G.O.D. IS GREAT

By the same author
LIFEBALL
MINDWORLDS
GODBLOGS
www.andyross.net

G.O.D. IS GREAT

How To Build A Global Organism

Andy Ross

A Ross International publication
Corrected printing August 2010

ISBN 978-1-452-89792-9

A CIP catalogue record for this book is available from the
British Library and U.S. Library of Congress

The Message

You deserve at least a tee-shirt summary before you commit to reading a book by an author who may be new to you on the unusual subject of how we're building a global organism.

My aim is to help you develop a lively understanding of what's happening to us all this century. The world is changing fast, and I think most of us could use a new and improved concept of where we're going. My way works fine for me. As an evangelist from the computer world with a background in science and philosophy, I think my insight is worth sharing.

Through our efforts to develop new technology and globalize its industrial application, we humans are working together to create something new on this planet. We're in the process of creating a global organization so integrated that we become parts of a single living organism. I call this organism *Globorg*.

I see Globorg as the natural culmination of biological evolution on planet Earth. It will embrace humans as living parts. Its history stretches back over three billion years. Globorg will wake up soon after we've finished building its brain. We're doing that now, with the infrastructure of the global web and its server farms hosting cloud services.

We shall achieve a unified consciousness with a global dominion. We shall identify with Globorg. Or rather, that's what we'll do if we know what's good for us. We'll orchestrate our activities worldwide so far that on a clear day we can see and act as one. Globorg will have woken up. But we still have a long way to go. First, we need to win the war between science and religion. I propose a logical foundation for a new philosophy of life.

This book is a road map from here and now to Globorg. It's a guide for beginners. It offers not just information or even knowledge but new ideas and a dash of wisdom. Enjoy.

The Book

This book can tell you something useful. I believe it's something we all need to think about. But you won't just get it in a minute. It takes time to sink in. You need to play with it awhile and relate it to your own experience. That's why I wrote it as a book.

But books are an old technology. This presents problems. In a world of podcasts and five-minute videos, who has time to read them? In a world of glamorous graphics on every website, who wants to feast their eyes on monotone pages of text?

As a blogger and a surfer, I sympathize. By any standards I'm a book lover, but I have to admit there are better ways to get information and better ways to delight the sensorium. And here I am offering you a book. Why didn't I just blog it all?

A book is a milestone in history. It has a time and a place and it makes a statement. It puts its message right inside your head, where the author's thoughts become your thoughts. It's a good medium for a deep message that you need to think about. The meaning sinks in and takes root. Or at least it does if the author has done his or her job right. This takes some skill, but when it works you know it.

Videos, podcasts, and blogs are more ephemeral. Blogs are best for short takes, where you can get the point in a minute and move on. Podcasts suffer from their fixed timing. Boring stretches go by as slowly as the good bits, where you want to pause and hear it again. As for videos, who can avoid getting distracted by the haircut or the horror-movie lighting? No, books work better for deep thought.

Deep or not, this book is hard to classify. Is it technology, current affairs, philosophy, prophecy? If you've read my blog, you'll recognize some of the themes and topics.

The book is a manifesto for the next revolution in life on Earth.

The Author

Authors should be authoritative. So who am I to be telling you what I'm about to tell you?

I'm not famous yet, so Googling me will probably steer you to my blog, which is filled with technical details and trivia, so may be too much at once. In a sentence, I'm a philosopher who decided not to become an academic philosopher but to carve a path in science and technology. My journey through physics teaching, academic book publishing, and making business software reminds me in retrospect of Charles Darwin's voyage on HMS Beagle. It gave me a wealth of facts to weave a worldview of some novelty.

A long time ago, I collected four degrees in philosophy, three from Oxford and one from London, but I don't imagine they're tradable currency now. I was firmly in the Anglo-American analytic tradition, which is based on a German mathematical approach and claims to have inherited the genius of philosophy from classical times. But mathematics and physics had impressed me more than philosophy, so I went back to teaching them. I soon saw that mathematical logic had spawned the whole digital information revolution of our time.

In Germany, first editing computer science books and then assisting at the birth of new software, I reflected on philosophy, science, religion, and the onward rush of globalization. On the way, I wrote a science-fiction novel and published a volume of essays on the science of consciousness. At the end of 2009, aged 60, I resigned from my job to write this book.

Actions speak louder. Read on.

Germany, 2010

The Thanks

To take the top twenty, I'm grateful to:
– Martin Amis for coaching my authorial voice
– Tim Berners-Lee for the really useful web ideas
– David Chalmers for the hard problem and the parties
– Frances Cairncross for the lesson in economics
– Richard Dawkins for keeping me sharp on evolution
– Daniel Dennett for lots of good philosophical ideas
– Bill Gates for some visionary ideas I could use
– Sam Harris for encouraging my interest in religion
– Andrew Harvey for showing me how religion works
– Christopher Hitchens for robust style and a title
– Ted Honderich for his personal kind of philosophy
– Otto Kuehn for years of lunchtime conversations
– Ray Kurzweil for showing me how to be a prophet
– Angela Lahee for lots of great conversations
– Judy Mallaber for the early introduction to politics
– Thomas Metzinger for liking my avatar idea
– Hasso Plattner for sparking my spirit of enterprise
– Helen Ross for being a brilliantly supportive sister
– Stephen White for many fruitful conversations
– Hans Woessner for training my perfectionism

The Contents

A revolution is not a dinner party.
Mao Zedong

A Muslim has no nationality except his belief.
Sayyid Qutb

I'm king of the world!
James Cameron

00
Opportunities

In an evolutionary journey across the generations, every step from parent to child must pay off. You can't climb up Mount Fitness unless every step you take is secure. This is a hard constraint on possible histories of evolutionary ascent.

So when it comes to the utopian future, the hard part is getting from here to there. How do we, now, with a small planet, diminishing resources, economic problems, political confrontations, human animals of breathtaking irrationality, and new hazards that seem more horrific every year, move in the right direction? Do we even know which direction is right?

Utopia is a hard sell. Any buyer needs a payoff now. I don't want to sell the global organism as a utopia, because it won't be, but it must work for someone or we'll never get there. And again, it must work all the way. So my first challenge is to show how opportunities around us, here and now, will lead us in the proposed direction.

In the first quarter of this manifesto, in the next four chapters, I introduce four big opportunities that people with money to invest can exploit to make more money. Once we've all harvested the fruits of those opportunities, we'll move on to take more chances and make more improvements until the global organism emerges whether we like it or not. There's a delicate balance of free choice and inevitability about all this that I find fascinating.

At first sight, the opportunities are technical optimizations that arise as we apply and refine new technology. The business challenges are secondary. The business environment needn't evolve much at first to accommodate the new ideas. But exploiting the opportunities will change the global business environment, and these changes are the focus in the second quartet of my tract. The third quartet will then take on a big new wave of transformations, and the final quartet will bring it all together in Globorg.

First, we need to climb the foothills of the fitness landscape.

0000
Avatars

Twenty years ago I invented an electronic globe. I called it the Globall Hyperatlas and decided not to patent it. For the record, I defined the concept in a 1991 issue of a journal called *The Visual Computer*. Some years later I remembered the article and mailed copies to Bill Gates and Nathan Myhrvold at Microsoft. In 2008, Microsoft announced they had developed an electronic globe they called the Sphere. I posted a photo of Microsoft developer Andy Wilson showing off a prototype on my blog. But by then Google Earth had arrived and the whole idea of a globe as user interface was gone.

I still think building an electronic globe is an interesting technical challenge. You take a glass sphere with a hole in the base and cover it inside with millions of pixels. You program it to display rotating images of the Earth, generated either from stored files or from data collected via live feeds from imaging satellites, and to show days and nights, seasons, and weather and climate changes. You could also make it show continental drift, population and resource data, and a lot more besides. You could steer the imagery with a handheld device and have a lot of fun.

But you'd soon want more. Zoom is already a problem, since the curvature of the sphere sets the scale, so you'd need a second monitor to show fine detail. But then you could forget the sphere and just zoom in on a normal screen image, as in Google Earth.

What this little story illustrates, for me at least, is the headlong nature of the slide or the rush to virtualization. The display device, be it spherical or flat, is just a barrier between the virtual image and my consciousness. I want it gone, transparent, and that means I don't care whether it's round or flat. If it limits my experience in any way, for example by setting a scale with its radius of curvature, it's already suboptimal. I just want to shoot the desired bit-stream to visual cortex. In practice, this means that maximally generic hardware, like flat screens, supplemented by maximally functional software, for doing scaling and rendering and so on, is the way to go. Virtualize as much as possible and reduce shipping costs.

But the rush to virtualization takes us a lot further than this. Why ship a big screen when you can get by with a pair of eyeglasses fitted with a HUD projector? Head-up displays in aircraft are a mature technology and they can be miniaturized via nanotech as well as any other devices. At present they've just become small enough for use in cars, but soon they'll shrink to eyeglass size and we'll all go around with smart shades that project 3D images of anything we like into our field of view.

More yet, why not cut the eyeball out of the loop and just shoot the bits straight to the visual cortex at the back of the brain via neural implants made of carbon nanotubes? The technology for this step is still schematic, and because the retina already performs some pre-processing of incoming images it may not be easy, but there's something there. The whole point of the process chain, from the input imagery coded in bits to the electrochemical buzz at the back of the brain, is to go from end to end as smoothly and directly as possible, in order to enjoy a minimally fudged experience of whatever it is that the imagery represents.

This is still not the end of the virtualization story, because there's still a real you getting shot with all this stuff and deciding what to do with it. This is of course suboptimal. Who wants to be forced to make decisions in real time, again and again, endlessly, for excessively trivial streams of input that your own personal online virtual agent – your *avatar* – should be able to handle for you? To add some metaphysical woo-woo here, we can call your avatar your virtual self. It, or he, or she, still acts in your name and keeps you in the loop just as long as you want to be in it. But you can delegate the boring stuff. Or rather your avatar can delegate the boring stuff, to a series of further virtual agents that can be summoned like genies from a bottle by just starting the relevant threads in the underlying software.

Your avatar is your virtual self

Like me, you may find this move toward virtualization hard to get a grip on. As an innocent consumer of new technology, you just wanted to interact with an app, say, to check the weather in Kansas, make a hotel reservation in Athens, or have fun observing nature on a tropical

beach, whatever, and here we are talking about avatars. So let's be a little more cautious as we set up this future vision.

People go online to do things, not just to admire the imagery. When things get done, there's a doer and a there's a thing done. We're all used to the idea that the thing done can be represented in software, or virtualized in some fashion, so as to make your task easier. Instead of scratching inky squiggles on paper, you let your fingertips dance over the keyboard and see little images of letters dancing to your tune on the screen. When you've completed your virtual document, you click on the printer icon and get a piece of paper bearing tidy testament to your thoughts. If, like me, your trade is words, you soon learn to cut the paper, save the trees, and keep your thoughts in bits for as long as possible. You do much the same – you still write words – but they've been virtualized. Soon we'll all get hooked on the descendants of Amazon Kindle readers and Apple iPads, and paper-based books and magazines will go the way of horse-drawn carriages. As reading and writing get more virtualized, we step closer to virtualizing ourselves.

Returning to avatars, the point is that you, the flesh-and-blood creature with wishes that may be fulfilled online, have an online presence. Your presence can be tricked out any way you like, as a Na'vi perhaps, as in James Cameron's wonderful movie, and you can inhabit this avatar with as much conviction, with as much feeling and relish, as you like. Others can experience your avatar as you. The avatar becomes a proxy for the real you. In just the same way, for example, I can experience the e-book rendering of *War and Peace* as Leo Tolstoy's immortal masterpiece, and interact with the proxy well enough, by reading it, to feel I've enjoyed the masterpiece itself. Similarly, with a Globall, I can experience the moving image of the Earth as my home planet in all its physical beauty, though here the proxy allows less in the way of interaction with the original. But in all such cases, the image stands in for the real thing. This is how avatars represent us.

As a philosopher who enjoys going to metaphysical extremes, I want to push the idea further. We live in our interaction with each other, and that interaction is still mostly in the flesh. But as online living takes up more of our identity, especially as it relates to the worlds of business and government, so the avatars we construct online encompass more of what makes each of us special. My online presence is defined by a lot of facts that are coded as bits. Anyone wishing to

steal my online identity would only need to copy a sufficiently descriptive subset of those bits to reconstruct a plausible avatar of me that could do bad things on my behalf. So we know the bits go a long way toward defining us. Now the philosopher in me says that our innermost identity as people is largely made up from layers of such imagery inside our heads, as our own proxies for ourselves, so in the end we construct our own avatars.

