclear arsenal has been surreptitiously left beside your bed; etc. The point is not that these consequences are at all likely but that they are possible. The same kind of consideration shows that virtually every action could lead to a catastrophe and so would be prohibited by the Doomsday Principle. And this shows that the principle is inconsistent and so incoherent. In any given situation, you have to perform one action or another; even doing "nothing" is really doing something (keeping things the same). And since any action could lead to a catastrophe, you will be violating the Doomsday Principle no matter what you do. For example, the consequences of *not* staying in bed that morning could be

[6] This principle and some of its problems to our knowledge were first discussed by Stich (1978).

just as disastrous as the consequences of staying in bed. Or, more to the point, there could be some catastrophic consequence that society could avert only by developing artificial cells, if for example artificial cells could be used to cure heart disease, which is far and away the leading cause of death in the United States (Ropeik and Gray, 2002). So the Doomsday Principle prohibits your action no matter what you do; and, as we said, you always have to do something. So the principle is incoherent.

The likelihood of triggering a nuclear reaction by getting out of bed is negligible, of course, while the likelihood of self-replicating nanomachines wreaking havoc might not be. With this in mind, one might try to resuscitate the Doomsday Principle by modifying it so that it is triggered only when the likelihood of catastrophe is non-negligible. But there are two problems with implementing such a principle. First, the threshold of negligible likelihood is vague and could be applied only after being converted into some precise threshold (e.g., probability 0.001). But any such precise threshold would be arbitrary and so hard to justify. Second, it will often be impossible to ascertain the probability of an action causing a catastrophe with anything like the requisite precision. For example, we have no way at present of even estimating if the likelihood of self-replicating nanomachines causing a catastrophe is above or below 0.001. Estimates of risks are typically based on three kinds of evidence: toxicological studies of harms to laboratory animals, epidemiological studies of correlations in existing populations and environments, and statistical analyses of morbidity and mortality data (Ropeik and Gray 2002). We lack even a shred of any of these kinds of evidence concerning self-replicating nanomachines, because none of them yet exist.

When someone proposes to engage in a new kind of activity or to develop a new technology today, typically this is permitted unless and until it has been shown that some serious harm would result.[7] Think of the use of cell phones, the genetic modification of foods, the feeding of offal to cattle. In other words, the new activity is innocent until proven guilty. The party who proposes the new activity need not first prove that it is safe; rather, the party who questions its safety bears the burden of proof for showing that it is unsafe. Furthermore, this burden of proof can be met only with scientifically credible evidence that establishes a causal connection between the new activity and the supposed harm. It's insufficient if someone suspects or worries there might be such a connection, or even if there is scientific evidence that there *might* be such a connection. The causal connection must be credibly established before the new activity can be curtailed.

This approach to societal decision making has in the eyes of many lead to serious problems. New activities have sometimes caused great damage to human

---

[7] One notable exception to this pattern is the development of new drugs, which must be proven to be safe and effective before being allowed into the public market.

health or the environment before sufficient evidence of the cause of these damages had accumulated; One notorious case is thalidomide, which was introduced in the 1950s as a sleeping pill and to combat morning sickness, and was later discovered to cause severe birth defects [REF]. This perception has fueled a growing attraction to the idea of shifting the burden of proof and exercising more caution before allowing new and untested activities, i.e., in treating them as guilty until proven innocent. This approach to decision making is now widely known as the *Precautionary Principle*, and it proposes that *we should not pursue a course of action that might cause significant harm even if we are uncertain whether the risk is genuine*. The precise details of the Precautionary Principle vary in different formulations of it, but a representative statement is the following (Geiser 1999, xxiii):[8]

> The Precautionary Principle asserts that parties should take measures to protect public health and the environment, even in the absence of clear, scientific evidence of harm. It provides for two conditions. First, in the face of scientific uncertainties, parties should refrain from actions that might harm the environment, and, second, that the burden or proof for assuring the safety of an action falls on those who propose it.

The Precautionary Principle is playing an increasing role in decision making around the world. For example, the contract creating the European Union appeals to the Principle, and it governs international legal arrangements such as the United Nations Biosafety Protocol [NEED REFS]. With the increasing visibility and influence of the Precautionary Principle is coming growing controversy.[9]

The authors are among those skeptical of the Precautionary Principle. It is only common sense to exercise due caution when developing new technologies, but we find that the Precautionary Principle too insensitive to the complexities that must be confronted when making decisions in the dark. One can think of the Precautionary Principle as a principle of inaction. It embodies the intuition that, when in doubt about, one should leave well enough alone. The advice to leave well enough alone is sensible if two presumptions are true: that things are well at present, and that they will remain so if the status quo is preserved. But these

[8] The formulation of the Precautionary Principle adopted at the Wingspread Conference (Raffensperger and Tickner 1999, pp. 353f) and the ETC Group's formulation (2003, p. 72), are very similar to the formulation quoted in the text. Our discussion of the Principle will be directed to formulations of this type. Since the Principle has been defined in various more or less similar ways, our discussion does not apply equally to all of them.

