


# Slate

future tense

## Debating Extreme Human Enhancement

There are some places our "wetware" bodies just can't go.

By Nicholas Agar, Brad Allenby, and Kyle Munkittrick  
Updated Friday, Sept. 16, 2011, at 7:18 AM ET



*This conversation is part of a Future Tense, a partnership between Slate, the New America Foundation, and Arizona State ... on the boundaries between humans and machines and on "Is Our Techno-Human Marriage in Need of Counseling?"*

**From:** Kyle Munkittrick

**To:** Nicholas Agar and Brad Allenby

**Subject:** Should We Use Nanotech, Genetics, Pharmaceuticals, and Augmentations To Go Above and Beyond Our Biology?

Posted Tuesday, Sept. 13, 2011, at 5:38 AM ET

Gentleman, start your exobrains.

Nick, Brad, thank you for joining me in this discussion about the future of human beings. On Sept. 15, Future Tense is hosting an event on the "Techno-Human Condition" in Washington, D.C. As such, it seems appropriate to spend the next few days discussing the convergence of human beings and technology. Brad, you and Daniel Sarewitz address the associated ethical and moral dilemmas your book [The Techno-Human Condition](#); Nick, the title of your most recent book, [Humanity's End: Why We Should Reject Radical Life Enhancement](#), would seem to sum up your perspective pretty neatly. But then again, your previous book, [Liberal Eugenics: In Defense of Human Enhancement](#), suggest your view might not be so clear cut.

With our introductions out of the way, let's talk transhumanism.

First things first: What is transhumanism? Transhumanism as a philosophy argues we can become better than human through technology. Unguided, natural evolution has done all it could hope to do. Transhumanists believe that from here on out, humans

should take up the reins and craft the evolution of our species using nanotech, genetics, pharmaceuticals, and augmentations to go above and beyond our biology.

Transhumanism was everywhere this summer, from the cyborg anti-hero in the video game [Deus Ex: Human Revolution](#) to the ill-fated brain-boosting drug experiments of [Rise of the Planet of the Apes](#). The promise and peril of radical enhancement seems to be weighing on our collective minds.

But the big question is: Do humans want to be transhuman? That is, presuming that genetic engineering, cybernetic augmentation, cognitive enhancement, and the cure for aging are all technologically possible in the next 50, 100, or 200 years, are they something people will want to use? And is enhancement something people *should* want?

Short answer: Yes, absolutely. I think humans do want to transcend biological limitations and become better than our bodies and genetics currently allow.

The long answer is, well, longer and more complex. Were the question "Does the average person want to be transhuman, *right now*?" I would answer, "Probably not." The phrase "human enhancement" conjures [Gattaca](#), Frankenstein's creature, and the social engineering of Huxley's [Brave New World](#). The religious have their hells and their demons, while those of us with a more scientific disposition have our dystopias and Big Brothers. Popular entertainment argues for us to be exceedingly wary of anyone promising bigger brains, longer lives, and stronger bodies. Western Civilization, as a whole, has trained itself to fear the Promethean hubris of stealing our evolutionary fire from the gods of nature.

Or perhaps this isn't fear of technology or progress but instead a recognition of the value of our current human experience. Our caution toward transhumanism is an act of self-preservation against alienation. It's not that we don't want to be transhuman; it's that we know better than to succumb to those temptations. Nick, in [Humanity's End](#), you argued in favor of moderate enhancement, wherein people can be enhanced *up to* the species' typical best but no further. That is, through enhancement, everyone would be about as smart as the smartest people usually are or live to be as old as the oldest people usually live—a species composed of Albert Einsteins, Michael Jordans, and William Shakespeares. Still human, but to the maximum.

Why shouldn't we want to go further? Is the pursuit of thousand-year youthful life spans, of genetic profiles allowing IQs above 200, and of prosthetics for every possible purpose something that we can responsibly want for humanity?

The argument against transhumanism is that limits and restrictions create meaning but are also liberating. Nick, you argue that those with extended life spans would become more fearful of death, because they would have so much more to lose. As a result, people would opt for safe but shallow digital experiences, leading to long, ultimately empty lives. You also argue that those who would be enhanced beyond human capacity would no longer be demonstrably human. Transhumans would be capable of experiences far beyond those of the average human, such that to describe those who are enhanced as human would be a misnomer. The threat there is that those who are human would feel pressure to no longer be human—and, in the end, society may no longer value

humanness itself. In short, the loss of human biological limits is the loss of our humanity.

I don't doubt that moving beyond mere humanity is a fundamental shift in our values. However, I don't see why humanism as a philosophy is the be-all-end-all of secular values. I don't think transcending our biology through enhancement would be any more ethically disorienting than maximizing our biology through enhancement. I find it difficult to believe the argument against radical enhancement for three reasons.

First, I simply don't see the difference extending a person's 70-year life span to 115 or 150. One is human, as you define it, and the other is transhuman. What is it about the first number that makes it less alienating and more desirable?