This topic is well aired by an eminent philosopher called Daniel C. Dennett. Since Dennett will reappear in our story, let's take the time to introduce him. With his domed baldpate and long silver beard, he'll surely be the model for a superbly iconic avatar one day, but once he was a young man. After studying in the Sixties under the logician Willard Van Orman Quine at Harvard and under Gilbert Ryle at Oxford, Dennett came up with some unsettling brain-in-a-vat thought experiments that got philosophers and cognitive scientists alike thinking more carefully about how our mental lives play out in our brains and bodies. Based at Tufts University near Boston, he's now one of the world's greatest philosophers, as well as a working farm owner and an accomplished sailor. Not only that, he's as friendly and approachable as your ideal Santa Claus.

Dennett related his brain ideas to selves in his big 1991 book *Consciousness Explained*. In short, he says the self is like a fictional character that we create for ourselves by telling and retelling the tale of our own lives to ourselves. This tale goes through multiple drafts as we spin out the story in our cerebral word processors, in what he calls the Joycean virtual machine that runs in our brains. The reference is to James Joyce, who in his amazing novel *Ulysses* introduced Molly Bloom through a stream of consciousness represented as an unbroken string of words flowing through the virtual machine software running in her brain. Virtual machines were unknown when Joyce wrote his epic, which he published in 1922, but the idea that we define ourselves with our words is as old as the poetic arts. My new step is to generalize from words to anything that can be coded as bits and say we define ourselves as avatars.

As in the movie *Avatar*, the question of which is the primary reality can be important. We're used to the idea that our physical body, the one that eats and sleeps and has sex, is the real and primary one. All the rest is made of bits that evaporate when you turn off the power.

But many physicists are happy to describe the universe as a colossal quantum field defined by a few septillion googols of bits, which form our world like a vast simulation, much bigger than in the Wachowski brothers' movie *The Matrix*. So our physical reality as human apes is like an avatar reality. To use a handy religious metaphor, we're all avatars in a game played by God.

The past is fixed but the future is not

All this technomania can easily diverge too far, so let's tighten up the story. The virtualization of more and more of the interactions that are important to us in our daily lives is an ongoing process with no end in sight. The Globall saga was just an example. Not only the things we do but also the things we are as we do these things admit virtualization. Behind the pixelated faces of our avatars, we shall go online to do our business, at work and play, and this code-bitten life will take up more and more of our time and energy. So much is fact and easy extrapolation. The big question I want to explore in this book is this: Where are we going?

As our future unfolds, new opportunities emerge and we make new choices. Soon the iteration of steps gets too complicated to track and we have to fall back on possibilities and contingencies that confuse everyone. But I want to keep it simple. If things get too complicated, I shall just back up and stay with the big picture, where the view is clearer. There is still a tale to tell, as we shall see, and a clear enough message to keep us busy for years working out what to do before the deeper future I glimpse begins to crystallize around us.

My metaphysical picture here is that the past is fixed but the future is not. As we move on in time, ranges of possible futures in a still-fluid state are crystallized or compacted down, in part by our deliberate actions, to form new strata in the layered bedrock of the past. How much of what goes down we can decide freely, in our attempts to choose or at least fine-tune our fate, and how much happens willy-nilly, independently of our conscious will, is a question for metaphysicians to debate among themselves. Many philosophers, Dan Dennett among them, have written good books on the topic, so I don't need to get side-tracked discussing it here, fascinating as it is.

To be more concrete, I shall restrict my attention mainly to what seems likely to occur this century, up to the year 2100 CE, give or take a decade or two. And I shall look mostly at practical issues impacting directly on humans, such as business and economics, technology and applied science, demographics and the politics of identity, and climate change and the environment. That looks like a big enough agenda for one time-pressed author, even if it is only a small step in the ongoing public debate about all this. In fact it looks like such a big agenda that you may be worried. Is it not too much for a mere philosopher, however ambitious he is? Here I aim to reassure you, if you have the faith to read on.

Agent + Angel = Avatar

So let's get started. The concept of avatars strikes me as a good one to launch the proceedings. As we organize our societies increasingly around online services, citizens who wish to keep on good terms with the burocrats and take due advantage of business opportunities will need to organize their own online presence. Reading and writing emails and wheeling and dealing via online forms on the web will soon take more time than it's worth. Just as maintaining a homepage, posting photos and videos, socializing on chat sites, and tweeting can easily swallow up what remained of a nerd's social life, so running a business life online can all too easily drown out the vital tasks of making the rounds and pressing the flesh. We need help.

As humans, we have a physical life and a virtual life. We cut down on the physics as the online life takes over. Speaking for myself as a typical victim of this lifestyle, I could use more and better software to help out whenever I get fed up with online chores, and I'd like more tools to represent me whenever someone wants to deal with me online. We can unite the role of a software agent that performs delegated tasks for me with the role of a screen angel that assists and advises me and call the result an avatar.

My avatar represents me online and helps me manage my online work. It's my shadow self, the bit-image of me. It's my humble servant and my proud ambassador. Perhaps it's even a first prototype of my informatic soul.

Late in the last century, as an amusing exercise to mortify some Mormon missionaries, I developed a theory of the soul as the bitmap of a person. Imagine how the *Star Trek* teleporter for beaming Kirk and Spock down to a planetary surface might work. The machine reads the precise quantum state of their bodies as a discrete set of bits (or rather qubits, but let's not fuss), destroying the bodies in the process, and then runs the bits through a sort of 3D printer to print clones of Kirk and Spock at the other end. The bit-stream is the informatic soul. That was my theory. In the cyberspace (the omega-dimensional unit hypercube, to give it a mathematical dress, where lowercase omega stands for a countable infinity) of all possible bit strings (qubit strings), our souls are represented by a vanishingly sparse dust of points, each with a finite but rather long address vector.

Given such a fantasy, an avatar in my new sense is analogous to an informatic soul. Naturally, it's vast orders of magnitude simpler, since in the current state of the art it's hardly more complicated than the bit string for a Facebook page, but the idea is pregnant with the deeper potential of the *Star Trek* teleporter.

In the gaming world, my avatar is just the screen hero with whom I identify and through whom I act. In the movie *Avatar* the hero's avatar is, logically enough, the Na'vi body he operates remotely via brain-waves, rather like a radio-controlled lab rat with a wi-fi chip in its brain. In both senses of the term, my avatar is my incarnation in another world, and in both cases the real me is a bag of flesh that gets temporarily ignored as "I" go out to infect my avatar with the gift of my life. This is the sense I want to zoom in on, for the reason that I think avatars thus conceived will be big in the coming decades.

Avatars will be big

Avatars will be big because we shall all need them. We shall soon be struggling to keep a handle on our online presence, and our avatars will serve as the lifesavers that represent us whenever we need the help or want to take time out. The feature that distinguishes avatars from agents, angels, ambassadors, bots, clones, and drudges generally is that my avatar, when active and on duty, incarnates me.

One big thing here is the business opportunity. Imagine you're the Bill Gates of your generation, looking for a new software challenge big

enough to use you up, a killer app so bursting with must-have promise that your initial capitalization can go to the billions. Robots are interesting, but that's a hardware challenge. Genetic engineering is interesting too, but that's for lab freaks. No, it's avatars. Like the first Microsoft products, avatar software addresses a new need, in this case to outfit and equip my online presence to appear and act in just the way I want and need.

So this seems to me to be a billion-dollar idea. It's not really new, of course, but the name is. The name hit me when I saw the movie. Gamers have given the word a resonance that just happens to work for the new app. For me at least, the concepts behind agents and angels were promising but the words didn't quite hit the spot. Technically, whatever we call them, they're just finite automata or finite state machines or something similar, and the theory behind them was already clear to Alan Turing and other computing pioneers many decades ago.

To update the background ideas, if you want to dig deep here and struggle with a big book, my tip for ambitious beginners is to have a go at Stephen Wolfram's *A New Kind of Science*. Wolfram is the software genius who coded just about the whole of mathematics into his software package *Mathematica*, which made him enough millions to take years off to write his big book. Based on Turing's ideas, Wolfram's work on cellular automata suggests that we can model the universe, people, and complex processes generally as such automata. So let me define avatars as automata that have a me-shaped hole in them, such that when I fill the hole they can do my will online.

Avatars will leverage the semantic web

Tim Berners-Lee invented the World Wide Web. Fresh from Oxford as a young physicist, he adapted the then-new Internet protocols, created the hypertext markup language HTML, wrote the world's first browser, and put them all into service at the giant CERN particle accelerator in Switzerland to handle the data traffic from their particle experiments. (The accelerator was called the Large Electron–Positron collider and was later rebuilt as the Large Hadron Collider.) To publish the data, Tim also put the first-ever website online in 1991, on an Apple NeXT computer at CERN. After founding the World Wide Web

Consortium (W3C) and encouraging the viral spread of his brainchild, he sees the semantic web as the next phase in the evolution of the web, as he explained to me at the annual WWW conference in New York in 2004.

Technically, the semantic web will be a lot more complicated than the present web. Any given node *N* will refer to other nodes that define the collections of entities invoked to build up the functionality of node *N* (in a nod to philosophers, the collections are called ontologies). So when you want to do business at a semantic website, your app will need to access all the reference sites to understand the semantics of all the stuff you need to use to do the business.

The big benefit is that the new web will support understanding. Because understanding is reduced to machine mappings, your avatar can do it better than you can. For example, say you want to spend a week in Athens. You avatar can check flight schedules and hotel listings and make plans on your behalf, and when all is good and ready offer you a screen with a few options to browse, so all your tired human body needs to do is click a box. The avatar then reaches out online to seal the deal. Avatars will leverage the semantic web.

If you're as habitually skeptical as I am of such idealized scenarios, you'll now be asking what the difference is between this avatar vision and a setup with quite ordinary software agents and wizards to do the hard work. The difference is this. I shall identify with my avatar but there's no way I can identify with the agents and wizards that clutter today's world. Maybe an agent or wizard does what I want and maybe it doesn't. Maybe it just annoys me, like that notorious Microsoft animated paper clip that popped up to ask me if I was writing a letter whenever I typed the word "dear". For such an agent, I can sit back and judge both the execution and the result and feel no more emotional reaction than a flicker of pleasure or pain.

By contrast, the promise in the concept of an avatar is that it unites agency and wizardry with the semantic smarts that come with instant access to the world's ontologies. On top of this – and this is the key feature – it comes with a me-shaped hole for my personal data. It's the online incarnation of my face, my fingers, my thoughts, my data trail, and in the fullness of time my plans and hopes and dreams. One day my avatar may even gently embrace my informatic soul, so that I may sit back from the keyboard and rest in peace.

If the business opportunity here is big enough for another Microsoft or Google to emerge, then a large part of our future is defined. In the business reality we now inhabit, if sufficiently many people worldwide can be made eager to extend their online presence, we can be sure that a new generation of entrepreneurs will join the chase and build up companies to develop the new avatar apps. As with the desktop PC and web search industries, a huge new ecosystem of companies offering subsidiary goods and services will arise to round out the revolution.

There's plenty of room at the bottom

The information revolution was a tectonic shift that caused three to five big waves of innovation, depending on how you count.

First, IBM and other pioneers gave us mainframes. These were centralized computers – "big iron" – that users could timeshare to write and run their batch jobs. Multinationals and other well-heeled corporations used mainframes to crunch their core data, and this allowed the companies to grow and consolidate their presence in ways that were previously impossible.

Then came personal computers, thanks to printed circuits and in particular the microprocessor. The earliest machines were mainly for hobbyists but progress was fast. The 16-bit Intel 8086 of 1978, for example, was a milestone. Bill Gates and Paul Allen implemented the high-level programming language BASIC – "beginner's all-purpose symbolic instruction code" – on it, which enabled the kids at Microsoft to develop killer apps like spreadsheets and word processors for the sweet new machines, such as the IBM PC of 1990. Soon the whole industry took off and there were millions of PCs and countless new apps and games to run on them.

A secondary wave here was the emergence of client-server architectures for business computing. What had by then become the "tyranny" of the rigid mainframe was replaced by a new freedom. In 1972, a small group of former IBM employees including Hasso Plattner founded SAP, which developed its early products into the 1992 R/3 system. This three-tier system featured "real-time" client-server architecture with users and their desktop clients on top, apps running on servers in the middle, and a big database at the bottom. R/3 bundled a suite of productivity apps for enterprise resource planning

(ERP) and SAP became the king of ERP. To keep you in the picture, I should disclose that I spent almost eleven happy years at SAP, from 1999 to 2009.