[9] Attempts to defend the Precautionary Principle and make it applicable in practice are collected in Raffensperger and Tickner (1999), while Morris (2000) gathers skeptical voices.

presumptions are often false, and this points to two problems for the Precautionary Principle. We will consider them in reverse order.

Leaving well enough alone might make sense if the world were unchanging, but this is manifestly false. The world's population is continuing to grow, especially in poor and relatively underdeveloped countries, and this is creating problems that will not be solved simply by being ignored. In the developed world average longevity has been steadily increasing; over the last hundred years in the U.S., for example, life expectancy has increased more than 50% (Wilson and Crouch 2001). This has changed in the major threats to human health, so that today heart disease and cancer are far and away the two leading causes of death in the United States (Ropeik and Gray 2002). Pollution of drinking water is another growing problem. There are estimated to be a hundred thousand leaking underground fuel storage tanks in the United States, and a fifth of these are known to have contaminated groundwater; and a third of the wells in California's San Joaquin Valley have been shown to contain ten times the allowable level of the pesticide DBCP (Ropeik and Gray 2002). These few examples illustrate that there is a continual evolution in the key issues that society must confront. Now, the Precautionary Principle does not require us to stand immobile in the fact of such problems. It allows us to adopt any of a number of tried and true methods when addressing them. What it prevents is using a method that has not been shown to be safe. But there is no guarantee that these and all future problems will succumb to tried and true methods. So the Precautionary Principle ties our hands when we try to address new challenges.

This leads to a second, deeper problem with the Precautionary Principle. New procedures and technologies often offer significant benefits to society, and many of these are new and unique. Cell phones free long-distance communication from the tether of land lines, and genetic engineering opens the door to biological opportunities that would never occur without human intervention. We are not saying that the benefits of these technologies outweigh the risks they pose; we are saying simply that they do have benefits. And the Precautionary Principle simply ignores such benefits when prohibiting unproven procedures and technologies. To forego these benefits is to bring about a harm—what one might call a "harm of inaction." These harms of inaction are opportunity costs created by the lost opportunities to bring about certain new kinds of benefits. Again, our point is not that these opportunity costs are so high that we should always develop new procedures and technologies, but that the Precautionary Principle prevents them from being considered at all.

These considerations surfaced in the dawn of genetic engineering. The biologists who were developing recombinant DNA methods suspected that their new technology might pose various new kinds of risks to society, so the National Academy of Science, U.S.A., empanelled a group of experts to look into the matter. Within a year this group published their findings in *Science*, in what has come to be known as the "Moratorium" letter. They recommended suspending all recombinant DNA studies "until the potential hazards … have been better evaluated" [Berg et al. 1974, p. XXX]. This is an example of precautionary reasoning (well before the Precautionary Principle had been formulated as such). Recombinant DNA studies were suspended even though no specific risks had been scientifically documented. Rather, it was thought that there *might* be such risks, and that was enough to halt activity. In fact, the Moratorium letter proposed an extreme form of precaution, since it temporarily blocked even *research* on recombinant DNA.

The Moratorium letter provoked the National Academy of Science to organize a conference at Asilomar the following year, with the aim of determining under what conditions various kinds of recombinant DNA research could be safely conducted. James Watson, who signed the Moratorium letter and participated in the Asilomar conference, began to have serious misgivings about the degree of precaution that was being advocated. In a recent account of those times, he writes that "I now felt that it was more irresponsible to defer research on the basis of unknown and unquantifiable dangers. There were desperately sick people out there, people with cancer or cystic fibrosis—what gave us the right to deny them perhaps their only hope?" [Watson 2003]. In other words, Watson was concerned about the harms of inaction brought about by extreme precautionary measures.