Second, there is no reason to believe transhumanism will happen at a flick of a switch. It is far more likely that we'll see slow, piecemeal progress in which enhancement first inches toward human biological limits and then, just as slowly, breaks through those limitations into transhuman ability. Why wouldn't a slow, incremental process of enhancement over generations allow for our human ethics and culture to shift and grow with the technology?

Third, and finally, I don't believe human meaning is created by having limits but by breaking limits. We are at our most human when we are being better; humanity was better for our breaking the hold of Earth's gravity to orbit the moon. Humans want to go beyond ourselves. So, I ask, what exempts the limits of human biology from our human drive to exceed them?

I think people want transhumanism, but not without some time to get acquainted with their new selves.

I must say, I am quite excited for this discussion.  
Kyle

---

**From:** Brad Allenby

**To:** Nicholas Agar and Kyle Munkittrick

**Subject:** Even If We Don't Want Enhancement, We Might Not Be Able To Stop It

Posted Tuesday, Sept. 13, 2011, at 12:05 PM ET

---

Kyle, thanks for the start. I look forward to hearing you folks laugh hysterically (but virtually) as you read my responses.

The "What is transhumanism?" question needs to be asked, but that question can lead one off-scent. I'm willing to stipulate for the court that transhumanism is a philosophy of human progress through primarily technological evolution, which—reflecting attitudes that go back to the Enlightenment—some people are for, and some people are against. If you are skeptical of authority, as many Americans are, you probably like some sort of

fairly libertarian approach to enhancement—"it's my body, and I'll enhance it if I want!" If you're more conservative—a touch of Burke, perhaps—you probably are comfortable with the idea that there is a "human nature," and that, all in all, it would be a good idea to respect that. I think both these positions, reflecting a simple system perspective on a complex adaptive system, probably only capture part of the overall gestalt. I'd rather focus on some additional questions that Kyle's piece raises for me.

First, though, I would suggest that readers reacquaint themselves with (or meet for the first time) [the opening scene](#) of the film [2001: A Space Odyssey](#). The apes jitter about the monolith that has appeared among them, then one goes off and plays with some bones. Picking one up, he first plays with it, then starts swinging it, then realizes that it is not a bone; it is a weapon. He begins crushing things with it, including other apes. My point: An ape with a bone is a very different creature than an ape with a bone who knows it's a weapon. The culture and psychology of being human in the world is being connected to such things, whether it's a bone, a '67 Corvette 427 Sting Ray, or Google (if you don't think the Corvette belongs in there, you weren't a male adolescent in the late '60s). In other words, it is at least arguable that to be human is to be transhuman. To Kyle's point: It is not that we're transcending our biology; rather, we're fulfilling our biology. That argument is over, and has been since hunter-gatherer times. So what has changed?

Two things, I think. First, technological change is far more rapid and pervasive than it has been in the past. We've always had technologies that restructured society, culture, economies, and psychology—the steam engine did, railroads did, cars did, airplanes did, and search engines that increasingly substitute for memory do. But depending on how you count, we have five foundational technologies now—nanotech, biotech, robotics, information and communication tech, and applied cognitive science—all of which are not only evolving in interesting and unpredictable ways; they are actually accelerating in their evolution. Moreover, they're doing that against the backdrop of a world in which systems we've always framed as "natural"—the climate, the nitrogen and phosphorous cycles, biology and biodiversity, and others—are increasingly products of human intervention, intentional or not. We are [terraforming](#) everything, from our planet to one another ... and it's all connected, of course.

Critically for our purposes, the human being is more and more becoming a design space. Per our ape friend above, this is not new. But the rate of change and concomitant (and mainly still potential) psychological, cultural, and social dislocations are. So in some ways transhumanism is not, as we'd like to frame it, about us; it's about reality as we know it becoming our design space—including, and especially, that part of reality that we have heretofore ring-fenced as "human." What we are really seeing is the human equivalent of the Great Divergence, that period in history when economic development touched the lives of some cultures and led them to exponential economic growth, while others remained behind. I'm not saying this is a good thing. But it is where we're headed. Already, people in developed countries live almost twice as long as others, similarly human, in some developing countries. That's a pretty profound divergence, and we seem to accept that *en passant*.

So to Kyle's points. First, I agree that humans like to enhance: We buy stuff that enhances us all the time; we couple into media, computers, and distributed cognition networks such as Google with abandon; and we use enhancement drugs pretty freely when we can get them. Students do, athletes do, workers do. And even where some risk is involved, we enhance just for cosmetic purposes. (I live in Scottsdale, Ariz., and sometimes it seems as if half the local GDP out here comes from cosmetic surgery that enhances virtually every part of the human form—or, if there be too much of a particular form, de-enhances it.) And we choose [life-extension enhancements](#) just as freely—that's what vaccines are, that's what a lot of modern medical care is. We're already seriously enhanced over previous generations, and there's little reason to assume we won't continue to extend our lifespan if we can—and handle the resource implications the same way we always have, with blood and iron.

