



# Singularities and Nightmares: the Range of Our Futures<sup>1</sup>

**David Brin**\*

*Futures Unlimited, California, USA*

In order to give you pleasant dreams tonight, let me offer a few possibilities about the days that lie ahead—changes that may occur within the next twenty or so years, roughly a single human generation. Possibilities that are taken seriously by some of today's best minds. Potential transformations of human life on Earth and, perhaps, even what it means to be human.

For example, what if biologists and organic chemists manage to do to *their* laboratories the same thing that cyberneticists did to computers? *Shrinking* their vast biochemical labs from building-sized behemoths down to units that are utterly compact, making them smaller, cheaper, and more powerful than anyone imagined. Isn't that what happened to those gigantic computers of yesteryear? Until, today, your pocket cell phone contains as much processing power and sophistication as NASA owned during the moon shots. People who foresaw this change were able to ride this technological wave. Some of them made a lot of money.

Biologists have come a long way already toward achieving a similar transformation. Take, for example, the Human Genome Project, which sped up the sequencing of DNA by so many orders of magnitude that much of it is now automated and miniaturized. Speed has skyrocketed, while prices plummet, promising that each of us may soon be able to have our own genetic mappings done, while-you-wait, for the same price as a simple EKG. Imagine extending this trend, by simple extrapolation, compressing a complete biochemical laboratory the size of a house down to something that fits cheaply on your desktop. A *MolecuMac*, if you will. The possibilities are both marvelous and frightening.

When designer drugs and therapies are swiftly modifiable by skilled medical workers, we should all benefit.

But then, won't there also be the biochemical equivalent of 'hackers'? What are we going to do when kids all over the world can analyze and synthesize any organic compound, at will? In that event, we had better hope for accompanying advances in artificial intelligence and robotics... at least to serve our fast food burgers. *I'm* not about to eat at any restaurant that hires resentful human adolescents, who swap fancy recipes for their home molecular synthesizers over the Internet. Would you?

---

\* E-mail: davidbrin@sbcglobal.net; <http://www.davidbrin.com/>.

<sup>1</sup> Parts of this essay were transcribed from a speech before the conference Accelerating Change 2004: "Horizons of Perception in an Era of Change" November 2004 at Stanford University.

Now don't get me wrong. If we ever do have MolecuMacs on our desktops, I'll wager that 99 percent of the products will be neutral or profoundly positive, just like most of the *software* creativity flowing from young innovators today. But if we're already worried about a malicious one percent in the world of bits and bytes—hackers and cyber-saboteurs—then what happens when this kind of 'creativity' moves to the very stuff of life itself? Nor have we mentioned the possibility of intentional abuse by larger entities—terror cabals, scheming dictatorships, or rogue corporations.

These fears start to get even more worrisome when we ponder the next stage, beyond biotech. Deep concerns are already circulating about what will happen when *nanotechnology*—ultrashall machines building products atom-by-atom to precise specifications—finally hits its stride. Molecular manufacturing could result in super-efficient factories that create wealth at staggering rates of efficiency. Nano-maintenance systems may enter your bloodstream to cure disease or fine-tune bodily functions. Visionaries foresee this technology helping to save the planet from earlier human errors, for instance by catalyzing the recycling of obstinate pollutants. Those desktop units eventually may become *universal fabricators that turn almost any raw material into almost any product you might desire...*

... or else (some worry), nanomachines might break loose to become the *ultimate pollution*. A self-replicating disease, gobbling up everything in sight, conceivably turning the world's surface into *gray goo*.<sup>2</sup>

Others have raised this issue before, some of them in very colorful ways. Take the sensationalist novel *Prey* by Michael Crichton, which portrays a secretive agency hubristically pushing an arrogant new technology, heedless of possible drawbacks or consequences. Crichton's typical worried scenario about nanotechnology follows a pattern nearly identical to his earlier thrillers about unleashed dinosaurs, robots, and dozens of other techie perils, all of them viewed with reflexive suspicious loathing. (Of course, in every situation, the perilous excess happens to result from *secrecy*, a topic that we will return to later.) A much earlier and better novel, *Blood Music* by Greg Bear, presented the up and downside possibilities of nanotech with profound vividness. Especially the possibility that most worries even optimists within the nanotechnology community—that the pace of innovation may outstrip our ability to cope.

Now, at one level, this is an ancient fear. If you want to pick a single cliché that is nearly universally held, across all our surface boundaries of ideology and belief—e.g. left-versus-right, or even religious-vs-secular—the most common of all would probably be: "*Isn't it a shame that our wisdom has not kept pace with technology?*"

---

<sup>2</sup> In his article, "*Molecular Manufacturing: Too Dangerous to Allow?*" Robert A. Freitas Jr. describes this scenario. One common argument against pursuing a molecular assembler or nanofactory design effort is that the end results are too dangerous. According to this argument, any research into molecular manufacturing (MM) should be blocked because this technology might be used to build systems that could cause extraordinary damage. The kinds of concerns that nanoweapons systems might create have been discussed elsewhere, in both the nonfictional and fictional literature. Perhaps the earliest-recognized and best-known danger of molecular nanotechnology is the risk that self-replicating nanorobots capable of functioning autonomously in the natural environment could quickly convert that natural environment (e.g., "biomass") into replicas of themselves (e.g., "nanomass") on a global basis, a scenario often referred to as the "gray goo problem" but more accurately termed *global ecophagy*". In explaining this scenario, Freitas does not endorse it.

While this cliché is clearly true at the level of solitary human beings, and even mass-entities like corporations, agencies or political parties, I could argue that things aren't anywhere near as clear at the higher level of human *civilization*. Elsewhere I have suggested that 'wisdom' needs to be defined according to outcomes and processes, not the perception or sagacity of any particular individual guru or sage. Take the *outcome* of the Cold War... the first known example of humanity acquiring a means of massive violence, and then mostly turning away from that precipice. Yes, that means of self-destruction is still with us. But two generations of unprecedented restraint suggest that we have made a little progress in at least one kind of 'wisdom'. That is, when the means of destruction are controlled by a few narrowly selected elite officials on both sides of a simple divide.

But are we ready for a new era, when the dilemmas are nowhere near as simple? In times to come, the worst dangers to civilization may not come from clearly identifiable and accountable adversaries—who want to win an explicit, set-piece competition—as much as from a *general democratization of the means to do harm*. New technologies, distributed by the Internet and effectuated by cheaply affordable tools, will offer increasing numbers of angry people access to modalities of destructive power—means that will be used because of justified grievance, avarice, indignant anger, or simply because they are there.

