



What Price Freedom?

Robert A. Freitas Jr^{*}

Institute for Molecular Manufacturing, Palo Alto, California, USA

Criminals vs. terrorists

In an attempt to make David Brin's¹ privacy-free 'transparent society' more palatable to civil libertarians, Robert Sawyer² has proposed an "Alibi Archive" in which everyone's activities are meticulously recorded in a centralized, judicially controlled archive, with the archives legally accessible only under court order and only upon request of the person whose activities were recorded. In a criminal investigation, this person would be able to access (and make public from the archives) those records of his activities that would definitively establish an alibi for him, thus conclusively proving that he was elsewhere when the crime was being committed. Potential criminals would know that they would not be able to establish an alibi in this manner, and thus would be deterred from committing crimes.

Regardless of the merits of this idea (and there are many aspects that can be debated), it seems that it is workable only with respect to perpetrators who actually care if their illegal activities are discovered. In the unique case of suicidal terrorists who plan to kill themselves during the achievement of their objectives, the alibi archive simply won't work as a deterrent. Suicidal perpetrators plan on being dead after the commission of their crimes. They won't care what, or whether, anything can be proven after they're gone. We need some additional ways of deterring them and disabling their ability to act. Perhaps some sort of highly intrusive and actively monitored nanotechnology-enabled omnipresent recording system could be employed to this end.

Freedom fighters vs. terrorists

But we must be careful not to throw the baby out with the bathwater. In view of the recurring emergence of oppressive governments throughout human history, we must ask: Do we really want to make 'freedom fighting' as impossible as 'terrorism'? Consider some future age in which the United States (or pick your favorite alternative technologically sophisticated

^{*} E-mail: see www.rfreitas.com.

¹ Brin, David, "*The Transparent Society: Will Technology Force Us to Choose Between Privacy and Freedom?*" (1998) (Perseus Books Group).

² Robert J. Sawyer, "Privacy: Who Needs It?" *Maclean's*, 7 October 2002; <http://www.sfwriter.com/privacy.htm>.

developed nation) becomes dominated by a totalitarian dictator (whether nanotechnology-enabled or otherwise). Imagine secret police crashing through the doors of private homes in the dead of night; the arrest and torture of citizens as a purposeful government policy for the suppression of dissent; gulags to warehouse troublemakers; and even summary executions. As responsible citizens and humanitarians, we may wish to retain the right, and the ability, to overthrow such an oppressive government by force if necessary, even at the risk of our own lives. An argument for the morality of this idea is made in the preamble to the U.S. Declaration of Independence. Furthermore, the U.S. Bill of Rights (Second Amendment) enshrines the right of U.S. citizens to bear arms in part to ensure that no future usurper (whether foreign or domestic) would dare think that he could establish a tyranny on these shores, knowing that the citizens here were sufficiently well armed to contest his unwanted rule.

The word 'freedom' hides many rhetorical landmines. In this essay, 'freedom' will refer to the minimal possible level of control imposed by external governmental authority on the thoughts and actions of individual people, consistent with the stable and thriving existence of a civilized society. But there is another fundamental definitional problem that we must now address: Are there any actionable distinctions between 'freedom fighters' and 'terrorists'? The distinctions are not clear-cut, but one approach might be to assess the differences between these two types of actors in terms of their goals and methods.

In terms of goals, a freedom fighter typically is focused on destroying what he regards as an oppressive government, including its leaders, its functional appendages, and its supporters. The terrorist often has more diffuse objectives. He may be seeking to overthrow a government, but he might also be trying to displace an extant social or economic order with which he disagrees but is otherwise powerless to influence directly. Or, he may be trying to alter the culture, including the religious preferences or practices of local or larger regional populations, or to affect outcomes in territorial or other disputes, perhaps based on religious or ethnic differences.

In terms of methods, freedom fighters are not above employing dirty tactics, including assassinating individual key civilians who are viewed as indirectly supporting, or at least acquiescing to, the oppressive government. But these tactics will mostly be directed at the oppressive government or its specific supporters or physical plant, and not at the general population. In contrast, the terrorist often prefers to target civilians and otherwise innocent parties, in large indiscriminate numbers, regardless of whether they are a part of (or support in any way) the oppressive regime or unwanted cultural milieu that the terrorist seeks to displace. The terrorist employs a kind of trickle-up theory of political action—rather than attacking an otherwise impregnable entity directly, he attacks an innocent civilian population in the hopes that this population will become restless enough to demand changes in the government or society in order to get the terrorists to stop. In effect, terrorists engage in social blackmail against innocents. Since, in the nanotechnology-rich future, a terrorist could decide that huge blocks of innocent humanity should be sacrificed for political, racial or religious reasons, we should not, in good conscience, allow this capability to emerge unchallenged.