Some harms of inaction are real and broad in scope, as the threat of rampant antibiotic-resistant disease can illustrate. The overuse and misapplication of antibiotics during the last century has undermined their effectiveness in the future. By 2000 as many as 70% of pneumonia samples were found to be resistant to at least one first-line antibiotic, and multi-drug-resistant strains of *Salmonella typhi* (the bacterium that causes cholera and typhoid) have become endemic throughout South America and Africa, as well as many parts of South and East Asia (WHO 2000). Hospital-acquired *Staphylococcus aureus* infections have already become widely resistant to antibiotics, even in the wealthiest countries. Many strains remain susceptible only to the last-resort antibiotic vancomycin, and even this drug is now diminishing in effectiveness (Enright et al. 2002; WHO 2000). The longer we go on employing the same old tactics in the problem of antibiotic resistance , the less effective they become. It takes on average between 12 and 24 years for a new antibiotic to be developed and approved for human use, and pathogens begin developing resistances to these drugs in just a fraction of this time (WHO 2000). Our present strategy for fighting bacterial infections is not viable in the long run.

One key to weaning ourselves from antibiotics is developing effective preventative medicine. Probiotics, the practice of cultivating our own natural microbial flora, is one such program that has begun gaining popularity among nutritionists, immunologists and pathologists. The cultivation of health-promoting bacterial symbiotes has been shown to enhance the immune system, provide essential nutrition to the host, and decrease the likelihood of colonization by hostile microbes (Bocci 1992, Erickson and Hubbard 2000, and Mai and Morris 2004). Even HIV rates have been shown to decrease among people hosting healthy microbial populations (Miller 2000). Microbes live on virtually every external surface of our bodies, including the surface of our gastrointestinal tract and on all of our mucous membranes, so they are natural first lines of defense against disease.

Our natural microbial flora fail to defend us against many diseases, but wherever a colony of pathogenic organism could thrive, so conceivably could the right innocuous species So we would have a defense against the pathogen if the innocuous microbe prevented the pathogen from thriving, by competitive exclusion. Genetic engineering and artificial cell technologies offer two ways we could develop novel probiotics to compete against specific pathogens.These solutions, moreover, would be viable in the long term and more beneficial than present techniques, as a method of curing as well as preventing disease (leaving aside the many other potential benefits of probiotic nutrition and medicine).

However, the Precautionary Principle would bar us from taking advantage of these new long-term weapons against disease. It is impossible to be certain that a new probiotic would cause no problems in the future, especially when human testing is out of the question (again, because of the Precautionary Principle). Though releasing new probiotics into human microbial ecosystems is undoubtedly risky and must be done cautiously, it may prove the only way to prevent deadly global epidemics in the future. So it is not at all implausible that the consequences of inaction far outweigh the potential risks of these technologies. Thus, the Precautionary Principle may leave us no safer than a deer in the headlights. The initial decisions concerning artificial cells that society will face will have to be made in the dark. The potential benefits of artificial cells seem enormous, but so do their potential, and without gathering a lot more basic knowledge we will be unable to say anything much more precise about those benefits and risks. We will be unable to determine with any confidence even the main alternative kinds of consequences that might ensue, much less the probabilities of their occurrence. Given this lack of concrete knowledge about risks, an optimist would counsel that we take action to pursue the benefits and have confidence that science will be able to find a way to handle any negative consequences when they arise. The Precautionary Principle is a reaction against precisely this kind of blind optimism. Where the optimist sees ignorance about risks as propping the door to action wide open, the Precautionary thinker sees that same ignorance about risks as closing the door to action.

We think that each of these positions is an over reaction to the other position. We propose adopting a more moderate approach to decisions in the dark, one that we hope avoids the excesses of both extremes. We dub this approach the *"Cautious Courage" Principle*, and it involves *cautiously but proactively pursuing new courses of action and new technologies even if the risks are uncertain*. As its name indicates, the Cautious Courage Principle involves the exercise of caution, as does the Precautionary Principle. It recommends that we should pursue new courses of action only if we have vigilantly sought to identify and understand any risks that might arise. But uncertainty about outcomes and possible risks should not necessarily block action. This is where the courage in our principle enters. We should also have an eye on the possible benefits that might ensue, and compare the risks and benefits of various alternative courses of action (including the "action" of doing nothing), and have the courage to make a leap in the dark if on balance that seems to make the most sense.

The Cautious Courage Principle is vague, of course; it is no mechanical algorithm for generating decisions. The Precautionary Principle is sometimes criticized on these very grounds, and if this is a criticism it would apply equally to the Cautious Courage Principle. But we are unsympathetic with this criticism. For the reasons outlined in the previous section, we think no mechanical algorithm for making decisions in the dark will be sensible. So there will be no avoiding the exercise of judgment, and that means that guiding principles will necessarily be somewhat vague.