#### THE RETRO PRESCRIPTION—RENUNCIATION

Faced with onrushing technologies in biotech, nanotech, artificial intelligence, and so on, some bright people—like Bill Joy, former chief scientist of Sun Computers—see little hope for survival of a vigorously open society. You may have read Joy's unhappy manifesto in *Wired Magazine*,<sup>3</sup> in which he quoted the Unabomber (of all people), in support of a proposal that is both ancient and new—that our sole hope for survival may be to renounce, squelch, or relinquish several classes of technological progress.

This notion of *renunciation* has gained credence all across the political and philosophical map, especially at the farther wings of both right and left. Take the novels and pronouncements of Margaret Atwood, whose fundamental plot premises seem almost identical to those of Michael Crichton, despite their differences over superficial politics. Both authors routinely express worry that often spills into outright loathing for the overweening arrogance of hubristic technological innovators who just cannot leave nature well enough alone.

At the other end of the left-right spectrum stands Francis Fukuyama, who is Bernard L. Schwartz Professor of International Political Economy at the Paul H. Nitze School of Advanced International Studies of Johns Hopkins University. Dr Fukuyama's best-known book, *The End of History and the Last Man* (1992) triumphantly viewed the collapse of communism as likely to be the final stirring event worthy of major chronicling by historians. From that point on, we would see liberal democracy bloom as the sole path for human societies, without significant competition or incident. No more "interesting times".<sup>4</sup> But this sanguine view did not last, as

<sup>3</sup> "Why the future doesn't need us." *Wired Magazine*, Issue 8.04, April 2000.

<sup>4</sup> While my description of *The End of History* oversimplifies a bit, one can wish that predictions in social science were as well tracked for credibility as they are in physics. Back in 1986, at the height of Reagan-era confrontations, I forecast an approaching fall of the Berlin Wall, to be followed by several decades of tense confrontation with "one or another branch of macho culture, probably Islamic."

Fukuyama began to see potentially calamitous ‘history’ in the disruptive effects of new technology. As a Bush Administration ‘court intellectual’ and a member of the President’s Council on Bioethics, he now condemns a wide range of biological science as disruptive and even immoral. People cannot, according to Fukuyama, be trusted to make good decisions about the use of—for example—genetic therapy. Human ‘improvability’ is so perilous a concept that it should be dismissed, almost across-the-board. In *Our Posthuman Future: Consequences of the Biotechnology Revolution* (2002), Fukuyama prescribes paternalistic government industry panels to control or ban whole avenues of scientific investigation, doling out those advances that are deemed suitable.

You may surmise that I am dubious. For one thing, shall we enforce this research ban worldwide? Can such tools be squelched forever? From *elites*, as well as the masses? If so, *how*?

Although some of the failure modes mentioned by Bill Joy, Ralph Peters, Francis Fukuyama and the brightest renunciators seem plausible and worth investigating, it’s hard to grasp how we can accomplish anything by becoming neo-Luddites. Laws that seek to limit technological advancement will certainly be disobeyed by groups that simmer at the social extreme, where the worst dangers lie. Even if ferocious repression is enacted—perhaps augmented with near-omniscient and universal surveillance—this will not prevent exploration and exploitation of such technologies by social elites. (Corporate, governmental, aristocratic, criminal, foreign... choose your own favorite bogeymen of unaccountable power.) For years, I have defied renunciators to cite *one example*, amid all of human history, when the mighty allowed such a thing to happen. Especially when they plausibly stood to benefit from something new.

While unable to answer that challenge, some renunciators have countered that all of the new mega-technologies—including biotech and nanotechnology—may be *best* utilized and advanced if control is restricted to knowing elites, even in secret. With so much at stake, should not the best and brightest make decisions for the good of all? Indeed, in fairness, I should concede that the one historical example I gave earlier—that of nuclear weaponry—lends a little support to this notion. Certainly, in that case, one thing that helped to save us was the *limited number of decision-makers* who could launch calamitous war.

Still, weren’t the political *processes* constantly under public scrutiny, during that era? Weren’t those leaders supervised by the public, at least on one side? Moreover, decisions about atom bombs were not corrupted very much by matters of self-interest. (Howard Hughes did not seek to own and use a private nuclear arsenal.) But self-interest will certainly influence controlling elites when they weigh the vast benefits and potential costs of biotech and nanotechnology.

Besides, isn’t elitist secrecy *precisely* the error-generating mode that Crichton, Atwood and so many others portray so vividly, time and again, while preaching against technological hubris? History is rife with examples of delusional cabals of self-assured gentry, telling each other just-so stories while evading any criticism that might reveal flaws in The Plan. By prescribing a return to paternalism—control by elites who remain aloof and unaccountable—aren’t renunciators ultimately proposing the very scenario that everybody—rightfully—fears most?

Perhaps this is one reason why the renunciators—while wordy and specific about possible failure modes—are seldom very clear on which controlling entities should do the dirty work of squelching technological progress. Or how this relinquishment could be enforced across the

board. Indeed, supporters can point to *no historical examples* when knowledge-suppression led to anything but greater human suffering. No proposal that's been offered so far even addresses the core issue of how to prevent some group of elites from cheating. Perhaps *all* elites.

In effect, only the vast pool of normal people would be excluded, eliminating their myriad eyes, ears and prefrontal lobes from civilization's error-detecting network.

Above all, renunciation seems a rather desperate measure, completely out of character with this optimistic, pragmatic, can-do culture.

#### THE SELDOM-MENTIONED ALTERNATIVE—RECIPROCAL ACCOUNTABILITY

And yet, despite all this criticism, I am actually much more approving of Joy, Atwood, Fukuyama et al., than some might expect. In *The Transparent Society*, I speak well of social critics who shout when they see potential danger along the road.

In a world of rapid change, we can only maximize the benefits of scientific advancement—and minimize inevitable harm—by using the great tools of openness and accountability. Above all, by acknowledging that vigorous *criticism is the only known antidote to error*. This collective version of 'wisdom' is what almost surely has saved us so far. It bears little or no resemblance to the kind of individual sagacity that we are used to associating with priests, gurus, and grandmothers... but it is also less dependent upon perfection. Less prone to catastrophe when the anointed Center of Wisdom makes some inevitable blunder.