Considering these differences in goals and methods, we could seek to design selective nanotechnology-based defenses against social blackmail by terrorists that will not at the same time forestall freedom fighting. Such defenses could be keyed to the differences in goals—say, protecting from murder all individuals except those belonging to government entities. Or such defenses could be keyed to differences in methods—say, systems that allow the murder of one

person at a time, but actively prevent the perpetration of simultaneous mass murders. However, a corrupt government would not allow itself to remain vulnerable in this manner and might seek to turn the tables and make the subject population exclusively vulnerable instead. It is difficult to see how to implement such defenses in a reliable and incorruptible manner without employing an executive artificial intelligence (AI) that is capable of informed judgment and independent action (e.g., the robot policeman scenario from the 1951 movie, *The Day the Earth Stood Still*), which raises a host of new difficulties and issues.

Another class of terrorist-selective defenses could be keyed to the intentions, rather than to the actions, of potential actors. Future medical nanotechnology should enable intrusive involuntary brain scans of sufficient fidelity to accurately measure and report internal psychological states and motives. But here too there are several difficulties. First, all human beings on Earth would have to be continuously monitored for ‘terrorist’ intentions. This monitoring duty would probably fall to some government (or related institutional) entity, and a corrupt government entity could not be prevented from scanning for ‘freedom fighter’ intentions as well. Such scanning would elevate Brin’s ‘transparent society’ to a new level to intrusiveness—we might call it the ‘transparent mind’—which would be even more anathematic to civil libertarians and would offer even greater potential for abuse. Second, the amount of data to be processed might be so enormous as to require the intervention of an AI (as in the previous example) to sort it all out, whether the AI was a stand-alone system or embedded in a human/machine hybrid system. Third, it is but a small step from passively monitoring brain states to actively controlling those brain states using nanotechnology-based neural nanorobotics, which would enable the push-button disposal of critics by tyrants. Thus, the freedom fighters would again be disabled along with the terrorists.

If we conclude that it may not be possible either to reliably distinguish between freedom fighters and terrorists, or to reliably defend against one but not the other, then we may have to resign ourselves to the existence of both—or neither—of these types of actors in our world. If we choose to accept both (tolerating freedom fighters in order to avoid tyrants and tyrannical governments), then we are tacitly agreeing to accept the presence of terrorism. If we reject both, the way is open for tyranny. Which shall we choose?

Shall we accept both terrorism and freedom fighting?

If we agree that it is morally proper to allow freedom fighting (in those rare instances when it becomes necessary), then what is the potential harm that we risk by agreeing to accept the possibility of terrorism along with it? No one has yet presented a clear and comprehensive exposition of the actual dangers involved. An important recommendation is that a detailed and ruthlessly honest study of likely scenarios and consequences should be performed as soon as possible, perhaps incorporating a scenario gaming process that encourages the most imaginative informed challenges to be tested by intelligent motivated players. One common claim is that the situation may be asymmetrical in favor of the evildoers, who may choose the place and timing of attacks and may also employ the element of surprise. This claim seems somewhat naive because it ignores the following important factors.

First, the terrorists are unlikely to possess the most advanced technologies available. Weapons used by terrorists are often relatively low-tech because such means are cheaper to

obtain, simpler to operate, and less likely to fail when activated in rushed circumstances. Terrorists also tend to be less educated and less technically sophisticated than defenders. The most advanced technologies usually will be possessed by the defenders—typically government-funded police or military entities in the developed world. The cleverer and more multifunctional that future nanotechnology-based weapons are posited to be, the less likely terrorists are to have them; hence, sophisticated launch and dispersal scenarios will be less likely to be successfully accomplished by terrorists. Of course, there always can be exceptions—for example, terrorists could surreptitiously receive (or steal) advanced technologies from sources in developed nations.