The naive libertarian position on enhancement tends to overlook systemic effects. For instance, would the world really be better off if violent tyrants stayed awake longer and concentrated better? At what point, as Kyle notes, do individual enhancements add up to community-scale shifts, such that if I don't enhance, I'm the "new subnormal"? There is, of course, the other popular position: that there is a "human nature" that is core to being human, and that "human nature" cannot, and should not, be subject to human design. The trick here is, who gets to define "human nature"? Generally, the folks who use this argument define it very broadly, often based on specific, usually conservative, religious traditions. But, just as in the case of naive libertarianism, this "human nature" position raises difficulties: Since people seem to want to enhance, what I am necessarily arguing if I take the human-nature posture is that authoritarian force—the state, the church—must stifle research on or deployment of enhancements that people would otherwise want to use. And the justification usually implicitly or explicitly rests on my worldview, as opposed to other worldviews that might allow, or even embrace, such enhancement.

I think, however, there's a more basic question: Even if we don't like enhancement, can we stop it? If a particular technology gives a nation or a culture an advantage, won't it over time be developed anyway? Adam Smith famously created blue-collar division of labor with his pin factory. Railroads and other massive industrial capitalism firms created white-collar division of labor with their size and complexity. Human enhancement is simply extending this trend into the realm of the physically and psychologically human to a degree heretofore not feasible—the ultimate division of labor, as it were. And each resultant increase in efficiency and productivity feeds directly into power—economic and military and cultural.

I don't say that's good. I don't say that it won't cause significant problems, as religions begin to understand that humans have taken their own self-creation into their own too-fallible hands, and we struggle with the prejudice and hate that humans always display, especially under pressure. But I do think human enhancement—transhumanism—is going to be hard to govern and harder to stop. You want the future? You can't handle the future. But it's coming anyway.

Over to you, Nick!

Best,  
Brad

---

**From:** Nicholas Agar  
**To:** Brad Allenby and Kyle Munkittrick  
**Subject:** Why I Don't Want To Be a Cyborg  
Posted Wednesday, Sept. 14, 2011, at 10:18 AM ET

---

Hi Kyle and Brad. Excellent discussion we've got going here!

I hope you guys don't mind if I cast myself in the role of Dr. Who, Captain Kirk, or John Connor—a selfless defender of humanity against threats alien, cyborg, and self-inflicted.

I believe that some of the technologies advocated by transhumanists genuinely threaten our humanity (or values connected with it). The threat is different in kind and more nebulous than that posed by [Daleks](#) or Terminators. But it's no less real.

Do humans want radical enhancement? Let me start by saying why I, as a card-carrying human, am personally opposed to it.

First, I should clarify that I'm not really about telling people they can't do stuff. Like any old-fashioned liberal, I'm a fan of John Stuart Mill's harm principle: People should be allowed to do what they want so long as they don't hurt others. If you want to cover your entire body in tiny tattoos of Elmo and doing so doesn't hurt anyone else (or infringe [copyright laws](#)), then go ahead. I see the philosopher's principal job as one of informing rather than finger-waving.

Some things that seem extremely attractive at the outset are less so upon closer inspection. I'm concerned that we're being wowed by enhancement technologies. When I read [Kyle's blog](#) and other explorations of radical enlargement (billion-fold enlargements of our intellects, millennial life expectancies), it's hard not to feel dazzled. You come away with a one-sided picture of human enhancement. When you're buying a house, you need more to go on than the glowing appraisal of the agent who's trying to sell it to you. My book [Humanity's End](#) exposes some of the hidden costs of radical enhancement. (More on those costs in a minute.) Of course, it's possible that what we gain when we trade in our Alzheimer's-prone neurons more than makes up for what we lose. But I actually doubt it.

I'm really enjoying [Brad and Daniel's book](#). (I haven't quite got to the end yet.) It seems to me a very useful engagement with technologies of human enhancement. Brad, you and Daniel are skeptical about philosophers and philosophical reflection and as a consequence want the debate about human enhancement to have a more practical focus. I'm a professional philosopher, so it's not surprising that I don't agree. I think the most interesting and pressing current questions about human enhancement are philosophical. They address human nature and values tied to human nature. These

questions can be frustrating partly because the values associated with our humanity are complex. And the notion of turning ourselves into nonhuman beings—transhumans, posthumans, various types of cyborgs—is new, something explored up till now only in the most speculative science fiction. Some people think we've been headed down this path for a long while. We're certainly a technological species. But this means that we're a species that *uses* technology—not a species that's destined (as [Ray Kurzweil](#) seems to think) to *become* technology. I don't see how we signed up for posthumanity the minute one of us picked up a jaw bone and bashed another one of us on the head, a la [Brad's 2001: A Space Odyssey example](#). You can renounce a technology that hasn't been integrated into your brain or body. If you decide that life's better without a smartphone, you can toss it. When we start technologically upgrading our brains and bodies, renunciation's much harder. Not all ways of applying technology to human brains and bodies are bad. But they are dangerous in a distinctive kind of way.