Next came the online revolution. The Internet and on it the World Wide Web allowed global networking and encouraged the vision of all the computers in the world uniting in the "cloud" to form one humungous hypercomputer. Stanford students Larry Page and Sergey Brin founded Google in 1998, and the Google search service, powered by thousands of servers in huge distributed networks, became the first killer app in the cloud.

The latest follow-on wave in the new century brought us mobile devices such as notebooks, handheld organizers, and cell phones connected in wireless networks at what we call the edge of the cloud. In a biologically inspired vision, all the nodes, from the big root server nodes to handy little edge or leaf nodes, are rather like neurons in a global informatic neuronet.

The five waves are not the end of the story. Even if we suffer a climate apocalypse that forces us to divert our attention to relocating all our coastal cities, the waves will continue for a while yet. The inexorable laws of human inventiveness and desire to get rich on the push side and of our need for productive employment and efficient business processes on the pull side will endure. The ongoing crystallization of our lives around increasingly hi-tech information exchange will continue to press hard against the frontiers of tradition, ignorance, and vested interest. And the technology is still changing, evolving, developing, improving, always in the direction of allowing more data to be processed more quickly and more efficiently in terms of energy usage, on smaller and more robust devices, which will get cheaper, to ensure that the new platforms and apps become more pervasive and more central to our lives.

The physics prodigy Richard P. Feynman said at a year-end speech in 1959 that there's plenty of room at the bottom. He meant that machines could continue to get smaller and faster and more efficient for quite a while yet. In a decade or two, we shall hit the quantum limits, and then we shall need to do things differently to keep up the momentum. But since quantum computing is already a big and active research field, we can be fairly sure the revolution will just roll on as we enter the age of picotechnology.

Imagine a model home with a terascale box in the basement

Before we drift off into the hazy deeps of the future, we need to pin down the avatar idea. Essentially, we're looking at a big app that consumes more hardware than most people currently have installed at home. To be sure we're on the right page here, we can guess the sort of retooling we shall need to deploy avatars pervasively. Very crudely, we shall feel the need for gigabit broadband to all homes (to download movies in seconds) and into home servers with teraflops multicore processor arrays (to render movie-style animations) and terabyte-range installed memories (to run fat apps fast). In other words, we shall begin to migrate up a few orders of magnitude from the gigascale of current home setups to the terascale. Each order of magnitude will take ten years or so, depending on how distracted we all are by the clash of civilizations and global warming.

We shall build new houses with networks that put all the heavy data-crunching load of the household onto one home box, and we shall interact with the box via handheld devices or wall screens around the house. Done right, this could work out cheaper for householders than living in unboxed homes where random assortments of impulse purchases are jury-rigged together in the ad hoc constellations that many of us rely on now.

Now we need to be a tad more systematic. Imagine a model home from the future twenty or so years out. It has a big box in the basement that routes movies, music, phone calls, and Internet traffic of all kinds through a home server engine sitting in a terabyte of installed memory to feed screens and devices around the property. (To keep you fully informed here, my SAP team in 2009 was busy building a scalable multicore database engine that would run productivity apps entirely in terascale memories when I left to write this book.) The box is a standard commodity, an off-the-shelf utility.

The box pays for itself by managing energy usage. It negotiates cheap deals with energy providers and schedules home power usage to optimize pricing, for example by heating water overnight. It meters and sells power from rooftop solar collectors, again negotiating with local and remote customers to get the best deal. And it charges the household's electric car, perhaps exploiting the car's battery capacity as a buffer to store energy for home use whenever the pricing works

out well. In line with most present expectations, this is a future where a lot of energy comes from renewable sources such as solar, wind, and hydro rigs, and the pricing is dynamic, depending on the weather. The household is an active node in a smart grid along the lines that IBM and other companies are currently developing on a strategic scale. Over the financing lifetime for the house, which we can set arbitrarily at twenty years, a box costing a few thousand dollars could pay for itself several times over in such a smart grid.

That was the plumbing. But the householder need never lay hands on the box in the basement. He or she tells it what to do through a wall screen in the kitchen or living room. The box is smart enough to react to commands spoken loud and clear in an authenticated voice and to reply through a photorealistic majordomo avatar on the screen. Voice recognition and speech processing are heavy apps in terms of flops of data processing, but with online help where necessary (for example when parsing the bilingual idiolect of the au pair) the box will certainly have the raw power to handle the task.

The industry perspective here is that software developers regard unused processor cycles as a treasure trove. An installed base of terascale setups will quickly trigger rollout of flops-hungry apps that load them to capacity, such as voice-driven features. These will doubtless be marketed with as much zeal as we see devoted to the toy apps in the present generation of mobile phones. Energetic marketing of new apps and hardware is a hard fact of modern life. On the plus side, marketing campaigns can only reinforce the natural appetites of information consumers, and the progress we're contemplating can only unfold smoothly if it comes with effective marketing to ensure consistent funding for the latest technical novelties in the pipeline.

The model household will be well equipped to entertain its occupants. The majordomo will be sure to collect information about the tastes and habits of the humans in the house, and will presort a selection of movies and television shows to offer them on request. As a householder, I shall refuse to let myself be boxed in by this presumption. Regularly, as the mood takes me, I'll cruise the movie sites through my avatar to find new stuff that cuts across all the previous profiling in my home server. I shall be equally capricious about the ambient music and about the landscapes and portraits adorning the wall screens on standby around the house. The majordomo will

endeavor to preselect reading matter for me, based on my previous reading and current projects and interests, and again I shall exercise my caprice as I pick up and discard the tablets around the house. The cost of all this caprice in terms of online micropayments for digital content will likely be small compared with the one-time cost for setting up the home network.

The size of the present and future market for information products is worth boggling at for a moment. There are apparently some four billion information nodes in the global net at present, most of them phones, and many of these nodes have the potential to expand as the perceived needs of their users increase. With about a billion households in the world and no good reason not to do all we can to bring all of them up to something like our level, we can reasonably begin to plan for a giganode global network in which all nodes are at least megascale (in terms of bandwidth, flops, and bytes), most are gigascale, and some are terascale. And this is just the home network. The networks for business and industry, health and education, and government and security, not to mention science with its petascale workloads, suggests that company stocks for broadband fiber and routers and so on are likely to pay good dividends for a while yet.

The market for good avatar software has the potential to grow to billions of unit sales. Every human being on Earth could have his or her own avatar in the global neuronet.

Once you upload your stuff to the cloud, you're on show

A person's avatar will be more than just an animated face. The software will be a natural peg on which to hang all the information about a person's health, education, employment, and so on (all protected by the appropriate authorizations, we hope). Such personal information will grow without limit as we learn to save genome maps and ultrasound scans, snapshots of educational exposure to online pages and books, employment logs with second-by-second granularity for online work, bank statements with details down to the microcent, personal mail, recorded phone calls, and such machine-sorted items as video clips from security cameras. All this stuff will need to be linked to the avatar for easy access. Getting the avatar software right for this sort of future is an enormous challenge, but also a big opportunity.

Imagine, too, the goldmine that all this avatar data will be for future historians. In the near future, we shall be able to recreate any historic figures whose legacy includes a sufficiently fat data trail with enough detail for pixel-perfect animated biopics of any desired length from a few seconds to hours or even years. A student of a future celebrity with a fat trail, perhaps a future Einstein, will be able to experience recreated lectures and meetings as if the celebrity in person, fuzzy hair and all, were there in the same room. The computational challenge here is no harder than that of creating big blue Na'vis for cinema audiences, and the cost of that computation will plummet as we build more and more petascale machines and deploy terascale setups in homes worldwide.

These visions of the future can easily run off the rails, so again let's pause for a sober reality check. Many envisioned utopias in the twentieth century went badly wrong, and we need to steer well clear of any more disasters on the scale of megadeaths like those of the Soviet Union or the Third Reich or of a nuclear war (the latter still virtual in early 2010, but very close to reality in my childhood). The disaster scenarios that loom largest now, as the global war on terror shades into the global war on climate denial, are best seen as collateral damage from precisely the rosy future of colossally increased machine infrastructure and pervasive data collection and evaluation that we're now contemplating.

With present energy-hungry technology, a billion terascale compute nodes on top of what we now have would impose a burden on the same order of magnitude as a billion extra cars in the world, which is clearly unacceptable. But unlike car technology, which is improving only by small increments as we begin to drive compact hybrids and the like, the information technology we deploy in future will be orders of magnitude more efficient than it is today. So in fact the environmental footprint of a billion avatars will be a lot less than the combined footprints of all the people with those avatars.

Balanced against the cost are the benefits. Home infrastructure sufficient for video-conferencing will enable millions more people to learn and work at home, reducing commuting energy costs by substantial margins worldwide. Virtual vacations, where vacationers enjoy immersive telepresence at exotic locations by beaming their avatars into remote robots (on which more later) but leave their

physical bodies sitting back at home, will further reduce our environmental footprints by reducing the demand for commercial airline flights. By then, too, we may have mastered carbon-neutral energy technologies such as thermonuclear fusion (again, more later) to power the huge server farms we'll need to implement global telepresence. Altogether, we can be fairly sure the environmental cost of the entire avatar revolution will be sustainable.

But environmental unsustainability is not the worst problem here. Political unsustainability, for several reasons, may be worse. The big benefit of the information revolution is that you have all your information online, where it can be used more easily than when it's buried in physical stuff such as printed paper or fading magnetic tape. Once the data is put online in a suitable server landscape, it will be indexed and stored redundantly with remote backups and versioning and so on. Since the web is designed to be resilient enough to survive a nuclear war, the data is then safe for posterity, just as long as the global machine keeps ticking over. So that's one political dependency. The world has to keep running.

A trickier political issue is sustaining data privacy and security. The debate in early 2010 over body scanners at airports and how to handle the revealing images they generate gave just a hint of the sort of issues that will explode on us at the slightest provocation. There's no reasonable way to prevent revealing data, be it as images or voice or text, from being available online, protected only by potentially hackable authentication or encryption barriers. Once you upload your stuff to the cloud, you're on show, veiled only by the ultimately political protection of its privacy and security. This is as issue that won't go away.

Politically motivated invasions of privacy are all too easy to imagine. They were routine occurrences in the former Soviet regimes of Eastern Europe and are still common in dictatorial and totalitarian regimes where a state police force has unchecked powers. This danger adds urgency to the political project of promoting democratic structures worldwide, since we can be more confident that democratically accountable governments will subscribe to the tolerable privacy practices that we've worked out in the advanced industrial countries. Only then can all Earthlings begin to enjoy the benefits of

pervasive online technology in their daily lives without the crippling fear of data misuse ruining everything for them.

There's a more sensitive issue here that we've only just begun to address. People have deep moral issues when it comes to privacy and personal data. How can we expect moral extremists to welcome pervasive video surveillance of their every public move? Yet there are new risks in the technology we're installing as fast as we can, to enable anyone, anywhere to access any information they want online and use it instantly for any purpose they choose, subject only to authorization barriers. These risks force our public servants to limit the new freedom with pervasive surveillance and inroads on personal privacy. The new freedom brings new limits.

The problem is that such limits force us in turn to trust those servants. We must trust them to account properly for any inroads on privacy they impose upon us. In any case, the complications and indirections that bedevil online data handling reduce transparency and increase the risks of improper manipulation. We can't always be sure that a website we visit is not an elaborate spoof or that a personal image or email won't be retouched or edited. Given such uncertainties, trust is essential. That trust is hard to sustain if you're a political outsider. If you're being persecuted, paranoia is a reasonable response.

These are huge themes, but they're directly relevant to our enthusiasm for increasingly pervasive information infrastructure. To take a typical example, if I'm simple-minded enough to find erotic imagery or celebration of gay or libertarian lifestyles offensive and yet I live in society where these are practically unavoidable whenever I go online for news or social interaction, then I'm unlikely to enthuse at the prospect of diving freely into the data pool. If in addition I'm a member of a minority that's mistrusted by the surrounding majority, the fear that I may be persecuted if I relax my guard is never far away when I do business online. The trust we all need to march forth in lockstep into the brave new world of everything online is enormous, and those of us who can muster that trust without reservation are lucky indeed.

In this and similar ways, complications arising from the politics of identity impose a speed limit on the highway to the global gigacity of avatars having free intercourse with each other among transparent structures that offer no resistance to copious flows of data. Before we can enjoy life in the transparent gigacity, we must bring each and

every human whose avatar should inhabit the city up to speed on the elementary civilities that will govern intercourse in that city. In my considered opinion, this will be the hardest obstacle to overcome in the course of the present century. Online living requires high levels of trust and understanding.