Second, as molecular manufacturing pervades human society, there will occur numerous minor mishaps and relatively inconsequential accidents involving this new technology (as with any new technology that is introduced for the first time). Basic civilian defensive systems analogous to police and fire departments will gradually emerge that are specifically designed to cope with nanotechnology-based minor mishaps and emergencies on a local level.³ Hence, the future environment in which terrorists must operate will include ubiquitous nanotechnology-based protective civil defenses.

Third, the knowledge that a mass-murder terrorist threat scenario is plausible will induce responsible governments to put in place extensive external public event monitoring⁴ (not necessarily requiring the monitoring of internal brain states of individual citizens) and military-type responses to deal with larger-scale threats of mass destruction if and when they might occur.

The net effect of these factors is to moderate the possible negative impacts of a nanotechnology-era terrorist attack. Such attacks might therefore be deemed an acceptable risk if there has been a reasonable level of investment in civil defense by the government. Admittedly, this is only a tentative conclusion that will require a great deal of further study and considerable (possibly heated) debate.

Shall we reject both terrorism and freedom fighting?

On the other hand, if we agree that it is morally proper to make acts of terrorism upon innocent populations impossible to carry out, then what is the potential harm we risk by agreeing to reject any possibility of freedom fighting along with it? It appears that the harm we risk in this case could be far more severe. That's because the conclusions about terrorists that we reached in the previous section are all precisely reversed in the case of tyrants.

Specifically, the tyrant—especially one in control of a technologically sophisticated, highly developed nation—would be more likely to possess some of the most advanced technologies currently available. He probably would have access to the most multifunctional weapons and delivery systems, and these systems will be capable of numerous, diverse, and secret deployments. Since he may control (whether directly or indirectly) the governmental organs of civilian emergency and military response, he also could circumvent the normal

³ Robert A. Freitas Jr., Ralph C. Merkle, *Kinematic Self-Replicating Machines*, Landes Bioscience, Georgetown TX, 2004; <http://www.MolecularAssembler.com/KSRM/6.3.1.htm#p22>.

⁴ Robert A. Freitas Jr., "Some Limits to Global Ecophagy by Biovorous Nanoreplicators, with Public Policy Recommendations," Zyvex preprint, April 2000, Section 9 (recommendations 2 and 3); <http://www.rfreitas.com/Nano/Ecophagy.htm>.

protective programs of these systems, forcing them to react to external threats in an asymmetric manner to his own advantage. Even worse, the tyrant could corrupt these systems and redirect them as global threats, and thus aspire to global domination.

This analysis seems to suggest that preserving the ability to freedom fight against tyrants may be necessary to avoid a future of perpetual despotic thralldom, and the price we pay is the acceptance of the possibility of terrorism.

Better dead or red?

Just because one society initially chooses not to employ the necessary heavy-handed nanotechnological means to render both freedom fighting and terrorism effectively impossible, that does not mean other societies will make the same choice, nor even that the first society will not change its position over time. It appears quite likely, though perhaps not inevitable, that eventually, somewhere in the world, a tyrant will emerge who is equipped with some of the most sophisticated nanotechnological instrumentalities available. This tyrant would likely employ these advanced technical means to eliminate within his own borders any possibility of freedom fighting or terrorism, both of which he might rationally presume could be directed at him or his vassals. Other technically sophisticated societies might or might not have the will or the means to oppose this tyrant, and still other societies might decide to emulate or join him; therefore, his emergence and ascendancy cannot be ruled out.

Our analysis thus far has suggested that the existence of terrorism may be an acceptable price to pay, in order to keep alive the option of freedom fighting to contest the dismal consequences of a nanotechnology-enabled despotism. But even if we agree in principle that it is morally proper to allow freedom fighting in many situations, there is still one instance in which this conclusion becomes more difficult to defend, and which raises perhaps the most troubling question of this essay: *Is it right to engage in freedom fighting to dislodge a tyrant if such fighting might result in the destruction of all humankind?*

Note that to address this issue it is not necessary to completely resolve the interesting but somewhat ancillary technical military question of whether nanotechnology-enabled defensive or offensive instrumentalities would be inherently or presumptively paramount in effectiveness, or alternatively whether defense and offense will likely remain roughly equal in effectiveness during the unfolding of likely nanotechnology development pathways. Such rough equality has been observed throughout much of human history and also regularly occurs in stable biological, commercial, mimetic, and other freely-evolving competitive ecologies, though noteworthy (if temporary) imbalances have occurred sporadically in various times and places throughout history. Rather, the key issue here is whether any one party can achieve a sufficient nanotechnological military capability that would enable the destruction of all life on Earth, or of all human civilization.