So how are they dangerous? Here's a hypothesis. Enhancement places human meaning under threat. We're a social species; we connect with other humans. We find meaning in what they do and experience. But we have a weaker connection with radically enhanced beings. Kasparov's chess play is much more interesting to us than the objectively superior stratagems of Deep Blue. This interest in connecting explains why it's worth wasting(?) money to put humans on the moon and trying to send them to Mars. We take a pleasure in human experiences that we couldn't in machines that would gather the same data and collect samples of rock more efficiently. By sending Armstrong and co. to the moon, we as a species get a veridical sense of what it's like to step on to an alien world. Try as you might, it's difficult to relate to the information-gathering processes of the Mars Rover.

Kyle diagnoses a tension in my thinking. My 2004 book [Liberal Eugenics](#) defends human enhancement. It presents an argument that prospective parents should be granted a limited prerogative to enhance their children's characteristics. Yet in 2010's *Humanity's End*, I'm opposing enhancement. There's actually no contradiction here. Human enhancement is like exposure to the sun and quality New Zealand pinot noir. It's good—but one can overdo it.

Kyle wonders whether I can give a principled answer to the question "How much is too much?" At what age does human life extension flip from being a good thing to being something that's bad, or at least less good? I have to confess that I can't answer these questions. But here's another question that's difficult in the same kind of way. (Apologies for all the analogies—they're philosophers' main resource when addressing issues that are new and difficult to understand.) What's the right amount of money to spend on melanoma prevention? The two easy answers are clearly wrong. It can't be zero dollars; melanoma is an extremely nasty cancer. It also can't be the entire health budget; that would leave no money to spend on anything else. The right answer is somewhere between nothing and everything—it's difficult to work out exactly where, but it's certainly worth trying. I feel the same way about the right amount of human enhancement.

Kyle, you think we may have plenty of time. Maybe that's true. But I agree with Brad about the quickening pace of technological change. Who's to say that radical human

enhancement won't arrive (much) earlier than scheduled? If it does, then let's be prepared. Brad, you seem to want to go the other way. You say it's already too late. I'm not so sure. While humans can't uninvent a technology, we can be selective in its uses. I can think of quite a few wars since 1945 in which U.S. presidents would have found it really rather convenient to drop a nuclear bomb or two. So I think that we can refrain from certain uses of enhancement technologies if we decide that doing so is important enough. I'm hoping for a future in which people will select very cautiously from what could be a pretty bewildering menu of possible human enhancements.

Kyle, you're clearly a human-enhancement enthusiast. I see a lot of confidence from transhumanists about the future that radical enhancement will make. Do you see any need for caution in respect of the development or the roll-out of enhancement technologies? Do you see a case for avoiding or limiting certain potentially dangerous avenues of research such as AI, for example?

Brad, you pick up on some really interesting things by grouping technologies together—seeing posthumanizing technologies as continuous with, and not fundamentally different from, earlier technologies. Do you think that your focus on general trends might occasionally overlook important moral differences?

Nick

---

**From:** Brad Allenby

**To:** Nicholas Agar and Kyle Munkittrick

**Subject:** It's Impossible To Predict Whether a Technology Will Be a "Good" or "Bad" Thing

Posted Wednesday, Sept. 14, 2011, at 12:36 PM ET

---

Wow. So Nick and Kyle have raised some great points ... makes me want to write another book (not!). There are a couple of issues here I think are fairly easy, a number that are much harder, and one or two that I want to suggest require reframing.

To take a simple one first, an obvious question is whether people actually want enhancing, and if so, of what kinds. Here, I think the first-off answer is pretty easy—give them the opportunity and see what they do! The problem with that approach, of course, is that it allows one to avoid worrying about network effects: If one person enhances, I'm normal (OK, maybe *I* am not, but let's be generic here); if 25 percent enhance, I'm still competitive; if 75 percent, I'm the new subnormal. Moreover, it sidesteps the "abortion effect": What is appropriate when a majority perhaps accepts an enhancement, but to a minority it represents an absolutely reprehensible choice? And, conversely, when is it OK for society to stop people from doing what they want as individuals? I think most of the folks who are seriously thinking about these issues acknowledge that glib generalities are inadequate, no matter how we tussle over specifics and fight our corners based on our individual perspectives.

But here we begin to get into some of the harder issues. For one thing, we need to be very clear that technologies that are potent enough to worry about are also potent enough to destabilize existing institutional, cultural, economic, and social systems. This can be either good or bad and often involves significant distributional effects, which are very hard to generalize. Facebook helped power the Arab Spring—is that a "good" or a "bad" thing? Isn't it too early to tell, at least in terms of outcome? Railroads destroyed a rural America and much of the extant ecology of the American Midwest—again, is that "good" or "bad"? UAVs—unmanned aerial vehicles—used in combat in Af-Pak—"good" or "bad"? When considering these questions, which I think are generally not answerable in any rigorous sense, it is important to differentiate between what is already here and what is speculative. In most cases, the immediate effects of a technology are known and are positive, at least for a substantial number of people—otherwise, of course, why does one bother to introduce it? (We call this Level 1 in *The Techno-Human Condition*.)