Returning to the humble technology of avatars, we shall need to be able to guarantee their authenticity absolutely if we're ever to trust them. Imagine the horror you'd feel if someone were to use your avatar to empty your bank account or kidnap your kids. Sensitive interactions of all kinds would need to be guaranteed safe by means of sufficiently rigorous authentication and encryption protocols. Public key cryptography is probably the best present-day technique for such guarantees (but let's not digress here to explain how it works). The problem is that it relies on the difficulty of calculating the prime divisors of large numbers (with, say, 128 digits in binary notation), and that difficulty diminishes every year. With quantum computers, the difficulty will essentially vanish. This is just an example of the new dangers that lurk when we demand high security in the new world of web avatars. There's no absolute security. There never was and there never will be. But properly managed, online living can be safer than driving on public roads.

0001
Augmented Reality

The wimp interfaces we use today are already ripe for renewal. Windows, icons, menus, and pointing devices like the mouse are mature techniques for supplementing keyboard interactions but are too static for the new reality of living in the cloud. Developed at the Xerox Palo Alto Research Center (the famous PARC), wimp interfaces first hit the big time in 1984 with Steve Jobs' brainchild, the Apple Macintosh. When I first started using a Mac in 1989, it was whole new world compared with the command-line apps on my previous mainframe terminal. But times change.

For the future, for seamless mastery of cloud services via voice and gesture, we shall need multimedia, animated, contextual, hyperlinked overlays on real-world visual arrays. That is, we'll need macho interfaces built into augmented-reality (AR) headsets.

At present, my online life happens mostly when I'm sitting at my desk and using a keyboard and mouse to do things on a big screen. This is rather limiting because I spend only a few hours each day sitting at my desk. If you own a hi-spec handheld device, you can do things on the move, but still you're offline or only passively online (hearing music or podcasts) when driving or jogging. You're always offline when doing interesting things in the physical world that demand your full attention. The problem is that it's precisely when we're doing things that demand our full attention that we're most likely to want to grab information from the cloud to help us do them. So let's see how the emerging technology of AR headsets will help.

For a start, imagine you're about to set up your complicated new home server, complete with cinema screen and speakers, other wall screens and speakers around the house, and a few tablets tuned to your local area network. You don't know your way around this technology but the guy who sold it to you said any kid could do it. He gave you an AR kit containing a snazzy pair of big sunglasses with light-adaptive lenses and attached earbuds, wired to a neckband with a power-pack pendant, and said press the red button. So here you are

at home with several big boxes in front of you. You put on the headset and press the red button.

Immediately a glowing 3D frame appears in your field of view. As you turn your head, the frame hangs back. It seems to be attached to the objects in your view and extends to fill the volume. As you look at the big boxes, they light up with bright frames projected around their edges. An animated flap image opens on the top of a box and a fat arrow rises from the open square. A voice in your ears says lift the server from the box, taking care to grasp the handles. You do so. Now the phrase "power cord" appears in glowing red letters, suspended above the extracted server. You reach out a finger to the zone in space where the phrase hovers and it becomes a red arrow, pointing with a rhythmic thrusting movement to the power socket on the server. The voice in your ears tells you what to do next. And so on. In half an hour, the task sequence is completed and the AR show ends with a trumpet fanfare and a cute fireworks display. Now you have a fully functioning home installation for all your online needs.

In this case, the AR set played an offline recording and only had to recognize the positions of the expected items, but we can guess how it might work in more general scenarios running apps that reach for online data. We already have apps for mobile phones that you can call when you point the camera at a famous landmark. The app makes an online image search to identify the landmark and uses the returned data to project some helpful text about the landmark onto your field of view. We can extend such apps for headsets.

For example, picture yourself walking along the street admiring cars aimlessly. Your headset is waiting for you to speak a word of command. You say the word "cars" and your headset responds. Its camera recognizes that you're staring at a rather bodacious car parked near you and it goes online with the image. Less than a second later, it comes back with a hit for a classic 2005 speed-wagon called the Bugatti Veyron, and projects a 3D frame of hypertext around your field of view so you can use your finger to show animations and hear sound bites about the car's cool features, such as its monstrously incorrect thousand-horsepower gasoline engine. You stand there like an idiot pointing into the air at hovering hypertext visible only to yourself and enjoying the sound bites. Unknown to you, your AR set also reads the license plate and looks up the owner's avatar online, and unknown to

your headset the avatar alerts the human owner of this gas-guzzling monster. Moments later, you see a muscular young man with a stern expression walking in your direction. Your headset recognizes his face and points him out with a flashing arrow while the voice-over tells you he's the owner. You turn away wistfully and walk off deep in thought, reflecting on the electric subcompact you drive in solidarity with the Greens.

Here's another usage scenario. You're in the kitchen with a pile of supermarket purchases on the table and you put on your AR set. You say the magic word "recipes" and a kitchen app fires up in your field of view. Your magic glasses take a few seconds to identify your purchases and then project an array of glowing words into the view. You contemplate them for a moment, then reach out and touch the word "curry". This action triggers a swirl of arrows, pointing eagerly at the relevant items on the table, and a voice then asks you how you want to prepare the curry. You point to a jar of chicken tikka sauce and a glowing halo appears around the jar inviting you to pick it up. You pick it up and let your camera read the fine print on the Tikka City label expounding in rather florid English the meticulous care that went into the manufacture of the sauce. Next, a halo appears around a pack of chicken cubes on the table. You pick it up and let the camera read how the cubes were prepared from premium soybean paste by award-winning engineered bugs in a state-of-the-art Tikka City lab in India. The voice congratulates you on your gourmet taste and tells you what to do next. And so on, again. In half an hour you have a fine dinner for two, ready to eat.

We can see that a huge semantic web infrastructure is presupposed in scenarios like these, and we can't expect to realize them for at least a decade and probably more, but it seems clear to me that the benefits are strong enough here to justify the colossal effort required to set up that infrastructure. Each and every Earthling could benefit personally from the sort of augmentation of perceived reality that such headsets provide, so marketers could sell a billion or more AR kits. The potential for fashion iGlasses would surely be easy to pump up and would likely encourage premium pricing, so the commercial incentive here is strong.

The semantic web infrastructure will evolve first and it will drive the market penetration of headsets. Apps for the sets will take time to

develop, as they do now for handheld devices, and will simply follow the availability of the relevant semantic web infrastructure. For example, as ontologies for the field of auto mechanics are defined, software agents will be developed to navigate those ontologies and assemble action sequences for various standard tasks in car maintenance, with break points for manual input from macho displays to create branching trees of alternative maintenance procedures. Once all that detail has been defined for a model series, the car company will roll out the app for its professional mechanics. I drive a BMW, and I see that BMW has already announced plans to roll out AR technology for its service teams.

Soon after the rollout of such apps, you can be sure that a secondary market will grow up to offer the same macho power to amateur car mechanics and hobbyists. Pirated custom variants will spring up like mushrooms for all the mischief you can imagine, from hot-rodding to hijacking.

AR sets can be made small enough to sell as personal effects

Before we get carried away with the joy of all this, a reality check is called for. Do the laws of physics allow us to build the required level of functionality into little pieces of hardware that we can pass off as cool? As of early 2010, the state of the art for eyewear was a Vizux Wrap 920AR prototype unveiled at a consumer electronics show in Las Vegas, incorporating cameras and small displays with combined resolution of less than a megapixel and a projected image like watching a screen at three meters. The computer power came from a Windows or Mac machine attached via USB. No one described this rig as cool enough for the street. But, hey, it's a start.

As of now, only freaks or professionals who don't care how they look would wear such AR rigs for more than a few minutes. And the real added value from living in an augmented reality comes from having it available, on call, just about any time and any place. Unless we can make the hardware setup as small and elegant as a set of traditional personal effects like glasses and wristwatch and iPod, the hope of mass adoption will remain doomed.

So let's inventory the functional items. The spectacles will contain projectors that throw bright 3D images onto the lenses with megapixel

resolution for fine text. They'll also contain a pair of cameras for 3D vision and perhaps a pair of sensors to track the owner's eye movements. Modern optical technology already allows us to make fully functional versions of this setup, so long as we're generous about the size and weight of the glasses, and there's no physical barrier to a few rounds of refinement to shrink them to any desired size. The price would be too high right now, but with production runs of many millions the price will fall dramatically. The earbuds and voice mike are available now, as small and cheap as required. Wireless antennas can fit in the spectacle frames. The power pendant is easy to make too, and can be offered in different sizes for more or less battery life and for accessorizing with different outfits.

The biggest problem is what to do with the hardware that runs the image processing software and the cool apps. But here too the problem is soluble. Current implementations of AR run on commodity Intel processors with a gig or more of memory and the current toy precursors of the cool apps run on mobile phones mounting smaller chips. For a useful setup, we'd need a chipset today that added serious bulk to the pendant. Then there's the battery power to run it. You'd have a warm brick around your neck. But here too, as in so many battlefields of the information revolution, Moore's law comes to the rescue. With chip architectures halving in feature size and doubling in energy efficiency every eighteen months or so, we only need wait a few years to shrink the desired chipsets as far as we like. Since physics allows the law to run for another decade at least, enough to shrink the kit by two orders of magnitude, we're home and dry.

The bottom line is that we'll be able to pack the logic for the AR sets into components small enough to sell as personal effects. We could even add wireless links to wristbands and finger rings for motion and touch sensors to realize hyperfine arm and hand interaction. As for the price, don't forget we plan to ship a billion units.

Thinking caps may be optional extras for AR sets

You may recall from the avatar story that another step forward is possible in principle. All the AR kit mentioned so far is external to the skin surface of the body. We could dispense with the power-hungry image projectors if we could feed their imagery direct to the optic

nerves that conduct the retinal signals to the vision areas of the brain. We could splice a socket into the nerve cable behind each eyeball and plug the AR set into the sockets. Unfortunately for wannabe Frankensteins, preventing the body from mounting an immune reaction to the socket splicing is still an imperfectly solved problem, and anyway the need for complicated surgery makes this idea impractical at present. So there's no short-term prospect of cutting our eyeballs out of the loop, I'm glad to say.

Perhaps we can realize the imagery in contact lenses. Computer scientist Vernor Vinge mooted this idea in his 2007 science-fiction novel *Rainbows End*. Vinge is a geek visionary. He earned immortality in 1993 with an amazing essay called "The Coming Technological Singularity" predicting the historic moment when machines become more intelligent than humans and take the driving seat for all our future progress. This moved AI guru Ray Kurzweil to become what I can only call the prophet of the Singularity. Ray's 2005 book *The Singularity Is Near* impressed me greatly in the conception of the book you're now reading. And just when I was writing the concept in 2009, Ray delivered a keynote at an SAP conference where I tweeted with him. As for Vernor's contact lenses, maybe we'll learn how to shrink HUD projection circuitry down that far, but I must admit I'm skeptical. Anyway, iGlasses are no big burden, and it's easier to take them off when you want to go bare-eyed.

As it happens, there's another option here. Reading brainwaves is becoming an increasingly exact science. Sensors on the scalp that read the electromagnetic waves emitted by the brain's neocortex, where we do our thinking, are already good enough to let you exercise a basic level of conscious control in a suitably configured game or app. So instead of using your voice or waving hands to control your augmented reality, both of which could be socially undesirable, you could just think your way through it. Unfortunately, the skull is a rather effective Faraday cage to isolate what goes on inside, so scalp readings are fuzzy and hard to interpret.

A better approach here would be to insert brainwave readers under the skull, perhaps between the meningeal layers that cover the cortical tissue of the brain. For example, researchers at Tufts University (where Dan Dennett is based) have shown how a surgeon can lay a thin net of electronics woven into a fine silk mat over the cortical folds.

The mat is folded and pushed through a small hole in the skull, then unfolded and shrink-wrapped over a patch of cortex. The silk biodegrades and leaves the net of electronics as close to the neurons as possible short of penetrating cell membranes and triggering an immune reaction.

Perhaps such operations can become safe and routine enough for mass rollout. But this won't happen soon. If nothing else, the "yuck factor" is enough to exclude them for now. So let's ignore Borg implants and focus on trinkets that remain external to the skin surface and rely on the traditional senses of sight, hearing, and touch, with voice and hand waving for motor control.

As an optional extra for motor control, we'll be able to wear thinking caps featuring a net of sensors to read our brainwaves. Present prototypes are crude and clumsy, but these are early days. The real problem is that you have to think hard and clearly. The effort of concentration involved is out of proportion to the benefit for anyone who still has full command of the usual body parts and senses.