Hence, in fact, I find fretful worry-mongers invigorating! Their very presence helps progress along by challenging the gung-ho enthusiasts. It's a process called reciprocal accountability. *Without bright grouches, eager to point at potential failure modes, we might really be in the kind of danger that they claim we are*. Ironically, it is an open society—where the sourpuss Cassandras are well heard—that is unlikely to need renunciation, or the draconian styles of paternalism they prescribe.

Oh, I see the renunciators' general point. If society remains as stupid as some people think it is—or even if it is as smart as *I* think it is, but gets no smarter—then nothing that folks do or plan at a thousand well-intentioned futurist conferences will achieve very much. No more than delaying the inevitable.

In that case, we'll finally have the answer to an ongoing mystery of science—why there's been no genuine sign of extraterrestrial civilization amid the stars.<sup>5</sup> The answer will be simple. Whenever technological culture is tried, it always destroys itself. That possibility lurks, forever, in the corner of our eye, reminding us what's at stake.

On the other hand, I see every reason to believe we have a chance to disprove that dour worry. As members of an open and questioning civilization—one that uses reciprocal accountability to find and probe every possible failure-mode—we may be uniquely equipped to handle the challenges ahead.

Anyway, believing that is a lot more fun.

---

<sup>5</sup> For more on this quandary and its implications, see: <http://www.davidbrin.com/sciencearticles.html>.

**THE UPSIDE SCENARIO—THE SINGULARITY**

We've heard from the gloomy renunciators. Let's look at another future. The scenario of those who – literally – believe the sky's the limit. Among many of our greatest thinkers, there is a thought going around—a new 'meme' if you will—that says we're poised for take-off. The idea I'm referring to is that of a coming *Technological Singularity*.

Science fiction author Vernor Vinge has been touted as a chief popularizer of this notion, though it has been around, in many forms, for generations. More recently, Ray Kurzweil's book *The Singularity is Near* argues that our scientific competence and technologically-empowered creativity will soon skyrocket, propelling humanity into an entirely new age.

Call it a modern, high-tech version of Teilhard De Chardin's *noösphere apotheosis*—an approaching time when humanity may move, dramatically and decisively, to a higher state of awareness or being. Only, instead of achieving this transcendence through meditation, good works or nobility of spirit, the idea this time is that we may use an accelerating cycle of education, creativity and computer-mediated knowledge to achieve intelligent mastery over both the environment and our own primitive drives.

In other words, first taking control over Brahma's 'wheel of life', then learning to steer it wherever we choose. What else would you call it...

- When we start using nanotechnology to repair bodies at the cellular level?
- When catching up on the latest research is a mere matter of *desiring* information, whereupon autonomous software agents deliver it to you, as quickly and easily as your arm now moves wherever you wish it to?
- When on-demand production becomes so trivial that wealth and poverty become almost meaningless terms?
- When the virtual reality experience—say visiting a faraway planet—gets hard to distinguish from the real thing?
- When each of us can have as many 'servants'—either robotic or software-based—as we like, as loyal as your own right hand?
- When augmented human intelligence will soar and—trading insights with one another at light speed—helping us attain entirely new levels of thought?

Of course, it is worth pondering how this 'singularity' notion compares to the long tradition of contemplations about *human transcendence*. Indeed, the idea of rising to another plane of existence is hardly new! It makes up one of the most consistent themes in cultural history, as though arising from our basic natures.

Indeed, many *opponents* of science and technology clutch their own images of messianic transformation, images that—if truth be told—share many emotional currents with the tech-heavy version, even if they disagree over the means to achieve transformation. Throughout history, most of these musings dwelled upon the spiritual path, that human beings might achieve a higher state through prayer, moral behavior, mental discipline, or by reciting correct incantations. Perhaps because prayer and incantations were the only means available.

In the last century, an intellectual tradition that might be called 'techno-transcendentalism' added a fifth track. The notion that a new level of existence, or a more appealing state of being, might be achieved by means of knowledge and skill.

But which kinds of knowledge and skill? Depending on the era you happen to live in,

techno-transcendentalism has shifted from one fad to another, pinning fervent hopes upon the scientific flavor of the week. For example, a hundred years ago, *Marxists* and *Freudians* wove complex models of human society—or mind—predicting that rational application of these models and rules would result in far higher levels of general happiness.<sup>6</sup> Subsequently, with popular news about advances in agriculture and evolutionary biology, some groups grew captivated by *eugenics*—the allure of improving the human animal. On occasion, this resulted in misguided and even horrendous consequences. Yet, this recurring dream has lately revived in new forms, with the promise of genetic engineering and neurotechnology.

Enthusiasts for *nuclear power* in the 1950's promised energy too cheap to meter. Some of the same passion was seen in a widespread enthusiasm for *space colonies*, in the 1970s and 80s, and in today's ongoing *cyber-transcendentalism*, which promises ultimate freedom and privacy for everyone, if only we just start encrypting every Internet message, using anonymity online to perfectly mask the frail beings who are actually typing at a real keyboard. Over the long run, some hold out hope that human minds will be able to *download* into computers or the vast new frontier of mid-21st century cyberspace, freeing individuals of any remaining slavery to our crude and fallible organic bodies.

This long tradition—of bright people pouring faith and enthusiasm into transcendental dreams—tells us a lot about one aspect of our nature, a trait that crosses all cultures and all centuries. Quite often, this zealotry is accompanied by disdain for contemporary society—a belief that some kind of salvation can only be achieved outside of the normal cultural network... a network that is often unkind to bright philosophers—and nerds. Seldom is it ever discussed how much these enthusiasts have in common—at least emotionally—with believers in older, more traditional styles of apotheosis, styles that emphasize methods that are more purely mental or spiritual.

We need to keep this long history in mind, as we discuss the latest phase: a belief in the ultimately favorable effects of an exponential increase in the ability of our calculating engines. That their accelerating power of computation will offer commensurately profound magnifications of our knowledge and power. Our wisdom and happiness.

The challenge that I have repeatedly laid down is this: “Name one example, in all of history, when these beliefs actually bore fruit. In light of all the other generations who felt sure of *their* own transforming notion, should you not approach your newfangled variety with some caution... and maybe a little doubt?”