But technology systems of any power also will have social, institutional, economic, cultural, and other effects. These are very real, but they are seldom known *a priori*—and, in fact, can probably not be known *a priori*, because the systems involved are complex, adaptive, and rapidly shifting. Thus, for example, I doubt very much that Gottlieb Daimler looked at his new engine in 1885 and thought to himself, "Darn! If I go any further, global climate change will undoubtedly be a problem, and I really don't want to even think about the Middle East, oil, and the Israeli-Palestine conflict!" One must necessarily, then, treat the potential implications of human enhancements—and indeed all technology systems—as speculative, whether one is a utopian or dystopian.

This is an important point. These are all complex adaptive systems, and we really don't know what they'll do in the future and how they will play out. It is easy to spin utopian and dystopian hypotheticals, and it can be useful, to the extent we use them as aids to think about possible scenarios, including what institutions and society might do to respond. But it is a category mistake to attribute reality to a hypothetical future and use it as a basis for policy in the present. If we really could know how a human enhancement would play out, sure, a response now might be appropriate. But the difficult reality is that we don't have the faintest clue what the systemic institutional, cultural, and social implications of such interventions are liable to be. I think we need to get a lot better at perceiving and speculating about these systems; that's part of building our adaptive capability for an increasingly contingent, uncertain, and unstable future. But it is a step too far to try to argue for specific actions now based on hypothetical dangers or benefits that "will" accrue in the future—because we don't know how the future is going to play out. Thus, for example, some environmental groups argued strenuously a couple of years ago for a "ban on nanotechnology" based on a pretty dystopian set of hypotheticals that, in urging significant and immediate policy change, they wanted to be accepted as predictions of real futures. At the time, of course, virtually all electronics operated within the accepted nanotechnology scale (fewer than 100 nanometers), and new technologies—quantum computers, for example—were specifically being developed based on nanoscale physical structures and behaviors. (In recent months, an [anti-nanotechnology group or individual](#) has sent bombs to nanotechnology experts in Mexico.) Hypotheticals are very useful thinking tools, and scenarios are a powerful way to prepare for contingent futures, but in both cases it is critical to remember that they

are aids to more powerful thinking, not predictions about what is "really going to happen."

And Nick, you're right here: Each technology is not only going to be part of larger systems but also something that will raise its own little conundrums. A high-level discussion like this doesn't get into that depth, but every one—from, say, lethal autonomous robots to drugs that enable memory manipulation—has its own charms and snares. But in either case—the large tech systems, and the individual developments—I don't think we have nearly the knowledge, or understanding, or even perception of these systems that we think we do.

And there's a further caution, I think. We worry a lot about physical and cultural imperialism. But there is also such a creature as temporal imperialism. We like the world we live in, and its values, and our perspective on "right" and "wrong" and other moral issues. But suppose that someone early in the European Enlightenment, convinced that medicine was a dangerous anti-religious plot by atheistic scientists, had gained the power to stop medicine to ensure that we all died at 30 to 35, as God intended us to? (This example seems more relevant to me with every passing year.) Suppose that we froze the prevalent moral view throughout much of human history that there were "people" and there were "barbarians" and that treating the latter like the former was not only unnecessary, but uneconomic? If one is dealing with issues that live and die in several years, ethics and morality can usually be regarded as fixed (if not universal), but is it not a dangerous hubris to impose our ethics on the further future by trying to freeze and shape their potential before they even speak? This does not call for that straw man, relativism, but it does suggest that a certain humility that reflects our status as temporary representatives of a far longer collective might not be inappropriate.

But I'd like to return to a final observation. A lot of this dialog presupposes that 1) we have some sort of ability to control the evolutionary paths of these technologies, and 2) that we know enough about them to be able to do so rationally, ethically, and responsibly. Both of these questions require further research, but I think they are, at best, whistling in the dark. Do we really control where nanotechnology, or biotechnology, or information and communication technology, is going and what it will eventually give us the power to do? After all, many medical advances are "dual use" in the sense that, used for someone requiring medical attention, they cure; used for someone who is otherwise healthy, they enhance. Are we going to stop medical enhancements? Or regulate them? Our regulation of off-label cognitive enhancers is failing at every campus I'm familiar with, and not just with the students. And hasn't the EU has done its best to stop any agricultural genetically modified organisms from being planted anywhere, even if, like "golden rice" it potentially would avoid a lot of harm and loss of life? And failed? And didn't the previous American administration try to limit stem cell research, only to create state-based programs and drive leading researchers to other countries, where they continued their work? Historically, did the ban on gunpowder that the Japanese and, more informally, the Chinese implemented—and which they no doubt saw as ethically justified—work, or did eschewing this technology just open the way for British (and American) gunboats to sail up the Yangtze in the 1840s onward, and Adm. Perry's "black ships" to force open Japan in 1853? Put another way: If the United States and the European Union decide not to pursue a particular enhancement technology for ethical

reasons, why is there any reason in the world to expect other competing cultures to do so as well—especially when their religious or philosophic traditions, and their balancing of the risks and benefits of enhancement, may well be very different? I realize this poses a difficult conundrum: Do we really want to give other countries a veto over our moral decision? And yet, to ignore the fact that a decision against an enhancement technology may do nothing but weaken a society over time—because even if "we" stop, "they" won't—cannot be immaterial.