The underlying problem is that the thinking cap gets such a fuzzy signal. But here technology will help. In fact, the signal is not so much fuzzy as complicated, with lots of overlapping signals from different parts of the cortex. The technical challenge is to record a large number of signals from a set of sensors dotted over the scalp and reverse-engineer the cortical activation pattern from the similarities and differences among the signals. This is a well-defined physics problem and solutions are already available. Terry Sejnowski at the Salk Institute in California has done pioneering work here with his algorithms for independent component analysis (ICA), which I heard him describe at a conference in New York in 2002. He described the problem by invoking a brain as big as New York City, reaching miles into the sky and with several billion inhabitants, where the challenge was to listen in from outside the city to conversations between individual citizens. It's not easy! As you can guess, we may have to wait a decade or two before we can build systems that not only disentangle useful signals from the number of scalp sensors we could reasonably build into a lightweight skullcap but also do this on processors small enough to hang around your neck. In any case, hair would get in the way, so cap-wearers would do well to shave their heads.

The big social payoff with using a thinking cap for headset control is that online navigation via brainwaves would essentially replace speech navigation. This would be a great relief to people like me who dislike having to share public space with people talking to virtual companions, whether those companions are at the other end of a phone connection or mediated by an avatar whose virtual existence is delocalized in the cloud.

There's a deeper problem with the thinking-cap idea. Solving the physics problem of mapping exactly where on the cortex the various components of the electromagnetic signal are located is no use unless you know exactly where the various thoughts you want to use as triggers have their neural homes. And this differs from brain to brain. So there'd have to be an extensive training phase before the thinking cap was any use to you. This is the thinking-cap version of an ancient philosophical problem usually called the mind-body problem. I want to use the thoughts in my mind to navigate in the cloud, but all I have is a piece of kit that uses neural excitations in precisely located parts of my brain to trigger the navigation. We all know that the thoughts in my mind correlate exactly with the neural excitations in my brain, with such exquisite precision that we no longer have much use for the old philosophical conundrum that thoughts are just not the same category of thing at all as patterns of electromagnetic energy in the cerebral neuronet. Yet something in that old conundrum still irritates us, namely that we can't read the *meaning* of a given excitation in the brain from the physics of the wave pulse.

Imagine, if you will, a thinking-cap enhancement of an AR kit delivered with the GoggleVision default settings. You put it on and hear the instructions. The voice tells you that to activate an online "car" search to identify something, which may or may not be a Bugatti Veyron, you just need to think of cars. So you dutifully stare at a car and think of cars. But the search result causes your view to be decorated with details about the surrounding city streetscape! You see callouts for the buildings and the shops and restaurants on their street frontages. Evidently your brain is wired up all "wrong" compared to the exemplary brain used to program the default settings. The neurons activated in the exemplary brain when its owner thinks about cars are the neurons activated in your brain when you think about street scenes.

This example may or may not be entirely realistic (in fact I'm sure it's not) but it illustrates a quite general problem. We all learn about our worlds in our own ways, and the result is that each brain wires itself up in a unique spaghetti jungle of axons and dendrites, with an exquisite correspondence that's unique to that brain and its owner. The laws governing the mappings are as universal and exception-free as any scientific laws, but their operation gives rise to different results in different brains in different environments. All this is pretty obvious, but it scuppers the thinking-cap idea for any scenario where you, as a potential user, have no desire to go through a lengthy training episode to change the GoggleVision defaults.

The way forward here is to make the training painless. You put on the GoggleVision gear, complete with thinking cap, and just go about your normal life (if you can!) for a few days while the device records what you do and what goes on in your brain and forwards all the data to an online modeling service. Then, in a colossal supercomputer computation, the service models your mind-brain mapping. It simply makes the best estimate it can of the semantic mappings between your thoughts and your neural excitations. Needless to say, given the hundred trillion synapses in a human brain and the zillions of environmental scenarios that trigger unique permutations of synaptic firings in the brain, this is all science fiction. We're not so far along yet, and won't be for a few decades at least. Somewhen before the end of the century, assuming no great catastrophe intervenes, we're likely to be able to generate such mind maps, but it will always be a major modeling effort involving huge computing resources, and my guess is that it just won't be worth the bother for the mass market.

So thinking caps may be marketed instead as optional extras for AR sets. If they are, users will either just have to get used to the defaults or go through a potentially tedious training ritual, which in turn will either be a crude manual effort where you just map a handful of salient correlations or an elaborate cloud computation that could be rather expensive. In any case, then, the net benefits will be less than we might have hoped. Just as many of us are content to interact with present computers via Qwerty keyboards whose ergonomics are from the nineteenth century, when mechanical typewriters needed key layouts that would minimize jamming, so many of us for most of this century will probably interact with the cloud via headsets that use voice and

gesture input, whose ergonomics are from the days before mind-brain mapping became an exact science.

The science-fiction idea of plug-in neural interaction via Borg sockets in the back of the neck is crippled by the same requirement for expensive customization. The medical hazards of implant rejection and the yuck factor are only the most obvious problems. Once you've inserted a bundle of fibers into the brain, what then? Initially, you'll get random readouts and only after a tediously extended period of probing and mapping will you begin to develop any ability to send and read signals that carry useful semantics.

The big benefit of Borg sockets is on the provider side. Once your lab has mapped a customer's neurons, you can hijack your customer's brain and make it do anything you want. You'd effectively have turned the person into a selfless Borg drone. Forget about their desires, emotions, and feelings – you can override all that stuff with the cloud input to the neurons. Your customer would become your zombie. For this reason, I'm pretty sure that this sort of approach will be banned as unethical. The brain will be sacred. The only exceptions will be for medical interventions with specific therapeutic purposes.

By the way, the medical implications of this technology are interesting independently of Borg implants. A headset supplemented with a tiny heart sensor and a hair-thin blood-pressure band could offer good data for online health monitoring. Regular measurements over many years, correlated with recorded activities such as getting drunk or jogging, would be enough to diagnose medical conditions far more accurately than today. Also, a headset could augment the wearer's vision or hearing quite massively without calling attention to the wearer's deficits. If everyone wears headsets anyway, a person who'd be functionally blind or deaf without one would hardly stand out. So the medical community will use headsets to deliver better and more comprehensive health care.

AR headsets may need headtop antennas

At first, AR headsets will be relatively simple devices. They'll be hardly more complicated than a present-day gaming station. And they'll be hardly more intrusive for users than present-day iPods worn with sunglasses. But the revolution they trigger will be big. This has the

potential to become the main paradigm for human interaction with the cloud. Why fill the world with big, clumsy, energy-hungry screens when you can use virtual screens beamed from projectors implemented in picotech that can sit on the bridge of your nose? But then you have to go wireless.

Think about it. Each headset will have a broadband wireless connection to the cloud. Our present woes with connectivity for phones and organizers will be seriously exacerbated if everyone in a public space is continuously downloading megapixel imagery for an AR set. Only rigorous policing of spectrum resources can prevent this issue from becoming a nightmare.

One way to approach the problem is to go from microwave to infrared signals. Base stations can offer gigabit connectivity if they use line-of-sight infrared rather than much fuzzier microwaves. But line-of-sight signals will need to be tightly zoned, with automatic frequency hopping and elaborate packet switching for zone handoffs. Headsets may need headtop antennas that infrared lasers can target via wall-mounted mirrors and so on – a nightmare of complexity.

Any telecoms engineer who prefers to avoid the nightmare will recommend instead a pervasive public infrastructure of sockets and big screens where people can do their business via cable. Headset people will either plug their pendants into the sockets or communicate wirelessly via their headsets with sensors just centimeters away. In either case, their avatars will log them on automatically and securely. Big screens will offer online services to people who for any reason don't have headsets.

A general public acceptance of headsets poses further problems. Will it be legal to drive wearing them? Surely not, for they'll be seriously distracting, and the data they generally display will be irrelevant for driving. Rather, cars in future will offer their own built-in reality augmentation, custom-designed to enhance the driving experience and "blinkered" with software sidewalls to ensure that they never show obviously irrelevant data.

What about children? Will they be allowed to wear headsets when young or will there be a legal age cutoff? Educationalists might reasonably insist that kids learn to get along without such enhancements until they reach a sufficiently ripe old age. Certainly, students will be forbidden from wearing such sets in school exams. And this will

exacerbate debates about the relevance of exams. In a world of pervasive connectivity, it may be absurd to insist that students have their abilities tested offline. But let's skip that debate – the future of school education is a big and separate issue.

AR headsets add value to semantic web infrastructure

Shelving all such policy questions for now, the biggest practical issue facing us as we contemplate an AR future is the semantic web infrastructure it presupposes. The infrastructure enables the apps that make the headsets worthwhile. Conversely, the headsets promise added value for semantic web infrastructure that can make all the difference to its feasibility in terms of generating revenue to pay for the infrastructure investment. So let's take a more serious look at the semantic web and its impact on AR apps.

Let's start with vision. When a computer processes an image from a camera, it can find surfaces and edges and make a 3D model. Perhaps it can even make inferences about mass and material and speculate about moving parts and functional features. This may be enough for tasks such as driving a vehicle through a standardized street scene with a finite inventory of possible obstacles. But the higher realm of semantics – what a given thing is and what use it is to humans – is likely to be too much for an isolated machine, at least if it's outside an artificially constrained environment.

The promise of the semantic web is that a computer can help itself to semantics by going online. A few salient object features can serve like a badge to proclaim membership in a particular ontology – a collection of things that form a class or category for certain purposes – and the computer can simply seek out the facts about that ontology online. In the worst case, an image of an unclassified object can be use to search for similar objects, which may already be allocated to ontologies. With the right cloud engines and good connectivity, all this can happen in a second. Once membership in a specific ontology is established, the computer can read off the relevant properties of the object in the image and react accordingly.

That's the principle as it relates to image processing. In the case of text processing, the principle is simpler. The computer simply looks up the word in a dictionary that assigns it to an ontology and then reads

off the properties from the ontology. Here the complication is that in most cases a given word can be in any of several ontologies, depending on its context. For example, each car manufacturer might maintain a specific ontology for each model in its product range, so a generic car word like "gearbox" would have its exact meaning spelled out in a hundred different ways, depending on manufacturer and model. The solution here is to build stacks of ontologies, from generic to specific, and link them in hierarchies. Now a search for an exact ontology for a given word goes in stages, where the other words forming the context of the given word are evaluated in parallel to help the app disambiguate the search as it navigates through a menu tree of increasingly specific ontologies. This may sound like a lot of work for the app, but such complications are routine for search engines today and normally add only a few milliseconds to their response times.

So machines can handle meaning. They can recognize objects via pattern matching and feature lookup. As a skeptic, you might complain that this is not real understanding. But what do we know? How do humans understand things? We recognize objects by becoming aware that they fall under familiar categories and we know their features via the classification, together with background or contextual knowledge to handle exceptional cases. All this is paralleled in the machine cases. There's nothing mystical about human understanding, as we can see from the sad history of human misunderstanding. Machines will understand things more consistently than humans do, until one day we're ready to concede defeat. Like machine chess, machine understanding may seem to be based on a primitive approach, but it will work.

That said, we shall certainly hold out for a long time – several decades, at least – before we concede defeat in the areas that mean most to us. It's easy to let the machine have the last word as to whether the car in front of me is a Bugatti Veyron or a Chinese lookalike, or whether the next corner is safe at fifty miles an hour, but we shall find it much harder to accept the verdict from the cloud as to whether Shakespeare's Hamlet was paranoid-schizoid or whether "God" and "Allah" denote the same concept. For me, I find it hard enough to accept that the Microsoft spellchecker is usually right – one sees all too readily the human errors behind the machine interface. However, once the global web of machines in the cloud begins to trace

the larger patterns in our all-too-human ontologies, it will deliver verdicts that are likely to be more defensible than the opinions of many human experts. The machines will learn from our experts and apply their hard-won wisdom consistently. We could hardly wish for more.

A standard philosophical complaint at this point is that machines are admittedly great at handling rules consistently, but the precedents on which the rules are based are unique cases, for which the machine approach falls flat. Hamlet is as good an example as any. Shakespeare invoked the mighty power of a fertile human imagination to create Hamlet. He did so freely, unconstrained by any deliberately applied rules. Whether the character thus conjured into fictional being falls under this or that psychological category in an index of psychiatric disorders is naturally a matter on which reasonable people can differ, and also an issue that's entirely secondary to the artistic and philosophical interest of Hamlet the archetype.

A machine verdict can carry no higher authority here. A developer familiar with the facts the machine used to calculate its verdict, and the rules it invoked as it did so, might be able to point to a dubious source or a poorly motivated application of a rule and discredit the machine verdict. The short history of software development includes plenty of cases where just this happened. Even the history of spell-checkers suffices. So the complaint here is that the machines can't do anything here without humans to boot the whole process into life.