#### IT MAY BE JUST A DREAM

Are *both* the singularity believers and the renunciators getting a bit carried away? Let's take that notion of *doubt* and give it some steam. Maybe all this talk of *dramatic transformation*, within our lifetimes, is just like those earlier episodes: based more on wishful (or fearful) thinking than upon anything provable or pragmatic.

Take Jonathan Huebner, a physicist who works at the Pentagon's Naval Air Warfare Center in China Lake, California. Questioning the whole notion of *accelerating technical progress*, he

<sup>6</sup> And more quasi-religious social-political mythologies followed, from the incantations of Ayn Rand to MaoZedong. All of them crafting “logical chains of cause and effect” that forecast utter human transformation, by political (as opposed to spiritual or technical) means.

studied the rate of “significant innovations per person”. Using as his sourcebook *The History of Science and Technology*, Huebner concluded that the rate of innovation peaked in 1873 and has been declining ever since. In fact, our current rate of innovation—which Huebner puts at seven important technological developments per billion people per year—is about the same as it was in 1600. By 2024, it will have slumped to the same level as it was in the Dark Ages, around 800 AD. “The number of advances wasn’t increasing exponentially, I hadn’t seen as many as I had expected.”

Huebner offers two possible explanations: economics and the size of the human brain. Either it’s just not worth pursuing certain innovations since they won’t pay off—one reason why space exploration has all but ground to a halt—or we already know most of what we can know, and so discovering new things is becoming increasingly difficult.

Ben Jones, of Northwestern University in Illinois, agrees with Huebner’s overall findings, comparing the problem to that of the Red Queen in *Through the Looking Glass*: we have to run faster and faster just to stay in the same place. Jones differs, however, as to *why* this happened. His first theory is that early innovators plucked the easiest-to-reach ideas, or ‘low-hanging fruit’, so later ones have to struggle to crack the harder problems. Or it may be that the massive accumulation of knowledge means that innovators have to stay in education longer to learn enough to invent something new and, as a result, less of their active life is spent innovating. “I’ve noticed that Nobel Prize winners are getting older,” he says.

In fact, it is easy to pick away at these four arguments by Huebner and Jones.<sup>7</sup> For example, it is only natural for innovations and breakthroughs to seem less obvious or apparent to the naked eye, as we have zoomed many of our research efforts down to the level of the quantum and out to the edges of the cosmos. In biology, only a few steps—like completion of the Human Genome Project—get explicit attention as ‘breakthroughs’. Such milestones are hard to track in a field that is fundamentally so complex and murky. But that does not mean biological advances aren’t either rapid or, overall, truly substantial. Moreover, while many researchers seem to gain their honors at an older age, is that not partly a reflexion of the fact that lifespans have improved, and fewer die off before getting consideration for prizes?

Oh, there is something to be said for the singularity-doubters. Indeed, even in the 1930s, there were some famous science fiction stories that prophesied a *slowdown in progress*, following a simple chain of logic. Because progress would *seem* to be its own worst enemy. As more becomes known, specialists in each field would have to absorb more and more about less and less—or about ever narrowing fields of endeavor—in order to advance knowledge by the tiniest increments. When I was a student at Caltech in the 1960s, we undergraduates discussed this problem at worried length. For example, every year the sheer size, on library shelves, of “Chemical Abstracts” grew dauntingly larger and more difficult for any individual to scan for relevant papers.

And yet, over subsequent decades, this trend never seemed to become the calamity we expected. In part, because Chemical Abstracts and its cousins have—in fact—*vanished* from library shelves altogether! The library space problem was solved by simply putting every

---

<sup>7</sup> For a detailed response to Huebner’s anti-innovation argument, see *Review of “A Possible Declining Trend for Worldwide Innovation” by Jonathan Huebner*, published by John Smart in the September 2005 issue of *Technological Forecasting and Social Change* (<http://accelerating.org/articles/huebnerinnovation.html>).



abstract on the Web. Certainly, literature searches—for relevant work in even distantly related fields—now take place faster and more efficiently than ever before, especially with the use of software agents and assistants that should grow even more effective in years to come.

That counter-force certainly has been impressive. Still, my own bias leans toward *another* trend that seems to have helped forestall a productivity collapse in science. This one (I will admit) is totally subjective. And yet, in my experience, it has seemed even more important than advances in online search technology. For it has seemed to me that the best and brightest scientists are *getting smarter*, even as the problems they address become more complex.

I cannot back this up with statistics or analyses. Only with my observation that many of the professors and investigators that I have known during my life now seem much livelier, more open-minded and more interested in fields outside their own—even as they advance in years—than they were when I first met them. In some cases, decades ago. Physicists seem to be more interested in biology, biologists in astronomy, engineers in cybernetics, and so on, than used to be the case. This seems in stark contrast to what you would expect, if specialties were steadily narrowing. But it is compatible with the notion that *culture* may heavily influence our ability to be creative. And a culture that loosens hoary old assumptions and guild boundaries may be one that's in the process of freeing-up mental resources, rather than shutting them down.

In fact, this trend—toward overcoming standard categories of discipline—is being fostered deliberately in many places. For example, the new Sixth College of the University of California at San Diego, whose official institutional mission is to “bridge the arts and sciences”, drives a nail in the coffin of C.P. Snow's old concept that the “two cultures” can never meet. Never before have there been so many collaborative efforts between tech-savvy artists and technologists who appreciate the aesthetic and creative sides of life.<sup>8</sup>

What Huebner and Jones appear to miss is that complex obstacles tend best to be overcome by complex entities. Even if Einstein and others picked all the low hanging fruit within reach to individuals, that does not prevent groups—institutions and teams and entrepreneurial startups—from forming collaborative human pyramids to go after goodies that are higher in the tree. Especially when those pyramids and teams include new kinds of members, software agents and search methodologies, worldwide associative networks and even open-source participation by interested amateurs. Or when a myriad fields of endeavor see their loci of creativity get dispersed onto a multitude of inexpensive desktops, the way software has been.<sup>9</sup>

Dutch-American economic historian Joel Mokyr, in *The Lever of Riches* and *The Gifts of Athena*, supports this progressive view that we are indeed doing something right, something that makes our liberal-democratic civilization uniquely able to generate continuous progress. Mokyr believes that, since the 18th-century Enlightenment, a new factor has entered the human

<sup>8</sup> The *Exorarium Project* proposes to achieve all this and more, by inviting both museum visitors and online participants to enter a unique learning environment. Combining state-of-the-art simulation and visualization systems, plus the very best ideas from astronomy, physics, chemistry, and ecology, the Exorarium will empower users to create vividly plausible extraterrestrials and then test them in realistic first contact scenarios. <http://www.exorarium.com/>.