On this query, we must regrettably conclude that the answer is most probably yes. Our existence proof for this claim derives from the Cold War doctrine of Mutual Assured Destruction that existed between the United States and the Soviet Union from the 1950s through the 1980s. This doctrine was implemented using a quantity of deliverable nuclear munitions on each side that was sufficient, if unleashed, to effectively extinguish all of modern human technological civilization on Earth, if not necessarily exterminate all individual human

beings. (It was recognized that some few might escape the direct effects of the nuclear holocaust by hiding in caves or in underground bunkers, and further assuming that a lengthy nuclear winter did not ensue following the detonation of tens of thousands of high-yield nuclear warheads over carefully selected technology-, fuel-, and people-rich targets.) So we already have an existence proof that Damoclean threats can and do exist in the case of nuclear technology. Much has been written of similar threats that may emerge from the fields of genetics and biotechnology. We must concede that the same situation cannot be ruled out in the case of advanced nanotechnology. Note that this existential threat can arise in many ways, ranging from despotic defenses so well entrenched as to require world-destroying countermeasures to defeat, to doomsday weapons constructed by the despot for the sole purpose of blackmailing humanity into submission.

Thus our final question can be rephrased more pointedly: Is a tyrannized humanity worth preserving, even at the expense of its freedom, in order to maintain the very existence of the human species? From a cosmic perspective, if we are the only sentience in the universe then it could be strongly argued that it would be immoral to take actions which have a high probability of leading to the extinction of our species, even in the name of freedom. That's because there is always the remote hope that in some future epoch the tyranny might fail, thus eventually returning the preserved humankind to a state of freedom, and because such resistance seems ultimately pointless if everyone is free but dead. To employ a familiar vernacularism, is humanity "better dead than red"?⁵

Some factors to consider

The answer to this difficult question will require careful thought and must deal with many fundamental issues concerning the possible uses of future advanced nanotechnologies. Let's set aside for now the basic tactical questions such as to what degree we should challenge a despot who merely *claims* to possess doomsday technologies, but offers no clear objective evidence supporting this claim beyond his boasts. We'll also ignore the many additional technical and ethical complexities introduced by the possibilities of migration into space, uploading to biology-free synthetic physical bodies, or emigrating to computer-generated virtual realities wherein dwell a 'virtual humanity'. Let us consider here only the simplest case, which is biological humanity here on Earth.

We might begin thinking about the answer by examining the implications of the decision to preserve biological humanity by complete submission to a tyrant. Armed with a sophisticated intrusive nanotechnology, a tyrant may undertake to rewrite the minds of his vassals, and perhaps even rewrite their human biology, to make it literally impossible for individuals to resist his will. As nano-lobotomized slaves, we might be allowed to retain most of our intelligence, but we would become integrally sheeplike, temperamentally and constitutively unable to rise up and fight for our freedom against the sovereign. We might be not just temporarily brainwashed but permanently braintailored, with much of our personality remaining intact but our minds utterly convinced beyond any doubt that the tyrant was good and should not be disobeyed. This could be implanted as an almost instinctual response, psychophysiologically

⁵ <http://www.answers.com/topic/better-dead-than-red-1>.

compelling in the same way that it is nearly impossible for us to disobey our primitive urges to breathe, eat or drink to stay alive. To further improve our efficiency, the tyrant might decide to relieve our minds of certain ‘silly’ distractions that would normally lead us to waste time fighting among ourselves over politics, religion, business, or sexual competition, rather than working as hard as possible on projects the tyrant deems useful. Even in this rather neutral scenario, many would say that the participants have become mere walking shells of former humanity who are no longer truly human in any meaningful sense. If the human race as we know it already has become functionally extinct, they would argue, then perhaps the race might no longer be viewed as being at risk of extinction if freedom fighters dare to oppose the tyrant—because the race *is* extinct. Yet even here, there is plenty of room for optimists to argue for continued inaction because “where there’s life, there’s hope”.