The right answer to this complaint is simply to agree. Of course the machines need us to get started. Semantics is essentially a pattern of deeply rooted human traditions about the use of words, as the Cambridge philosopher Ludwig Wittgenstein argued in the twentieth century. We can reduce the traditions to precedents and rules, then supply the precedents and rely on the machines to apply the rules. It's as easy as that. In the context of the semantic web, we codify the precedents into ontologies and invite the machines to apply simple associative logic to match their input to the fact base represented by the ontologies.

Seen at a high level like this, the project of preparing the way for the semantic web is much like what we're doing anyway to make the contents of the web more amenable to systematic search. Google and other web search providers index the web in many different ways, but

a lot of what they do is effectively classification. The engines map search terms to classes of potential hits, where each potential hit is a document containing the search term, and then sort the class according to such criteria as word frequency (a.k.a. TF-IDF, for term frequency – inverse document frequency, which is the default way to rank results in simple search engines) or number of incoming links (as in Google's secret sauce, the PageRank algorithm), until a ranked short list of actual hits is generated, which is just a neat little ordered class for the search term. As the precedents and rules for this enterprise become more established, the classifications created become more established too, and indeed they become established as substantial new extensions to the semantics of all the words involved. In this way, the semantic background of our use of words and images grows and changes.

Semantics is always a work in progress. The creation of the semantic web will in large part be the creation of a new hierarchy of ontologies that fixes the background of meaning in our lives in ways we never anticipated. When we start to leverage semantic web technologies to create augmented realities via macho interfaces at the edge of the cloud, we shall bring meaning into our lives in new ways.

The effect will sometimes be mind-blowing in the old hippy sense. Imagine the simple act of reading a paper book or magazine while wearing your magic glasses. Instead of seeing merely a page of text, you see the text framed by colored toolbars poised in 3D space. Your glasses recognize the text via an online search and use the search result to illuminate the text for you. You see words of interest in the text jiggling or glowing relative to the other words. You reach out and point to one of these words, and a colorful pop-up zooms up out of the word to offer a little video illustrating some aspect of the word. Perhaps you also hear a spoken sentence that explains the video. If the word refers in some way to a website, pointing at it takes you instantly to the website.

All this could get mightily distracting for a serious text, and many people will turn off all the effects and stick to the plain words. But the temptation to blow your mind will always be there to escape the tedium whenever you're unfortunate enough to be reading dull stuff. On the other hand, the illumination of novels involving sex and violence could become rather pornographic. In such cases, you may be

worried that your headset is reporting your deviant reading to a thought-crime database somewhere.

Privacy will be lost and gone when you wear a headset

Here we come back to the questions we asked about avatars a while back. The politics of all this threatens to become seriously disturbing. Privacy will be lost and gone when you wear a headset, because any cloud node you access as you wave your hands in your augmented reality will know that you, there, with that micropayment account, are receiving indexed and logged content. Before you sign on to that sort of nightmare, you'll want solid assurances that some level of personal privacy is being maintained, perhaps by automatic encryption of all sensitive information. All this will need a regulatory framework with legally binding safeguards.

Anyone who, for any reason, mistrusts the political authorities will mistrust headsets too, and indeed will likely mistrust the whole apparatus of the semantic web. As Tim Berners-Lee explains in his lectures about all this, the capstone of the whole edifice is trust. In the AR future, people who wear glasses shouldn't mistrust the system that augments them.

In practice, delivering satisfaction on the trust question will devolve to the corporations that provide the augmentation. Google's famous "don't be evil" mantra is a hint of the sort of image the service providers will have to cultivate ever more intensively. Any company with online responsibilities will have to work hard to keep its image well polished, primarily by training and retraining its employees to behave ever more correctly. Public trust in a company will become one of its most valuable corporate assets.

Big companies generally will devote ever more resources to guarding their brands. Employees will be ever more challenged to live up to the ethical and moral standards that best increase the brand value. In a world where religious or patriotic principles fail to provide the moral support they once did and where personal standards can be dismayingly eccentric, we can rely on company guidelines to take up the slack.

An emphasis on principles and corporate ethics is never out of place, even in mostly offline world. But in a future where customers

regularly bet their lives on the reliability of the services delivered through their headsets, the service providers will need to prove convincingly that they're doing the best they can to anyone at any time. This will be hard to do in a globalized yet diversified world, because customers in challenging circumstances can have challenging standards as to what they're prepared to trust. Too bad – we pay corporate leaders enough to expect them to work their butts off for us.

0010
Bod Pods

The twentieth century was the century of the automobile. No other technology so gloriously symbolizes the advance in human lifestyle achieved in that century of miraculous and atrocious developments. The internal combustion engine was just an enabling technology. Aircraft and spacecraft can be seen as exotic variants of the automobile for traveling in air and space. The computer was a technology for refining the car, from methodical design through mass production to safe and reliable operation. The two upstart powers of the twentieth century, Germany and Japan, grew rich through their car industries. Americans had a love affair with the "auto" that was still going strong in 2010.

When we cast around for a technology or device to symbolize the new century, this example or precedent looms large. Cars, in the more evolved forms that will appear in the coming decades, will continue to shape our lives in many ways, but increasingly they'll fade into the accepted background as appliances beside others. In part, this will be due to regulations governing their safe and responsible usage. It's hard to see a car as a symbol of freedom when using it is so surrounded with restrictions like licenses and speed limits and parking tickets and seat belts. We're sure to see increasingly restrictive limits on emissions and fuel mileage for new models, with the result that the fun will slowly leak out of the whole game. Cars will be more like washing machines or refrigerators.

In this century, the technology that promises to play a comparably decisive role is robotics. But before we get into that subject in all its fascinating complexity and threatening potential, let's dwell for a while on how cars are likely to evolve in the coming decades. The forecast reflects how our lifestyle more generally is likely to change. Roughly, as robots become more important, cars will become less important, largely for independent reasons but partly for the shared circumstance that the inventive engineers who would have worked on refining cars will find more career potential in creating new robots. Since any society has only a limited supply of gifted engineers, the result will be

like the one faced by the British car industry during my younger days, when the best engineers were otherwise engaged designing nuclear and military systems, aircraft and spacecraft, and mainframe computers. The indigenous car industry stagnated and then withered in face of competition from Germany and Japan. With robots as the new attractor worldwide, the car industry in all countries, China and India included, will lose some of its drive.

We shall see dramatic transformations in car design

First, we shall see dramatic transformations in car design. The General Motors Volt concept is a harbinger of things to come. The Volt is essentially a car-sized electric skateboard on which you can have any body you want, from speedster to pickup truck, either with a gasoline engine in a hybrid configuration or without gas, in an all-electric setup. The running gear is standard to an entire product line. At present, the concept is limited by battery technology. Either you weigh down the skateboard with a hybrid power train or enough batteries to give a decent range, or you skimp on batteries to get breakneck acceleration but risk running out of juice on a long journey.

Battery technology is improving slowly and incrementally. At present, lithium-ion power packs are the ones to beat. Soon, new designs, which may be based on lithium but also use nano-engineered electrodes, promise to multiply storage capacity and hence vehicle range, and ultracapacitors may supplement them by offering fast charge and discharge capability for hard acceleration on demand. Alternatively, the Swiss company ReVolt announced in late 2009 that its zinc-air batteries delivered three times the energy density of lithium-ion batteries by volume in safe and inexpensive packages. Also in late 2009, Arizona-based Fluidic Energy was developing a metal-air battery offering ten times the energy density of the best lithium-ion technologies for a third of the price. Ultracapacitors can augment battery packs. Although they can't store much energy, their fast charge and discharge enables designers to downsize the batteries, reducing their weight and cost. But all this will probably be an uphill slog. Any known technology for storing and delivering electrical energy is a compromise between expense, size, weight, capacity, speed, and safety. For some years yet, we shall follow the Maoist principle of

letting a hundred flowers bloom until we begin to see which technology option makes most sense in the new world of compact electric cars and smart electric power grids.

For a century we've been spoiled by our own good fortune. Gasoline and similar liquid fuels that store energy in the chemical bonds of combustible molecules are an impressively efficient and user-friendly way to deliver cataracts of power in wheeled vehicles. Because the fuel just squirted out of the ground in some places, and was available in huge quantities for the price of a drill rig, it was often cheaper than clean water. What a way to go! Almost every adult male in the Western world learned to deploy a hundred and more horsepower to satisfy the most fleeting whims as we drove hither and thither to beaches and bars, pumping out toxic and greenhouse gases as we went. This was the epiphany of the "me" civilization, where basically we all felt we had the right to do what we wanted, when we wanted, without forethought or regulation.

Those days are gone, at least so far as haring around the scenery at high speed is concerned. The traffic network in all developed countries will be increasingly regulated like a huge integrated computer network. Each car will probably soon be legally required to be permanently online, advertising its position, speed, license plate details, and possibly a whole lot more, such as driver identity and bank details, load or passenger information, anticipated route and destination, present and average fuel consumption, maintenance and safety check status, accident and insurance history, and history of contacts with known terrorist organizations. The entire network will be permanently monitored for traffic flow, journey times and speeds, and environmental impact, and will be regularly audited for various purposes, such as flow optimization by adjusting traffic light timing or road lane allocation, pollution control in city centers and speed control on freeways, and the like. We shall accept all this without protest as we learn to appreciate the benefits of driving cars that are always online.

What are those benefits? The obvious one is that our journeys will be smoother as the traffic lights change to green when we approach them. Also, breakdown and accident services will have a much easier time finding you if your car can report the bad news automatically to the right nodes in the network. At present, BMW offers head-up displays projecting 3D information through the windscreen for its

high-end models. Soon, we can be sure, all new cars will feature HUD info on the windscreen, with easy-to-read iconic representations of traffic information such as upcoming tailbacks, foggy or icy stretches, wide loads, and police controls or speed cameras. All this will be integrated with augmented satellite navigation via windscreen displays and bundled with animations to advertise cheap motels or burger bars, railway junctions with the barriers down, roadside hazards or suspicious packages, and warnings at road works alerting you to reversing dump trucks. And it will only take a few well-publicized cases of car-thieving terrorists turning their booty into bombs to convince us that letting the car do an online handshake check with the driver and passenger avatars is good for us too. If you weigh all these valuable benefits against a slight loss of privacy, you're not going to complain.

As I contemplate this wish list for car designers and traffic planners, I see a good case for giving cars their own avatars. Any car will need to be pretty smart to cope with all the complications we've just reviewed of online living, and it certainly won't do any harm if the car can organize itself at least in basic ways. Any self-respecting car will wish to be able to respond instantly if its owner decides to call it from the office, say, to enquire about its battery charge status or to check whether the next service has already been scheduled automatically (between car and dealer), or whether it (the car) feels up to a long trip tomorrow. If I call my car while strolling in my iGlasses, I'd like a cheeky cartoon car as a visually amusing avatar to appear in front of me, with animated headlights and grill to indicate its mood as I engage it in witty banter. I shall be aware that a cloud engine with no clue about moods will have synthesized any words the car directs at me, so the whole interaction will be a silly game, but it will succeed in making the extraction of basic information more fun. More seriously, the car's avatar may be smart enough to plan trips online, at least to the extent of checking traffic updates or scheduling a recharging stop on a long journey.

More seriously, too, a future electric car should be able to interface smoothly with long-distance travel infrastructure. I suspect that we'll soon be prohibited from driving non-stop for more than a couple of hours and be directed instead to trains.

Trains may convey our cars, with us in them, from city to city

Trains (yes, good old railway trains) may convey our cars, with us in them, from city to city or country to country. This arrangement will offer several big benefits. First, the environmental load of a hundred cars racing independently along the freeways from, say, Frankfurt to Rome is likely to be heavier than the load of a fast train with ten wagons, each with ten docking stations for the cars, which can be recharging the cars en route as a bonus. This depends on how efficient the trains can be, and doubtless the infrastructure costs for the railways will negate much of the benefit, but if advocates for rail travel are right that trains are the way of the future, this scenario will find favor on environmental grounds.

Let's think this through awhile. A (possibly robot) train with a hundred cars on board will be much heavier than a hundred cars, so moving its mass will need a lot of power, but trains travel in relatively straight lines and don't have to start and stop too much. If the trains are electric, their power comes from the grid as it does for the cars on a freeway, but the opportunities for optimization, such as putting rooftop solar panels over the car decks on the wagons for that extra two-kilowatt boost per wagon on sunny days, locating stations on hills to provide gravity assists for acceleration, and use of regenerative braking, may be significant.