<sup>9</sup> For a rather intense look at how ‘truth’ is determined in science, democracy, courts and markets, see the lead article in the American Bar Association's Journal on Dispute Resolution (Ohio State University), v.15, N.3, pp 597-618, Aug. 2000, “Disputation Arenas: Harnessing Conflict and Competition for Society's Benefit” or at: <http://www.davidbrin.com/disputationarticle1.html>.

equation: the accumulation of and a free market in knowledge. As Mokyr puts it, we no longer behead people for saying the wrong thing—we listen to them. This ‘social knowledge’ is progressive because it allows ideas to be tested and the most effective to survive; it is embodied in institutions, which, unlike individuals, can rise above our animal natures.

But Mokyr does worry that, though a society may progress, *human nature does not*. “Our aggressive, tribal nature is hard-wired, unreformed and unreformable. Individually we are animals and, as animals, incapable of progress.” The trick is to cage these animal natures in effective institutions: education, the law, government. But these can go wrong. “The thing that scares me,” he says, “is that these institutions can misfire.”

While I do not use words such as “caged”, I must agree that Mokyr captures the essential point of our recent, brief experiment with the Enlightenment: John Locke’s rejection of romantic oversimplification in favor of pragmatic institutions that work flexibly to maximize the effectiveness of our better efforts—the angels of our nature—enabling our creative forces to mutually reinforce. Meanwhile, those same institutions and processes would thwart our ‘devils’—the always-present human tendency towards self-delusion and cheating. Of course, human nature strives against these constraints. Self-deluders and cheaters are constantly trying to make up excuses to bypass the Enlightenment covenant and benefit by making these institutions less effective. Nothing is more likely to ensure the failure of any singularity than if we allow this to happen.

But then, swiveling the other way, what if it soon becomes possible not only to preserve and advance those creative enlightenment institutions, but also to do what Mokyr calls impossible? What if we actually can *improve human nature*?

Suppose the human *components* of societies and institutions can also be made better, even by a little bit? I have contended that this is already happening, on a modest scale. Imagine the effects of even a small upward-ratcheting in general human intelligence, whether inherent or just functional, by means of anything from education to smart drugs to technologically-assisted senses to new methods of self-conditioning.

It might not take much of an increase in effective human intelligence for markets and science and democracy etc. to start working much better than they already do. Certainly, this is one of the factors that singularity aficionados are counting on.

What we are left with is an image that belies the simple and pure notion of a singularity curve... one that rises inexorably skyward, as a simple mathematical function, with knowledge and skill perpetually leveraging against itself, as if ordained by natural law. Even the most widely touted example of this kind of curve, Moore’s Law—which successfully modeled the rapid increase of computational power available at plummeting cost—has never been anything like a smooth phenomenon. Crucial and timely decisions—some of them pure happenstance—saved Moore’s Law on many occasions from collision with either technological barriers or cruel market forces.

True, we seem to have been lucky, so far. Cybernetics and education and a myriad other factors have helped to overcome the specialization trap. But as we have seen in this section, past success is no guarantee of future behavior. Those who foresee upward curves continuing *ad infinitum*, almost as a matter of faith, are no better grounded than other transcendentalists, who confidently predicted other rapturist fulfillments in their own times.

## THE DAUNTING TASK OF CROSSING A MINEFIELD

Having said all of the above, let me hasten to add that I believe in the high likelihood of a coming singularity!

I believe in it because the alternatives are too awful to accept. Because, as we discussed before, the means of mass destruction, from A-bombs to germ warfare, are ‘democratizing’—spreading so rapidly among nations, groups, and individuals—that we had better see a rapid expansion in sanity and wisdom, or else we’re all doomed.

Indeed, bucking the utterly prevalent cliché of cynicism, I suggest that strong evidence does indicate some cause for tentative optimism. An upward trend is already well in place. Overall levels of education, knowledge and sagacity in Western Civilization—and its constituent citizenry—have never been higher, and these levels may continue to improve, rapidly, in the coming century. Possibly enough to rule out some of the most prevalent images of failure that we have grown up with. For example, we will not see a future that resembles *Blade Runner*, or any other cyberpunk dystopia. Such worlds—where massive technology is unmatched by improved wisdom or accountability—will simply not be able to sustain themselves.

The options before us appear to fall into four broad categories:

1. Self-destruction. Immolation or desolation or mass-death. Or ecological suicide. Or social collapse. Name your favorite poison. Followed by a long era when our few successors (if any) look back upon us with envy. For a wonderfully depressing and informative look at this option, see Jared Diamond’s *Collapse: How Societies Choose to Fail or Succeed*. (Note that Diamond restricts himself to ecological disasters that resonate with civilization-failures of the past; thus he only touches on the range of possible catastrophe modes.) We are used to imagining self-destruction happening as a result of mistakes by ruling elites. But in this article we have explored how it also could happen if society enters an age of universal democratization of the means of destruction—or, as Thomas Friedman puts it, “the super-empowerment of the angry young man”—without accompanying advances in social maturity and general wisdom.

2. Achieve some form of ‘Positive Singularity’—or at least a phase shift to a higher and more knowledgeable society (one that may have problems of its own that we can’t imagine.) Positive singularities would, in general, offer normal human beings every opportunity to participate in spectacular advances, experiencing voluntary, dramatic self-improvement, without anything being compulsory... or too much of a betrayal of the core values of decency we share.

3. Then there is the ‘Negative Singularity’—a version of self-destruction in which a skyrocket of technological progress *does* occur, but in ways that members of our generation would find unpalatable. Specific scenarios that fall into this category might include being abused by new, super-intelligent successors (as in *Terminator* or *The Matrix*), or simply being left behind by super-entities that pat us on the head and move on to great things that we can never understand. Even the softest and most benign version of such a ‘Negative Singularity’ is perceived as loathsome by some perceptive renunciators, like Bill Joy, who take a dour view of the prospect that humans may become a less-than-pinnacle form of life on Planet Earth.<sup>10</sup>

<sup>10</sup> In other places, I discuss various proposed ways to deal with the Problem of Loyalty, in some future age when machine intelligences might vastly excel the capabilities of mere organic brains. Older

4. Finally, there is the ultimate outcome that is implicit in every renunciation scenario: Retreat into some more traditional form of human society, like those that maintained static sameness under pyramidal hierarchies of control for at least four millennia. One that quashes the technologies that might lead to results 1 or 2 or 3. With four thousand years of experience at this process, hyper-conservative hierarchies could probably manage this agreeable task, if we give them the power. That is, they could do it for a while.