Nick is worried that I am assuming that it's already "too late." Well, yes and no. I think when that ape picked up the bone, and reconceptualized it as a weapon, and developed a culture in which bones and stones were used as weapons, it was in a meaningful sense already "too late," if you wanted to avoid human technological enhancement. On the other hand, "too late" is a normative judgment, and, frankly, I don't even get that far. I'm just trying to understand what's out there, and what it says to me is that a) rapid and accelerating technological evolution across the entire technological frontier is here, and it's already created psychological, social, and cultural changes we haven't begun to understand; b) because these systems are powerful, and grant personal and cultural authority, and have significant military and security implications, they are going to be hard to modify or stop; and c) even if we could modify them, they are sufficiently complex so at this point, with our existing institutions and worldviews, we are clueless as to whether we are doing something ethical and rational, or not. Tough world. But I'm not saying that it's good or bad. Just that it's already here, and it is, indeed, tough. I also don't say we can't adjust to do much better. But, again, I don't see that happening.

And finally, I do like philosophy—really! But I appreciate it more when it perceives what's actually out there, and speaks to real issues in ways that reflect the real world (yeah, I know ... obviously American vs. Continental). Too much of our current discussions—not this one, and not by you, Nick!—seem to revolve around worlds that might be, wistful dreams of sustainability and adroit defenses of past verities, or lapsing into speculation about angels, dancing with stars, and heads of pins. We do need philosophic framing—but we need it to be grounded in the world we have in front of us now, and to perceive that world, and to help us grapple with the complexity of that world, rather than to offer up erudite Kool-Aid for the cognoscenti.

Brad

---

**From:** Kyle Munkittrick

**To:** Nicholas Agar and Brad Allenby

**Subject:** How To Determine Whether a Proposed Enhancement Is Ethical

Posted Thursday, Sept. 15, 2011, at 10:09 AM ET

---

Nick, Brad, we're narrowing in on something vital about human enhancement. Transhumanism is necessarily speculative, but I think we can pin down why, precisely, this topic, which seems so science-fictional, matters so much for the rather near future.

We agree that we don't want to ban enhancement, nor do we want people to have unregulated, wanton license to recklessly modify themselves. So we are left with the question of how we balance the extremes and avoid the past horrors associated with projects that claimed to know how to make human beings better.

Nick, to answer your question, yes, I most certainly think we need a cautious approach. But, as you both point out, it is near impossible to accurately predict how most enhancement technologies will actually work and/or be used. I think we need an approach that presumes we have no idea what enhancement technologies will come into being. So we need a general philosophy, not a specific policy. Further, we need a philosophy that allows us to craft and adjust policy quickly in response to data, social shifts, and real-world events. Brad (you anti-philosopher, you), it is fitting you brought up Mill and Burke. I think both have something to contribute to our cautious, realistic approach to enhancement.

To determine if a given enhancement is ethical, we must focus on the individual. First ask, does this policy greatly harm individuals or small communities? Then ask, does this policy harm most individuals and groups? If both answers are no, we can move forward. What do I mean by harm? I tend toward Mill here, wherein a harm is a threat to life, limb, and liberty. But I also lean on Burke and know that no policy or technology exists in a vacuum; we must be realistic about what we allow. In short, eugenics and enhancement can and should be among the means we use to better society, but the primary end of policies around allowing or disallowing forms of enhancement should first and foremost be about individual well-being.

But does protecting the individual mean everyone must be enhanced equally? Of course not. Our second point of caution, around fairness and justice, is likely to be addressed by what we might call equality of access over time. Most technology starts expensive and exclusive but over time trends toward affordable and accessible. The rapid spread of cellular technology is a great example. I generally trust our current market economy to provide generally affordable access within a reasonable timeframe.

One might argue that those who have enhancements would act on the economy or legislature to prevent enhancements from being generally available, because lots of enhanced people means my enhancements won't be as valuable. Not so! Enhancement is not a zero-sum game. In fact, one could argue that without general intelligence enhancement humanity cannot survive the future.

Nick, you asked about my concerns around future threats like advanced AI. Let's add nanotech, global warming, nuclear exchange, and external threats like asteroids to the mix. There is a very good chance that our species is, as it stands, not smart enough to deal with these existential threats. If enhancements were zero-sum, we'd be doomed. But they're not. Enhancements, like most technologies, have network effects. Many enhancements, like intelligence and moral enhancement, get more beneficial as more people have them—for society, at least.