New technologies give us new powers, and these new powers confront us with new choices—principally, the choices whether and how to exercise the new powers. Thus, we bear new responsibilities, in particular the responsibility of making new kinds of choices wisely. The Precautionary Principle in effect says that responsible choices will be governed only by caution. This strikes us as timidly backing away from the new opportunities we have. The Cautious Courage Principle replies that we should be prepared to take some risks if the possible benefits are significant enough and the alternatives are comparatively unattractive. In other words, it counsels that we should courageously step up to the plate and be ready to swing.

Conclusions

Artificial cells are in our future, and that future could come within this decade. By harnessing the automatic regeneration and spontaneous adaptation of life, artificial cells promise society a wide variety of social and economic benefits. But their ability to self-replicate and unpredictably evolve create unforeseeable risks to human health and the environment. So it behooves us to start thinking through the implications of artificial cells today.

From the public discussion on genetic engineering and nanotechnology one can predict the outline of much of the debate that will ensue. One can expect the objections that creating artificial cells is unnatural, that it commoditizes life, that it fosters a reductionistic perspective, and that it is playing God, but we have explained why these kinds of considerations are all unpersuasive. Utilitarianism and decision theory offer scientifically objective and pragmatic methods for deciding what course to chart, but they are inapplicable when the decision must be made in the dark. No quantitative algorithm will guarantee sound policies as long as society is largely ignorant of the potential consequences of its actions. The Precautionary Principle is being increasingly used to cope with important decision in the dark, but the Principle fails to give due weight to potential benefits lost through inaction (what we called "harms of inaction"). We suggest that the Cautious Courage Principle might preserve the attractions of the Precautionary Principle while avoiding its weaknesses.

This suggestion is not offered as the definitive solution for how to think through the social and ethical implications of artificial cells. Our goal is primarily to initiate an informed and thoughtful public discussion of these issues. Public debate is already underway on related issues concerning genetic engineering and nanotechnology. We should now add a related thread on artificial cells. Indeed, our reasoning here applies to in general to technologies with the power to radically transform our lives, both those like genetic engineering and nanotechnology with which we are familiar, those like artificial cells of which we are now learning, and those of which we have not yet started to dream.

Acknowledgements

This research was supported by a Ruby Grant from Reed College. For helpful discussion, thanks to Steve Arkonovich, Todd Grantham, Scott Jenkins, Leighton Reed, and to audiences at the LANL-SFI Workshop on Bridging Non-Living and Living Matter (September 2003), at the ECAL'03 Workshop on Artificial Cells (September 2003), at a Philosophy Colloquium at Reed College (October 2003), and at a Philosophy seminar at the College of Charleston (November 2003). The promise of living technology in general and artificial cells in particular was developed in collaboration with John McCaskill, Norman Packard, and Steen Rasmussen, starting in the summer of 2001.

References

Bartel, D. P., and Unrau, P. J. Constructing an RNA world. *Trends in Cell Biol*ogy 9: M9-M13.

Berg, et al. 1974. Potential biohazards of recombinant DNA molecules. *Science* 185: 303.

Bocci, V. (1992). The neglected organ: bacterial flora has a crucial immunostimulatory role. *Perspectives in Biology and Medicine* 35(2), pp. 251-60.

Brannigan, M. C. 2001. Introduction. In *Ethical issues in human cloning: cross-disciplinary perspectives*, ed. M. C. Brannigan, pp. 1-4. New York: Seven Bridges Press.

Brooks, R. (2001) The relationship between matter and life. *Nature* 409, 409-411.

Cho, M. K., et al. 1999. (1999) Ethical considerations in synthesizing a minimal genome. *Science* 286, 2087.

Chrichton, M. 2000. *Prey*. New York: HarperCollins.

Comstock, Gary L. 2000. *Vexing nature? On the ethical case against agricultural biotechnology.* Boston: Kluwer Academic Publishers.

Crandall, B. C. 1992. The risks of nanotechnology. In *Nanotechnology research and perspectives*, ed. by B. C. Crandall and James Lewis, pp. 287-292. Cambridge: MIT Press.

Dobson, Andrew 1995. Biocentrism and genetic engineering. *Environmental Values* 4: 227.

Drexler, K. E. 1986. *Engines of creation: the coming era of nanotechnology*. New York: Doubleday.

Dyson, Freeman J. 2003. The future needs us! *New York Review of Books 50* (13 February), 11-13.

Ekland, E. H., and Baretl, D. P. 1996. RNA catalyzed nucleotide synthesis. *Nature* 382: 373-376.