When the various paths<sup>11</sup> are laid out in this way, it seems to be a daunting future that we face. Perhaps an era when all of human destiny will be decided. Certainly not one that's devoid of history. For a somewhat similar, though more detailed, examination of these paths, the reader might pick up Joel Garreau's fine book, *Radical Evolution*. It takes a good look at two extreme scenarios for the future — "Heaven" and Hell" — then posits a third — "Prevail" — as the one that rings most true.

So, which of these outcomes seem plausible? First off, despite the fact that it may look admirable and tempting to many, I have to express doubt that outcome 4 could succeed over an extended period. Yes, it resonates with the lurking tone that each of us feels inside, inherited from countless millennia of feudalism and unquestioning fealty to hierarchies, a tone that today is reflected in many popular fantasy stories and films. Even though we have been raised to hold some elites in suspicion, there is a remarkable tendency for each of us to turn a blind eye to other elites—or favorites—and to rationalize that *those* would rule wisely.

Certainly, the quasi-Confucian social pattern that is being pursued by the formerly Communist rulers of China seems to be an assertive, bold and innovative approach to updating authoritarian rule, incorporating many of the efficiencies of both capitalism and meritocracy.<sup>12</sup> This determined effort suggests that an updated and modernized version of hierarchism might succeed at suppressing whatever is worrisome, while allowing progress that's been properly vetted. It is also, manifestly, a rejection of the Enlightenment and everything that it stands for, including John Locke's wager that *processes of regulated but mostly free human interaction* can solve problems better than elite decision-making castes.

In fact, we have already seen, in just this one article, more than enough reasons to understand why retreat simply cannot work over the long run. Human nature ensures that there

---

proposals (e.g. Asimov's "laws of robotics") almost surely cannot work. It remains completely unknown whether humans can 'go along for the ride' by using cyborg enhancements or 'linking' with external processors. In the long run, I suggest that we might deal with this in the same way that all prior generations created new (and sometimes superior) beings without much shame or fear. By raising them to think of themselves as human beings, with our same values and goals. In other words, as our children. (See: <http://www.davidbrin.com/lungfish1.html>).

<sup>11</sup> Of course, there are other possibilities, indeed many others, or I would not be worth my salt as a science fiction author or futurist. Among the more sophomorically entertaining possibilities is the one positing that we all live in a simulation, in some already post-singularity "context" such as a vast computer. The range is limitless. But these four categories seem to lay down the starkness of our challenge: to become wise, or see everything fail within a single lifespan.

<sup>12</sup> This endeavor has been based upon earlier Asian success stories, in Japan and in Singapore, extrapolating from their mistakes. Most notable has been an apparent willingness to learn pragmatic lessons, to incorporate limited levels of criticism and democracy, accepting their value as error-correction mechanisms — while limiting their effectiveness as threats to hierarchical rule. One might imagine that this tightrope act must fail, once universal education rises beyond a certain point. But that is only a hypothesis. Certainly the neo-Confucians can point to the sweep of history, supporting their wager.

can never be successful rule by serene and dispassionately wise philosopher kings. That approach had its fair trial—at least forty centuries—and by almost any metric, it failed.

As for the other three roads, well, there is simply no way that anyone—from the most enthusiastic, “extropian” utopian-transcendentalists to the most skeptical and pessimistic doomsayers—can prove that one path is more likely than the others. (How can models, created within an earlier, cruder system, properly simulate and predict the behavior of a later and vastly more complex system?) All we can do is try to understand which processes may increase our odds of achieving better outcomes. More *robust* outcomes. These processes will almost certainly be as much *social* as technological. They will, to a large degree, depend upon improving our powers of *error-avoidance*.

My contention—running contrary to many prescriptions from both left and right—is that we should trust Locke a while longer. This civilization already has in place a number of unique methods for dealing with rapid change. If we pay close attention to how these methods work, they might be improved dramatically, perhaps enough to let us cope, and even thrive. Moreover, the *least* helpful modification would appear to be the one thing that the professional castes tell us we need—an increase in paternalistic control.<sup>13</sup>

In fact, when you look at our present culture from an historical perspective, it is already profoundly anomalous in its emphasis upon individualism, progress, and above all, *suspicion of authority* (SOA). These themes were actively and vigorously repressed in a vast majority of human cultures, because they threatened the stable equilibrium upon which ruling classes always depended. In Western civilization—by way of contrast—it would seem that every mass-media work of popular culture, from movies to novels to songs, promotes SOA as a central human value.<sup>14</sup> This may, indeed, be the most unique thing about our culture, even more than our wealth and technological prowess.

<sup>13</sup> See my essay on “Beleaguered Professionals vs. Disempowered Citizens” about a looming 21st century power struggle between average people and the sincere, skilled professionals who are paid to protect us: <http://www.amazon.com/gp/product/B000BY2PRQ/002-1071896-8741633>. In a related context, a ‘futurist essay’ points out a rather unnoticed aspect of the tragedy of 9/11/2001—that citizens themselves were most effective in our civilization’s defense. The only actions that actually saved lives and thwarted terrorism on that awful day were taken amid rapid, *ad hoc* decisions made by private individuals, reacting with both resiliency and initiative—our finest traits—and armed with the very same new technologies that dour pundits say will enslave us. Could this point to a trend for the 21st century, reversing what we’ve seen throughout the 20th—the ever-growing dependency on professionals to protect and guide and watch over us? See: [http://www.futurist.com/portal/future\\_trends/david\\_brin\\_empowerment.htm](http://www.futurist.com/portal/future_trends/david_brin_empowerment.htm).