Moral enhancement includes modifications such as genetically increasing one's capacity for altruism or willpower; taking mood-lifting drugs that make you less likely to become upset or more likely to bond with those around you; increasing "moral intelligence"—that

is, the part of your brain that allows you to understand why something is right or wrong, with other general cognitive enhancers. I bring them up alongside network effects because some argue that if only a small number of people were morally enhanced, that few would be victimized because they would be too nice and easily coerced. Thus, they argue that only with network effects in play will moral enhancement be beneficial to society. I contest that being moral actually requires more of a person's fortitude, willpower, and mental dexterity as well as requiring a predisposition to more altruistic and beneficent behavior, meaning network effects aren't necessary for morality to be beneficial but that network effects make moral enhancement far more beneficial for everyone. So it is with intelligence enhancement: The benefits of being smarter than everyone by being part of a minority of the enhanced would be insignificant in comparison with the massive benefits each person would receive if the vast majority of humanity had enhanced intelligence. We, as a society, have a major incentive to ensure as many individuals as possible have access to intelligence enhancement, should they desire it.

So, when will we be able to actually use the ethical criteria centered on individual well-being to evaluate these new technologies coming down the pipeline? Brad, you implore Nick and me to look at the world around us—sound advice. Yes, technology overall is advancing at an incredible pace, but it is not as if we will go from Ritalin being a controlled substance today to being able to purchase off-the-shelf neuro-implants tomorrow. We have a rigorous FDA approval process for almost any new biotechnology. The process is not rapid. So, while the FDA assesses scientific safety and efficacy, philosophers, scientists, and (ugh) politicians will be able to assess the pros and cons of that technology and rationally determine which policies to put in place to protect individual well-being before the new enhancement reaches the general market.

Nick, given your aversion to transcending natural human maximums, I'm curious, would you consider it acceptable to make the average and the maximum human capacity for certain traits, like intelligence or longevity, one and the same?

Kyle

---

**From:** Nicholas Agar

**To:** Brad Allenby and Kyle Munkittrick

**Subject:** Why Can't We Enjoy Being Human Just a Little Longer?

Posted Thursday, Sept. 15, 2011, at 5:31 PM ET

---

Kyle, you've alerted me to your [excellent blogs](#). It's been good crossing (philosophical) swords with a smart transhumanist. Brad, you've convinced me that philosophers of enhancement tend toward an overly static view of technology. We tend to downplay network effects. [Your book](#) not only brings these out but offers a systematic treatment of them. It's a really distinctive contribution to the enhancement debate, though I wish there were fewer references to Martin Heidegger—Kool-Aid for the (Nazi) cognoscenti.

Perhaps I can make a few final comments on my hopes for our species' future. First, I'm not opposed to change. I accept that we are an evolving species. I can imagine that those who really desired a totally changeless humanity might actually take a keen interest in certain biotechnologies. Geneticists are busy surveying the human species as it stands in the early 21<sup>st</sup> century—they're providing a picture of shared and variable elements of the human genome. Suppose there were Guardians of the Human Genome (Steven Spielberg, are you reading?) empowered to preserve the human species as of 2011. They'd seek to purge new genetic variants and maintain existing variants at their 2011 frequencies. These imaginary Guardians wouldn't be concerned with any absurd superpowers granted by genetic mutations as per the [X-Men movies](#). They'd view preservation of our current human genetic identity as a valuable end in itself.

Needless to say, this bioconservatism is absurd. Humans will change. But it still makes sense to have preferences about the nature of this change. I think that abrupt, large changes—the kind offered by transhumanists—tend to sever meaningful intergenerational connections. It's one thing to view Granny as hopelessly out of touch with iPads and Xboxes. It's another to view her as an idiot, barely getting by on yesteryear's genetic enhancements. I have similar feelings about my immediate descendants. After going through their teenage stage of despising everything I am and represent, I hope that my kids look back on my life and find just a couple of the things that I did meaningful (wipes away a tear).

Humanity won't last forever. But I don't see why we can't enjoy being human and relating to humans for just a little while longer.

I don't think this involves denying people things that they (really) want. Brad, I wonder about your proposed experiment or giving people the opportunity to choose enhancement and seeing what they do. A few years ago, lending institutions ran an experiment to see whether people with limited financial means would choose adjustable-rate mortgages. Lots did. But were these mortgages really in their best interests? Would people have chosen them if they were better informed? We now know that the answer to these questions for many borrowers was no.

People need information about values on both sides of the enhancement debate. It's relatively easy to see the appeal of life extension and cognitive enhancement. Their benefits are easily recognized and quantified. If Wal-Mart offers to set up in your town's historic district, it has the advantage of offering things whose value we can easily acknowledge and quantify—cheap consumer products and employment. Opponents have a more difficult time saying exactly what would be lost (it's not as if any of the buildings to be demolished appear on [UNESCO's list of World Heritage Sites](#)). Values that are difficult to express tend to get elbowed out by values that are easier to state and quantify. I think we need to put in more work expressing them.