One major optimization would be for the locomotives to generate their own electricity from a sufficiently benign source (other than diesel oil, that is). The source that comes to mind for me is a small nuclear reactor that can be inserted and exchanged as a sealed unit. As an example to get us started, physicist Otis Peterson at the U.S. Los Alamos National Laboratory invented a modular nuclear power reactor now marketed by a startup called Hyperion Power Generation. A Hyperion power module is sealed and self-contained, has a lifetime of up to ten years, generates 25 megawatts of electrical power, has a size that would just about fit into a locomotive, and weighs less than 50 tons. The fuel is enriched uranium, loaded at the factory. The design is inherently safe and includes redundant shutdown systems. We'd want to scale the module's size, weight, and power down by a factor of two or three, but I guess that would be easy enough. Still, I see the

wider risks any such units would introduce and I don't want to pre-judge the issues of political security and public acceptability.

Set against the risks, the benefit of letting locomotives generate their own electricity is that we avoid transmission losses, which can easily waste a third or more of the power in grids where the power is generated centrally and distributed via overhead lines or trackside live rails.

Altogether, it's hard to say how big the environmental benefit of using trains will be. We have a complicated equation with a lot of uncertain numbers in it. Still, it may be that the impact of taking cars off crowded roads and putting them on trains will be good.

The second big benefit of putting cars on trains is that their drivers can relax. I've lost count of how often over the years I've driven between Germany and Britain via the Channel Tunnel, where we drive onto a train and sit in the car for half an hour as we speed along under the waterway. It's a welcome break on a long journey. Since future cars with online services will be fun to sit in for hours if necessary, watching a movie or whatever, I see no reason at all not to prefer this mode of travel to the irritations and frustrations of airline travel when I think of underwear bombers and volcano eruptions.

My big vision for cars this century is that we'll come to see them as modular lifestyle capsules or bod pods for sustaining us in reasonable comfort under all external circumstances for hours at a time. Their mobility, such as it is in our traffic-choked times, will increasingly be seen as a given, obvious and unremarkable, and it will anyway be realized for many journeys via drive-on car trains or sea ferries, or perhaps even rail-gun tunnels like giant particle accelerators for super-fast travel.

Your car will be a package in a global federal express system

Let me explain. If you secure a car inside a tight-fitting metal box or can, you can fire the box or can (plus car plus driver) through a tunnel by subjecting it to crosswise and mutually orthogonal electric and magnetic forces. This is the principle of the electromagnetic rail gun, which can accelerate metal slugs to well-nigh relativistic velocities. People in cars don't want to be shot up to relativistic velocities, of course, but I can imagine speeds high enough to slash travel times for

intercity trips, so much so that there'll be no contest with airlines. Then the clumsy and inefficient use of aircraft to convey large numbers of people between cities can become a thing of the past.

An obvious prerequisite for these assisted-travel scenarios is that cars be more standardized in external dimensions than they are today. But this will trouble no one. Cars have been standardized in so many ways in recent years that another step in this direction, even one that enforces tight limits for trains and rail guns, won't bother many people. Anyone who wants to drive a giant land cruiser will simply accept as downside the loss of the assisted-travel options.

Another obvious prerequisite is that cars will be marketed with more emphasis than today on their interior fittings and trimmings. The car will be less of a rainproof motorbike and more of a space capsule, or your own personal and customized survival pod in the big, dangerous world of international mechanized courier delivery services. Your car will be a package in a global federal express system where it gets handled by robot forklifts and shot through tunnels and plugged into railway wagons and run through body scanners, all with you inside, trying to distract yourself in an online virtual reality of your own choosing.

We can take this "car as bod pod" idea further. Imagine a hideous future where numerous cities have been nuked and are now no-go zones with hazardous radiation levels. To protect its citizens from the marginal hazards in the less radioactive zones, the nuked state has mandated that all cars driving in its territories be built as sealed capsules with fallout filters and radiation meters. Your best hope of surviving in this state is to stay in your car. Perhaps your home is hermetically sealed, too, with airlocks for entry and egress and elaborate air and water filters and sensors to keep the interior clean. And you work in a sealed office. Your whole life would be either within sealed spaces or inside your car, traveling safely between them. Your car would be your spacesuit for all excursions outside the airlocked base stations.

You'll check in your car where you like and a robot will park it

Returning to the comfortable muddle of our times, one of the daily irritations of car ownership, at least in crowded Europe, is the need to

keep pocket change for car park fees. Wherever you drive, you need to find a legal parking space within a sensible distance of your destination and be ready to pay for it. Happily, car-parking rates are still cheaper than hotel rates for parking the human body, but still this is an aspect of life where big investments for a better future would pay off handsomely. In Germany, "park-houses" in all big cities are linked to street signs that show how many places are still free in the various houses. I can imagine robotized park-houses where you simply leave your car in the foyer and a robot forklift puts it into an elevator. The robot management keeps track of where the car is shelved and notes who will pick it up and pay the rent (respectively, you and your avatar, whom you'll command to pay the bill with a wave of your digitally enhanced finger). When you return to pick up the car, you wait for a few seconds in the foyer while the forklift valet fetches and delivers it to you.

That future may be only a decade away in German cities and a decade or two further away in more bucolic regions of Europe. Its key benefit is not the convenient valet service but the saving in real estate. Robots can stack cars without drivers much more densely. Human drivers need lots of ramp space and wiggle room plus space to open their doors. Deeper in the future, when city infrastructure is more intensively optimized to keep cars out of city streets, there may be rail-gun tunnels linking the park-houses in a city so that robots can shuttle empty cars between houses. Then you can check in your car where you like, preferably on the city outskirts, without having to cruise from house to house in search of a space. Once you've dismounted, you can take an instantly available robot taxi or shuttle bus to your destination. Later in the day, you can reclaim your car from any house you like, so long as you don't mind waiting a few seconds for your car to be shot to your location.

Robot drivers are the next topic we should consider. Cars that can drive themselves are still rather chancy. It's not yet safe to let them loose on our city streets. But these are early days, and it's not hard to imagine that robotics engineers will learn fast and build useful models within a decade or two. Before then, each new car model will feature another stealthy step forward in the automation of the driving task, for example by taking over for parallel parking, as in some current models. Because people usually like to drive themselves, and often

fancy themselves as pretty good drivers, the automation that cars offer is likely to remain discreet and unobtrusive for a while yet. Only when robot systems are so mature that the government regulators can mandate them for safety-critical aspects of the driving experience, such as keeping your distance from the car in front or checking other cars before changing lanes, will humans begin to notice that their sovereignty in the driving seat is an illusion.

It seems unlikely to me that robot drivers will take over completely. I'm sure they'll take over for driving big cargo trucks, if only because they'll be cheaper than human truckers and be able to drive in much tighter convoys in the slow lane. Similarly, I expect robots will soon be driving specialist vehicles on private land such as company stockyards or on airfields and military bases. But the challenge of building a robot that can impress the average human driver sufficiently to make him or her let go the wheel voluntarily seems daunting to me. Perhaps it's a generational thing. Perhaps future people will just lose interest in steering themselves through the scenery.

We can regard our cars as our physical avatars

Returning to cars as bod pods, in my youth it was a cliché about life in America that people were so wedded to their cars they'd soon lose the use of their legs. As late-onset diabetes lames more and more obese Americans, I guess this joke is more serious than it seemed. But the kernel of interesting truth here is that cars work for people as mechanical exoskeletons. They're our hard shells that give us the superhuman power to travel fast over long distances while remaining protected from inclement weather.

We can regard our cars as our physical avatars. They serve as our embodied presence in the great outdoors. Or to return to another metaphor, our cars are our spacesuits. They support our survival in the big spaces of our wider communities. Next time you see cars parked at natural beauty spots, where their owners make unsuited excursions into the raw elements of sun and wind and rain, think of humans letting their dogs off the leash, to let them make their doggy excursions into the raw elements. Cars are not yet autonomous life-forms, and probably never will be even when they have avatars and robot drivers, but humans inside their cars can often seem like a

different form of life than pedestrian humans. To see this for yourself, try hitchhiking alongside a major freeway where the cars zoom fast and furious past you. You'll feel like candidate road kill in a world of roaring metal monsters.

The post-nuclear apocalypse scenario suggests another way to see how bod-pod cars can help us. Imagine another apocalyptic future where we've somehow ruined the atmosphere. Perhaps it's just too much carbon dioxide and methane, perhaps it's high concentrations of sulfur or nitrous oxides or ozone, or perhaps it's carcinogens and heavy metals. In most such scenarios, the required breathing apparatus for "naked" humans would likely be more burdensome than most people would wish to endure. Jogging in a gas mask is not exactly an exhilarating experience. For our cars, it's no problem. Already most cars have air conditioners and microparticle filters, so adding sulfur scrubbers and so on would be no big deal.

We can take the idea further. Any cars that rely on combustion in our proposed future, for example models that burn hydrogen in fuel cells, will gulp in oxygen and pump out something else, for example water vapor. As they breathe the atmosphere, it would be a simple matter to mandate that they simultaneously pass their breathed air through a catalytic converter that corrects whatever atmospheric damage the government, in its wisdom, says is the most urgent danger. With a billion cars worldwide, each doing its part in the battle to clear the air, we could hope to undo some of the worst damage we did to our planetary gas mantle in the earlier days of the industrial revolution.

Cars will most likely be electric

A few years ago, the hydrogen economy was billed as the wave of the future. Transporting hydrogen was seen as more efficient than conveying electricity over long distances. It may be that hydrogen still has a big future. But there are drawbacks. Fuel cells for burning hydrogen to generate electricity locally, either in a car or in a neck pendant for your iGlasses, can apparently be made efficient enough to do the job, but distributing and delivering the hydrogen to the cells is still essentially an unsolved problem. Hydrogen is much bulkier for the energy it delivers than, say, gasoline, and no storage technology has yet emerged as a clear winner.

Adsorption in a porous matrix medium analogous to charcoal is the most promising approach, but the realistic options for the media here are much heavier than the adsorbed hydrogen, which excludes matrix solutions for mass-critical applications like aerospace. It even excludes them for neck pendants – who wants to lug a brick around to power their iGlasses?

The cryogenic approach of liquefying the hydrogen, as in the big old NASA rockets, isn't really workable for everyday scenarios involving cars and gas stations, as extended experiments by BMW and other carmakers have shown clearly enough. The deep-freeze tanks are bulky, heavy, and expensive. Worse, if the cryogenic insulation fails, the hydrogen evaporates and you risk explosions that recall the Hindenburg airship disaster. We don't want people driving around with hydrogen bombs in their cars.

No, the hydrogen story convinces me that cars will most likely be electric. Only off-road land cruisers will need combustion engines, and what they burn will doubtless be heavily taxed and regulated. Since electric cars, for all their potential, are best at being city skates, the result will be the domestication of the car.

Cars will become lifestyle modules, or pods, where people can be themselves differently than at home under the watchful eye of the majordomo avatar. Perhaps car avatars will be programmed to be especially indulgent of teenage sex.

Cars will certainly be set to crack down hard on drunk driving and the like. Your car will give you a health check whenever you take the wheel – in fact it might be your first line of defense against chronic health conditions that you're too lazy to monitor. (Actually, your humble loo at home will more likely get smart enough one day to do that job, but let's skip that topic for now.) A sensor in the car dash will smell your breath and detect any traces of alcohol, and sensors in the steering wheel will check your sweat and pulse readings. The seat will monitor your weight and nag you if the seat settings are bad for your back. The on-switch for the motor might demand that you enter a password or read a captcha to check that your brain is up and running. In all these ways and more, your car will become more personalized and less tolerant of your more dangerous or dodgy ways.

We should move to carbon-neutral transportation

When we finally stop burning petroleum spirits – the spirits of long-dead dinosaurs – in our cars, two big consequences beyond the technical questions we've been considering will follow. First, the environmental burden of our civilization will be ameliorated, and second, the oil-rich states in our world will need to find other ways to make money. Let's start with the environment.

Rising levels of atmospheric carbon dioxide are not only a great opportunity for governments to invent a new tax but also a threat to the people-friendliness of the global climate. If global warming triggers massive melting of polar icecaps, we suffer a climate catastrophe that puts most of my techno visions on hold. If polar icecaps melt, sea levels rise, perhaps by a meter this century, and we shall need to devote huge resources to reinforcing coastal city defenses or relocating cities. In addition, huge numbers of poor people living in low-lying coastal regions will need to relocate their entire communities. If, at the same time, some tropical regions become deserts because rainfall patterns change, we suffer global upheaval on such a scale that business as usual is no longer an option. This could effectively delay all our future scenarios for this century by a decade or two.