Both lighter and darker flavors of the above scenario are readily imagined. If the sovereign is a genuinely benign autocrat, then we might find his rule to be acceptable, even welcome, as an antidote to a deepening sociocultural chaos driven by accelerating change. Perhaps His Benevolence is truly interested in peace, exploration, artistic endeavors, and improving the material and spiritual quality of our lives. In the ideal case, he would be a true humanitarian, managing our affairs to maximize happiness and responsible progress while minimizing discord in all areas. But darker versions abound as well. The sovereign could be a truly malevolent autocrat, abusing his vassals for cruel sport or to advance his own extravagant lusts for luxuries, exclusivities, casual whims and sexual desires without regard for the welfare of his victims and having programmed them to willingly accept the abuse. We might hope that such a malevolent person would be self-destructive, or would eventually grow bored with the evil pleasures of pure ruination. But these may be vain hopes. If the malefactor has used nanotechnological means to ensure human obedience to his demands, then he faces no meaningful external threat to his control. And with the same technology, he can rewrite his *own* brain: (1) to enhance his lusts and intellectual powers over those possessed by others, (2) to render himself incapable of seriously contemplating suicide, and (3) to make the boredom of repetitive destructive activities entirely tolerable, even enjoyable. Since nanotechnology also enables the conquest of natural mortality, and because both the environment and the physical body of the tyrant can be made assassination- and accident-resistant by similar means, then he cannot be counted upon to simply die off, an important natural safety valve that has ended many tyrannies in human history. What then could possibly unseat him, once he is thoroughly ensconced in power?

We are inevitably drawn back to considering whether individual freedom is worth risking collective nonexistence. This is a tough question, to say the least. Many of us who came of age in the 20th century North American culture that glorifies individualism may tend to assume that freedom is almost always worth the risks. But other human societies where the imperatives of the community are elevated above those of the individual may analyze the situation with different assumptions. This cultural Rorschach could have real consequences. Those cultures that value individual freedom most highly might prefer to fight to preserve it, and hence may have a greater probability of becoming extinct. Those cultures that value individual freedom less highly may have a greater probability of being usurped and dominated by a nanotechnology-enabled tyrant.

It remains to be seen whether a coherent strategy can be synthesized from the diametric worldviews of universalists (who are interventionists) and multiculturalists (who are isolationists). Recognizing that global tyranny is a logical end-state of the unchecked spread of nanotechnology-enabled dictatorships that are capable of employing perfect mind control, those who subscribe to the policy doctrine of preemption might rationally conclude that it is necessary to actively liberate other societies that have already decided to capitulate (“entrust their future”?) to a nanotechnology-enabled autocrat. But might not budding tyrants rationally conclude that any developed nation population that treasures individual freedom above most other moral values should be exterminated preemptively in order to eliminate the most obvious threat to their global ambitions? Consider that humanity may have survived the Cold War because at key moments of crisis, both sides opted for survival over domination. In future conflicts, if either side is significantly less dedicated to survival than to domination, then, like a terrorist, that side will not be deterred from seeking domination at all costs.

Could mere discussion of these issues create a self-fulfilling prophecy? It is true that if potential future tyrants come to believe that people in general are unlikely to have the desire or will to resist them, or that people will be so effectively disarmed of personal weaponry by their well-meaning but overprotective governments that individual armed resistance would become futile, then deterrence of nanotechnology-enabled tyrannies is minimized and the emergence of those regimes may be accelerated. But this should affect only the timing, and not the ultimate fact, of such emergence. If the technology allows it—and it does—then eventually some tyrant will seek to close his iron fist around the throat of humankind. We need to decide what, if anything, we ought to do about this.

About the author:

Robert A. Freitas Jr. is Senior Research Fellow at the Institute for Molecular Manufacturing in Palo Alto, California. He received B.S. degrees in Physics and Psychology from Harvey Mudd College in 1974 and a J.D. from University of Santa Clara in 1979. Freitas co-edited the 1980 NASA feasibility analysis of self-replicating space factories and in 1996 authored the first detailed technical design study of a medical nanorobot ever published in a peer-reviewed mainstream biomedical journal. More recently, Freitas is the author of *Nanomedicine*, the first book-length technical discussion of the potential medical applications of molecular nanotechnology and medical nanorobotics; the first two volumes of this 4-volume series were published in 1999 and 2003 by Landes Bioscience. Freitas has published 25 refereed journal publications, patents, or contributed chapters on nanotechnology and in 2004 co-authored the book *Kinematic Self-Replicating Machines*, another first-of-its-kind technical treatise.