Kyle, you think that enhancement of general intelligence may be required for humanity's survival. Such enhancement might protect us from the fate of the dinosaurs—consigned to extinction by a threat that we couldn't even begin to comprehend. Just a speculation—but I wonder whether our response (or lack thereof) to extinction threats might not be a matter of being more intelligent. I was impressed by an argument in David Deutsch's recent book [The Beginning of Infinity](#). According to Deutsch, there are no threats that

we, as children of the Enlightenment, are not capable of responding to. We differ from the dinosaurs in having science. Humans are essentially problem-solving creatures. And really—if Bruce Willis is clever enough to deflect a potentially humanity-destroying asteroid in [Armageddon](#), then why shouldn't the rest of us be?

My opposition to certain enhancement technologies is no Luddite rejection of technology. It may seem odd to ally myself with the Trekkies, but I don't see what's wrong with a vision of the future in which beings like Capts. Kirk and Picard—fully recognizable human beings—get to travel to the stars. Surely we all don't have to become Spocks or Datas to qualify for that privilege.

Nick

---

**From:** Brad Allenby

**To:** Nicholas Agar and Kyle Munkittrick

**Subject:** There Are Some Places Our "Wetware" Bodies Just Can't Go

Posted Friday, Sept. 16, 2011, at 7:18 AM ET

---

So many interesting directions here! Before we sign off, let me offer a few comments on Nick's thoughts. First, while I like the Capt. Kirk model as well as anyone, I have to think there's got to be a point when our attachment to our physical selves simply becomes too inefficient. As biological organisms—"wetware"—we're just bad fits with certain environments, such as combat or space. As you know, we've been working on computer-brain interfaces, in the form of chips implanted in brains, that enable smooth coupling between an individual and their (properly instrumented) environment. Experiments that began here at Arizona State University and have been continued at Duke and elsewhere have involved [monkeys learning to move mechanical arms](#) to which they are wirelessly connected as if they were part of themselves, using them effectively even when the arms (but not the monkey) are shifted up to MIT and elsewhere. More recently, monkeys with chips implanted in their brains at Duke University have kept a [robot wirelessly connected to their chip running in Japan](#). Similar technologies are being explored to enable paraplegics and other injured people to interact with their environments and to communicate effectively, as well. The upshot is that "the body" is becoming more than just a spatial presence; rather, it becomes a designed extended cognitive network. It is conceptually a simple step to build body extensions that are wired directly into our brains, but, because they are not wetware, but hardened mechanical systems, can easily and far more efficiently go where no body part has gone before: space.

This is a bemusing prospect because it raises a number of interesting questions. How does it affect our brain, for example, if the time lag between sensory perception and core brain reception of the information is not fractions of a second, as with our current bodies, but, because our "body" is, in fact, on Mars, a lag of between three and 20-plus minutes? But perhaps more fundamental, what happens to our definition of "the human" when it begins to include remote functionality that may or may not be biological, but

might increasingly be mixed (as in the existing robot with guidance provided by rat brain tissue, rather than silicon)? Is humanity on Mars when my robot extension is? What about when cognition is mixed, partially onboard the Martian Brad, and partially in my wetware brain here in Phoenix? And, something that perhaps offends our sense of specialness, what happens when it becomes clear that our robotic selves, and not our biological Cartesian selves, are the real inheritors of Space, the Final Frontier?

I also appreciate your question about Heidegger (who was, for those of you unfortunate enough to have escaped philosophy in university, deeply involved in National Socialism, although the extent to which he was active as opposed to a fellow traveler is somewhat unclear). It is always a problem with brilliant, but seriously flawed, individuals as to whether one should appropriate the thinking and the contributions, and ignore the imperfections in the vessel from whence it comes. In general, I favor the appropriation model: Let people at least contribute some good if they can. Wonderful discussion. I am going off to practice remote time-lagged cognition.

Brad

*Nicholas Agar is an associate professor at Victoria University of Wellington in New Zealand. He is the author, among other things, of [Humanity's End: Why We Should Reject Radical Enhancement](#) (2010) and [Liberal Eugenics: In Defense of Human Enhancement](#) (2004).*

*Brad Allenby is the Lincoln professor of engineering and ethics; a professor of civil, environmental, and sustainable engineering; and the founding director of the Center for Earth Systems Engineering and Management at Arizona State University. He is co-author with Daniel Sarewitz of [The Techno-Human Condition](#).*

*Kyle Munkittrick is a bioethicist and a program director at the Institute for Ethics and Emerging Technology. He blogs at [Pop Bioethics](#) and Discover magazine's [Science Not Fiction](#). Follow him on [Twitter](#).*

**Article URL:** <http://www.slate.com/id/2303277/>