<sup>14</sup> Take the essential difference between moderate members of the two major American political parties. This difference boils down to *which* elites you accuse of seeking to accumulate too much authority. A decent Republican fears snooty academics, ideologues, and faceless bureaucrats seeking to become paternalistic Big Brothers. A decent Democrat looks with worried eyes toward conspiratorial power grabs by conniving aristocrats, faceless corporations, and religious fanatics. (A decent Libertarian picks two from Column A and two from Column B!) I have my own opinions about which of these elites are presently most dangerous. (Hint: it is the same one that dominated most other urban cultures, for four thousand years.) But the startling irony, that is never discussed, is how much in common these fears really share. And the fact that—indeed—every one of them is right to worry. In fact, only universal SOA makes any sense. Instead of an ideologically blinkered focus on just one patch of horizon, should we not agree to watch *all* directions where tyranny or rationalized stupidity might arise? Again, reciprocal accountability appears to be the only possible solution.

Although we are proud of the resulting society—one that encourages eccentricity, appreciation of diversity, social mobility, and scientific progress—we have no right, as yet, to claim that this new way of doing things is especially sane or obvious. Many in other parts of the world consider Westerners to be quite mad! And with some reason. Indeed, only time will tell who is right about that. For example, if we take the suspicion of authority ethos to its extreme, and start paranoically mistrusting even our best institutions—as was the case with Oklahoma City bomber Timothy McVeigh—then it is quite possible that Western civilization may fly apart before ever achieving its vaunted aims, and lead rapidly to some of the many ways that we might achieve outcome 1.

Certainly, a positive singularity (outcome 2) cannot happen if only centrifugal forces operate and there are no compensating centripetal virtues to keep us together as a society of mutually respectful sovereign citizens.

Above all (as I point out in *The Transparent Society*), our greatest innovations, the accountability arenas<sup>15</sup> wherein issues of importance get decided—science, justice, democracy and free markets—are not arbitrary, nor are they based on whim or ideology. They all depend upon adversaries competing on specially designed playing fields, with hard-learned arrangements put in place to prevent the kinds of cheating that normally prevail whenever human beings are involved. Above all, science, justice, democracy, and free markets depend on the mutual accountability that comes from open flows of information.

Secrecy is the enemy that destroys each of them, and it could easily spread like an infection to spoil our frail renaissance.

#### THE BEST METHODS OF ERROR-AVOIDANCE

Clearly, our urgent goal is to find (and then avoid) a wide range of quicksand pits—potential failure modes—as we charge headlong into the future. At risk of repeating an oversimplification, we do this in two ways. One method is *anticipation*. The other is *resiliency*.

The first of these uses the famous prefrontal lobes—our most recent, and most spooky, neural organs—to peer ahead, perform *gedankenexperiments*, forecast problems, make models and devise countermeasures in advance. *Anticipation* can either be a lifesaver, or one of our most colorful paths to self-deception and delusion.<sup>16</sup>

The other approach—*resiliency*—involves establishing robust systems, reaction sets, tools and distributed strengths that can deal with almost any problem as it arises—even surprising problems the vaunted prefrontal lobes never imagined.

Now, of course, these two methods are compatible, even complementary. We have a better computer industry, overall, because part of it is centered in Boston and part in California, where different corporate cultures reign. Companies acculturated with a ‘northeast mentality’ try to

---

<sup>15</sup> For a rather intense look at how “truth” is determined in science, democracy, courts and markets, see the lead article in the American Bar Association’s Journal on Dispute Resolution (Ohio State University), v.15, N.3, pp 597-618, Aug. 2000, “Disputation Arenas: Harnessing Conflict and Competition for Society’s Benefit” or at: <http://www.davidbrin.com/disputationarticle1.html>.

<sup>16</sup> I say this as a prime practitioner of the art of anticipation, both in nonfiction and in fiction. Every futurist and novelist deals in creating convincing illusions of prescience, though at times these illusions can be helpful.

make perfect products. Employees stay in the same company, perhaps for decades. They feel responsible. They get the bugs out before releasing and shipping. These are people you want designing a banking program, or a defense radar, because we can't afford a lot of errors in even the beta version, let alone the nation's ATM machines! On the other hand, people who work in Silicon Valley seem to think almost like another species. They cry, "Let's get it out the door! Innovate first and catch the glitches later! Our customers will tell us what parts of the product to fix on the fly. They want the latest thing and to hell with perfection." Today's Internet arose from that kind of creative ferment, adapting quickly to emergent properties of a system that turned out to be far more complex and fertile than its original designers anticipated. Indeed, their greatest claim to fame comes from having anticipated that unknown opportunities might emerge!

Sometimes the best kind of planning involves leaving room for the unknown. This can be hard, especially when your duty is to prepare against potential failure modes that could harm or destroy a great nation. Government and military culture have always been anticipatory, seeking to analyze potential near-term threats and coming up with detailed plans to stymie them. This resulted in incremental approaches to thinking about the future. One classic cliché holds that generals are always planning to fight a modified version of the last war. History shows that underdogs—those who lost the last campaign or who bear a bitter grudge—often turn to innovative or resilient new strategies, while those who were recently successful are in grave danger of getting mired in irrelevant solutions from the past, often with disastrous consequences.<sup>17</sup>

At the opposite extreme is the genre of science fiction, whose attempts to anticipate the future are—when done well—part of a dance of resiliency. Whenever a future seems to gather a consensus around it, as happened to cyberpunk in the late eighties, the brightest SF authors become bored with such a trope and start exploring alternatives. Indeed, boredom could be considered one of the driving forces of ingenious invention, not only in science fiction, but in our rambunctious civilization as a whole.

Speaking as an author of speculative novels, I can tell you that it is wrong to think that science fiction authors try to predict the future. With our emphasis more on resiliency than anticipation, we are more interested in discovering possible failure modes and quicksand pits along the road ahead, than we are in providing a detailed and prophetic travel guide for the future.

Indeed, one could argue that the most powerful kind of science fiction tale is the *self-preventing prophecy*—any story or novel or film that portrays a dark future so vivid, frightening and plausible that millions are stirred to act against the scenario ever coming true. Examples in this noble (if terrifying) genre—which can also encompass visionary works of non-fiction—include *Fail-Safe*, *Brave New World*, *Soylent Green*, *Silent Spring*, *The China Syndrome*, *Das Kapital*, *The Hot Zone*, and greatest of all, George Orwell's *Nineteen Eighty-Four*, now celebrating 60 years of scaring readers half to death. Orwell showed us the pit awaiting any civilization that combines panic with technology and the dark, cynical tradition of tyranny. In so doing, he armed us against that horrible fate. By exploring the shadowy territory of the future

---

<sup>17</sup> It is worth noting that the present US military Officer Corps has tried strenuously to avoid this trap, endeavoring to institute processes of re-evaluation, by which a victorious and superior force actually thinks like one that has been defeated. In other words, with a perpetual eye open to innovation. And yet, despite this new and intelligent spirit of openness, military thinking remains rife with unwarranted assumptions. Almost as many as swarm through practitioners of politics.

with our minds and hearts, we can sometimes uncover failure-modes in time to evade them.