Enright, Mark C., Robinson, D. Ashley, and Randle, Gaynor (2002). The evolutionary history of methicillin-resistant *Staphylococcus aureus* (MSRA). *The Proceedings of the National Acadamy of Sciences of the United States of America* 99 (11), pp. 7687-92.
Erickson, Kent L. and Hubbard, Neil E. (2000). Probiotic immunomodulation in health and disease. *The Journal of Nutrition* 130(2), pp. 403S-409S.
ETC Group. 2003. *The big down: from genomes to atoms*. Available on the web at http://www.etcgroup.org [accessed June 2003].
Fraser, C.M., et al. (1995) The minimal gene component of *Mycoplasma genetalium*. *Science* 270, 397-403.
Geiser, Ken. 1999. Establishing a general duty of precaution in environmental protection policies in the United States: a proposal. In Raffensberger and Tickner (1999), pp. xxi-xxvi.
Gillis, J. (2002) Scientists planning to make new form of life. *Washington Post*, November 21, A01.
Hutchison, C.A., et al. (1999) Global transposon mutagenesis and a minimal Mycoplasma genome. *Science* 286, 2165-2169.
Joy, B. 2000. Why the future does not need us. *Wired* 8 (April). Available on the web at http://www.wired.com/wired/archive/8.04/joy.html [accessed June 2003].
Kass, Leon R. 2002. *Life, liberty, and the defense of dignity: the challenge for bioethics*. San Francisco: Encounter Books.
Mai, Volker, and Morris, J. Glenn, Jr. (2004). Colonic bacterial flora: changing understandings in a molecular age. *The Journal of Nutrition* 134(2), pp. 459-64.
Merkle, R. 1992. The risks of nanotechnology. In B. Crandall & J. Lewis (Eds.), *Nanotechnology research and perspectives* (pp. 287-294). Cambridge, MA: MIT Press.
Miller, Karl E. (2000). Can vaginal lactobacilli reduce the risk of STDs? *American Family Physician* 61(10), pp. 3139-40.
Mooney, Chris. 2002. Nothing wrong with a little Frankenstein. *Washington Post*, December 1, B01.
Morris, Julian, ed. 2000. *Rethinking risk and the Precautionary Principle*. Oxford: Butterworth-Heinemann.
Nielsen, P., et al. (1991) Sequence-selective recognition of DNA by strand displacement with a thymine-substituted polyamide. *Science* 254, 1497.
Oberholzer, T. et al. 1995. Polymerase chain reaction in liposomes. *Chem. Biol*. 2: 677-682.
Pohorille, A., Deamer, D. (2002) Artificial cells: prospects for biotechnology. *Trends in Biotechnology* 20, 123-128.
Raffensperger, Carolyn, and Tickner, Joel, eds. 1999. *Protecting public health and the environment: implementing the Precautionary Principle*. Washington, D.C.: Island Press.
Rasmussen, S. et al. (2003) Bridging nonliving and living matter. *Artificial Life* 9, 269-316.
Reiss, Michael J. and Straughan, Roger. 1996. *Improving nature? The science and ethics of genetic engineering*. Cambridge: Cambridge University Press.
Resnick, Michael. (1987) *Choices: an introduction to decision theory*. Minneapolis: University of Minnesota Press.
Ropeik, David, and Gray, George. 2002. *Risk: a practical guide to deciding what's really safe and what's really dangerous in the world around you*. Boston: Houghton Mifflin.
Silver, L. M. 1998. Cloning, ethics, and religion. *Cambridge Quarterly of Healthcare Ethics* 7: 168-172. Reprinted in *Ethical issues in human cloning: cross-disciplinary perspectives*, ed. M. C. Brannigan, pp. 100-105. New York: Seven Bridges Press, 2001.

Stich, Stephen P. 1978. The recombinant DNA debate. *Philosophy and Public Affairs* 7: 187-205.
Stikeman, Alexandra. 2002. Nanobiotech makes the diagnosis. *Technology Review* (May): 60-66.
Szostak, J.W. et al. (2001) Synthesizing life. *Nature* 409, 387–390.
Walde, P., et al. (1994) Autopoietic self-reproduction of fatty acid vesicles. *J. Am. Chem. Soc.* 116, 11649-116454.
Watson, James D. 2003. *DNA: the secret of life*. New York: Alfred A. Knopf.
Wilson, Richard, and Crouch, Edmund A. C. 2001. *Risk-benefit analysis*. Cambridge: Harvard University Press.
World Health Organization (2000). World Health Organization Report on Infectious Diseases 2000: Overcoming Antimicrobial Resistance. Available on the Web at http://www.who.int/infectious-disease-report/2000/ [accessed May 2004].