Summing up, this process of *gedanken* or thought experimentation is applicable to both anticipation and resiliency. But it is only most effective when it is engendered *en masse*, in markets and other arenas where open competition among countless well-informed minds can foster the unique synergy that has made our civilization so different from hierarchy-led cultures that came before. A synergy that withers the *bad* notions under criticism, while allowing good ones to combine and multiply.

I cannot guarantee that this scenario will work over the dangerous ground ahead. An open civilization filled with vastly educated, empowered, and fully-knowledgeable citizens may be able to apply the cleansing light of reciprocal accountability so thoroughly that onrushing technologies cannot be horribly abused by *either* secretive elites or disgruntled AYMs (angry young men).

Or else...perhaps...that solution, which brought us so far in the 20th century, will not suffice in the accelerating 21st. Perhaps nothing can work. Maybe this explains the Great Silence, out there among the stars.

What I do know is this. No *other* prescription has even a snowball's chance of working. Open knowledge and reciprocal accountability seem, at least, to be worth betting on. They are the tricks that got us this far, in contrast to 4,000 years of near utter failure by systems of hierarchical command.

Anyone who says that we should suddenly veer back in that direction, down discredited and failure-riven paths of secrecy and hierarchy, should bear a steep burden of proof.

#### VARIETIES OF SINGULARITY EXPERIENCE

All right, what if we do stay on course, and achieve something like the *Positive Singularity*?

There is plenty of room to argue over what *type* would be beneficial or even desirable. For example, might we trade in our bodies—and brains—for successively better models, while retaining a core of humanity...of soul?

If organic humans seem destined to be replaced by artificial beings who are vastly more capable than we souped-up apes, can we design those successors to at least *think of themselves as human*? (This unusual notion is one that I've explored in a few short stories.) In that case, are you so prejudiced that you would begrudge your great-grandchild a body made of silicon, so long as she visits you regularly, tells good jokes, exhibits kindness, and is good to her own kids? Or will they simply move on, sparing a moment to help us come to terms with our genteel obsolescence?

Some people remain big fans of Teilhard de Chardin's apotheosis—the notion that we will all combine into a single macro-entity, almost literally godlike in its knowledge and perception. Physicist Frank Tipler speaks of such a destiny in his book, *The Physics of Immortality*, and Isaac Asimov offered a similar prescription as mankind's long-range goal in *Foundation's Edge*. I have never found this notion particularly appealing—at least in its standard presentation, by which some macro-being simply subsumes all lesser individuals within it, and then proceeds to think deep thoughts. In *Earth*, I talk about a variation on this theme that might be far more palatable, in which we all remain individuals, while at the same time contributing to a new of planetary consciousness. In other words, we could possibly get to have our cake and eat it too.



At the opposite extreme, in *Foundation's Triumph*, my sequel to Asimov's famous universe, I make more explicit something that Isaac had been alluding to all along—the possibility that conservative robots might *dread* human transcendence, and for that reason actively work to prevent a human singularity. Fearing that it could bring us harm. Or enable us to compete with them. Or empower *us* to leave *them* behind.

In any event, the singularity is a fascinating variation on all those other transcendental notions that seem to have bubbled, naturally and spontaneously, out of human nature since before records were kept. Even more than all the others, this one can be rather frustrating at times. After all, a good parent wants the best for his or her children—for them to do and be better. And yet, it can be poignant to imagine them (or perhaps *their* grandchildren) living almost like gods, with nearly omniscient knowledge and perception—and near immortality—taken for granted.

It's tempting to grumble, "Why not *me*? Why can't I be a god, too?"<sup>18</sup>

But then, when has human existence been anything but poignant? Anyway, what is more impressive? To *be* godlike? Or to be natural creatures, products of grunt evolution, who are barely risen from the caves... who nevertheless manage to learn nature's rules, revere them, and then *use* them to create good things, good descendants, good destinies? *Even godlike ones.*

All of our speculations and musings (including this one) may eventually seem amusing and naive to those dazzling descendants. But I hope they will also experience moments of respect, when they look back at us. They may even pause and realize that we were really pretty good...for souped-up cavemen. After all, what miracle could be more impressive than for such flawed creatures as us to *design and sire gods*?

There may be no higher goal. Or any that better typifies arrogant hubris.

Or else...perhaps...the fulfillment of our purpose and the reason for all that pain. To have learned the compassion and wisdom that we'll need, more than anything else, when bright apprentices take over the Master's workroom. Hopefully winning merit and approval, at last, as we resume the process of creation.

### About the author:

David Brin's popular science fiction novels have been translated into more than twenty languages, including *New York Times* best-sellers that won *Hugo*, *Nebula* and other awards. His 1989 ecological thriller, *Earth*, foreshadowed global warming, cyberwarfare, the World Wide Web and Gulf Coast flooding. A 1998 movie, starring Kevin Costner, was loosely adapted from his Campbell Award winner *The Postman*. DC/Wildstorm also released a groundbreaking hardcover graphic novel "*The Life Eaters*". *Kiln People* portrays a coming era when a simple advance in technology allows anyone to achieve the ancient dream of being in two places at once. *Foundation's Triumph* brought a grand finale to Isaac Asimov's famed Foundation Universe.

Brin is also a noted scientist and speaker/consultant who also appears frequently on television, discussing trends in the near future. He serves on advisory committees dealing with

<sup>18</sup> Of course, there are some Singularitarians—true believers in a looming singularity—who expect it to rush upon us so rapidly that even fellows my age (in my fifties) will get to ride the immortality wave. Yeah, right. And they call *me* a dreamer.

subjects as diverse as astronomy and space exploration, SETI, nanotechnology and national defense. His non-fiction book—*The Transparent Society: Will Technology Make Us Choose Between Freedom and Privacy?*—deals with issues of openness, security and liberty in the new wired-age. It won the 2000 *Freedom of Speech Award* of the American Library Association and a prize from the McGannon Foundation for public service in communications.