The history of human folly in face of long-term doom is sobering. I have no confidence that we shall do the right thing in time to head off a climate catastrophe. But I do have the opportunity, here and now, to remind us to do whatever we can to limit the impact of any such catastrophe. One way is for governments worldwide to mandate a move to carbon-neutral transportation as quickly as possible. This covers cars, trucks, trains, ships, and planes.

The simplest way to make the move is to make land transport electric, sea transport nuclear, and air transport solar. Making land transport electric is a no-brainer so long as the power nodes in the electricity grid are environmentally benign. Depending on their typical missions, the cars and trucks might carry on-board generators burning gasoline or diesel oil synthesized from air and water by engineered bugs living on solar farms. Similar bugs would also synthesize rubber for their tires to reduce the carbon footprint of manufacturing them.

As for the rest, nuclear-powered ships (and trains too) will remain a political hot potato for at least as long as terrorism and piracy

remain threats, and it may be more expedient in the shorter term to let ships burn synthesized fossil fuel. Solar power for aircraft seems absurd until you recall that it would be indirect, via kerosene made by bugs on solar farms.

As you see, there's a lot of room for negotiation in the details here, but the overall goal is both clear and reachable. For our purposes here and now, the main thing is that the car industry will go electric, if only to insure itself against the risk of getting hobbled later when a climate catastrophe begins to bite.

As for the oil and coal industries, we shall learn to see them as environmental disasters. We should leave the spirits of the dinosaurs underground where they belong. The carbon is neatly sequestered and needn't be touched, except perhaps to feed the global appetite for plastic and similar highly refined goods. Oil-rich states will need to find other ways to pay their way.

The Islamic belt of huge desert regions across North Africa and South Asia has another natural resource that can be tapped for profit, namely solar energy. The solar power that we could collect from vast solar arrays across the deserts dwarfs the power we now derive from burning oil. If only a small fraction of the Sahara Desert were paved over with solar panels, for example, we'd have more than enough power to keep the wheels of European industry turning through the present century and beyond. A pilot project in this direction called Desertec was formally kicked off in October 2009. The idea is that solar thermal arrays for generating steam to run turbines will generate electricity in North Africa, and this electricity will then be conveyed efficiently as direct current under the Mediterranean Sea to the big European power grids.

Later, as the technology matures and prices fall, we can supplement such arrays with big solar voltaic arrays. At present, the best photo-electric arrays convert about 20 percent of the incident photons into electricity, but they're expensive. To give some ballpark figures, a square-meter panel delivering 100 watts might cost 600 dollars, so it could cost you 100 grand to cover the sunny side of your house roof and generate enough power to stay cool and keep a few cars charged. But we can hope for progress here.

Generalizing the Desertec concept for other desert regions should be easy enough. Since Israel currently has strong expertise in such

solar technologies, the global community should press Islamic belt countries to do deals with Israel to bring the wider vision to fruition. Of course, there's a political deadlock between Israel and the Arab states to settle first, but we need a big vision to break out of the deadlock.

We've arrived at the topic of geoengineering, where the scale of the projects becomes planetary and global political coordination becomes essential to their success. Happily, the car industry itself sets a good precedent for how such coordination can be ramped up steadily without frightening away the partners. Looking at the global organization of car manufacture and the effective globalization of standards governing cars, for example regarding safety or emissions, we see that at the technical level globalization is already more of a fact than an ideal.

That effective globalization of the car industry was achieved in part as a result of global sharing of opinion or sentiment on features or issues relating to cars in the consumer marketplace. Online sentiment tracking is now big business, with giant machines sifting through the world's emails and web pages for references to products and for hints as to how they're going down. If there's a problem, the companies that caused it can find out pretty quickly. This sort of democratic feedback is a fine thing to have in a globalized world.

A global grid can bring economies of scale

A global grid for powering homes, cars, trains, cities, and industry generally can bring economies of scale that outweigh the political risks, for example the risk of deals coming unstuck with corrupt or unreliable regimes, which seem all too common in the energy business. It can do so most emphatically under the assumption that we need to take due account of global climate developments as we build out the grid. For in that case, we shall need to harvest solar power where and when the sun shines, wind power where and when the wind blows, and hydro power where the water flows or when waves and tides can be tapped for power. And then we'll have to convey the electricity to where it's needed.

We should plan to build a huge global infrastructure of giant direct-current power cables. Perhaps we'll master the art of building

economically viable superconducting cables for megascale (thousands of kilometers) power distribution. In either case, in a globalized world, national boundaries will be little more than excuses to impose taxes and licensing restrictions, which is a polite way of saying they'll be fronts for protection rackets. Progressive people everywhere will begin to see the sense of overthrowing the nationalists wherever they pop up their ugly heads and in place of their rackets erecting a global legal and fiscal environment.

Many of us see good reasons to encourage widespread investment in nuclear generating facilities. With modern technology for inherently safe fuel cycles and redundant backup safety and security systems, the technical challenges here have become relatively easy to overcome. Even the waste problem is soluble. New reactors that run hot can burn almost all of it, and total world volumes of the rest are tiny. The main challenge is political. First, we need to work harder to sell the idea to skeptical consumers who still have nightmares of nuclear terrorism and reactor meltdowns. Second, we need to be sure we can neutralize the threat of rogue regimes breaking out of the regulatory framework to threaten neighbors. The 2010 issues with Iran and North Korea are useful test cases here.

If we can meet the political challenge posed by nuclear power, nations or regions can become more self-sufficient in the power sector. If someone needs more power, they simply buy or build a new reactor. This may be easier than negotiating carbon credits or signing a deal for gas or solar power with an unreliable neighbor.

The global regulatory framework for nuclear power will be as restrictive as the framework for eco-friendly power. We'll find it takes only take a few more cases of non-cooperation, blackmail, piracy, or terrorism to sour the prospects for the whole sector as badly as they were soured by the threat of nuclear annihilation during the Cold War or by the Chernobyl meltdown in 1986. In a nuclear future, we'll quickly see the need for regulation drastic enough to head off any such threats in advance. This will probably be enough to strangle any nascent homegrown nuclear power industries outside a select club of prosperous and stable states. But poor and unstable states have more urgent things to do than grow their own nuclear technology anyway. The whole issue is entangled with that of nuclear weapons. We'll come back to those later.

On the topic of power generation, a longer-term possibility is the prospect of cheap power from giant fusion reactors based on the "tokamak" design of the International Thermonuclear Experimental Reactor (ITER). In such reactors, heat for powering turbines is generated by the nuclear fusion of ionized isotopes of hydrogen and lithium into stable helium nuclei. Once such tokamak reactors become viable, perhaps in hybrid form where fission reactions in the walls sustain the fusion reaction and so let us build smaller installations, they'll be safer and cleaner than traditional fission reactors and will essentially solve all our large-scale energy problems for centuries to come. Even the downside is an upside. They'll always be big and expensive, and the only way to deploy them globally will be as part of a single technical infrastructure run globally. The worldwide cooperation of physicists working with the CERN Large Hadron Collider is a good precedent here. So a tokamak future will be a globalized future in which rogue states and suicidal jihadists have been exterminated and forgotten.

Cars and motorbikes will invade the bike space

Let's return to the humble topic of cars and their technical horizons. One dream from my youth seems to have evaporated – mass rollout of flying cars is no longer an option. Air traffic problems we could have solved with enough computer power, but the world fuel burn would terrify greens and the flying bomb problem would terrorize the rest of us. So bye-bye to that idea. To a first approximation, the planet's airspace will be reserved for airlines, corporate jets, and military drones (on which more later).

Cars will look more and more like survival capsules. In some cases, cars may be little more than wheelchairs with built-in life support systems. These would support owners who were unsteady on foot and would be small and quiet enough for indoor use. As the proportion of elderly people in the world increases, the number of people who'd like to drive single-seat cars that are small and clean enough to cruise shopping malls and pedestrian spaces will be high enough to justify big investment and targeted marketing. The cars for this sector could be as much fun as you like but could also pack as much medical gear as an intensive-care station.

What about bikes? Whenever people are jammed together, bicycles are a good way to get around. Yet the technology they incorporate is often dismayingly primitive. Compared with cars and mobile phones, bikes look like holdovers from a century ago. Even the supposedly hi-tech models with absurdly many gears and carbon-fiber frames are just aching, in my engineering view, for a massive makeover. Why is all that greasy gear-and-chain "greeble" exposed to mud and grit? Why do wires and cables hang loose on the frame? Why are the tires so easy to puncture? Anyone who can market a bike with a sealed and stepless power train, properly integrated electrical and electronic subsystems, puncture-proof wheel-rim technology, and good anti-theft technology can make millions.

But already the wish list shows the problem. Cheap and robust devices can't offer much in the way of techno-bling. Any progress here is bound to be disappointing. Single-seat cars and motorbikes will invade the bike space. People will pay more for cars or motorbikes, even if they're less healthy.

So cars will evolve in various directions, for mall cruising by wobbly-legged oldsters, for rail-gun tunneling by goal-oriented city commuters, for treks into the wilderness by rich adventurers, and so on. Many cars will be one-seaters, more will be two-seaters, and most will be either pure electric or hybrid, depending on their envisaged usage scenarios.

Cars may become give-away items

The consumer market for cars will fragment. But since the underlying technologies will be standard, the economies of scale will only get more convincing over the years. So the number of carmakers world-wide will remain small, and dealership networks will likewise become even more globally integrated than today.

In contrast, the global infrastructure for supporting cars, from street systems and freeways, to smart traffic-light networks and roadside spy robots, to drive-on trains and rail-gun tunnel systems, to robot park-houses and car-hardening stations for nuclear apocalypse scenarios, will ramify endlessly. Soon the infrastructure will be a lot more complex and technically interesting than the cars themselves.

They'll just be the cheapo pods we use to project our bodies between nodes in the network as smoothly as possible.

Cars may even become give-away items, like some mobile phones today, to people who sign up for a power and network deal. Israeli entrepreneur Shai Agassi (who as an SAP director was my boss for a while at SAP) piloted this idea in 2007 with his California-based startup Better Place, which aims to become a global provider of electric vehicle networks and services. In cooperation with Renault and Nissan, Better Place will roll out electric cars in Israel and Denmark (two countries small enough for limited-range cars) together with nationwide infrastructures of stations for swapping and re-charging their battery packs using robots. In early 2010, the company was valued at well over a billion dollars.

The Better Place scenario is the natural way to shape the future for occasional or uncommitted drivers. But even highly personalized cars with faithful avatars that respond only to their owners' touch (or avatar) will be cheaper than the lifetime rental for the networks they zip around in. Buying a car will be signing up for membership in a mobility club, complete with park-house and rail-gun deals and automatic overnight updates for the car software.

So cars will become bod pods for zipping around in the shared infrastructure. The century-long love affair will be over. Instead, we shall learn to love our robots. They'll do anything for us.

0011
Virtual Bodies

The traditional Christian doctrine of resurrection, according to which following the Second Coming believers arise in "perfect bodies" to enjoy eternal lives in paradise, is an intriguing psychotrip. We all hanker after perfect bodies, defined either in our own weird ways or in shameless conformity to prevailing celebrity fashion, and the idea of enjoying them for eternity, or at least a long time, is naturally attractive for human apes like us, who survive for decades beyond their years of relative perfection.

This hankering is the pull that makes robots the technology of the century. We're not developing robots just so that we can build cars more cheaply or vacuum carpets more easily. We hold out the hope of building better bodies that we can somehow inhabit one day. This is a long historical process that will still be very crudely realized at the end of the century, but by then we'll probably agree that we've made huge progress in that direction.

For decades yet, the enabling technologies will be the main focus of attention. Avatar technology will enable us to do better what we're beginning to do already, namely to project salient aspects of ourselves into machine environments, for example to do business online, to tele-operate simple robots for taking on dangerous or dirty manual jobs, or to fly aircraft and spacecraft. In all these cases, we'll have an online presence that others can detect and recognize as representing our own persons. Once macho AR headsets start replacing wimp screens and joysticks, they'll probably become the default tools to steer avatars in remote or virtual embodiments. And our cars, in their pod-like future incarnations, will sometimes be natural places to sit while we do so.

There are two main development strands to pursue here. First, we need to get better at controlling remote hardware through our avatars. It seems likely to me that we'll do this mainly by controlling virtual objects through our avatars, as in present gaming software. The transition from virtual to real can be as seamless as we like from the inside, so that we need hardly know whether we're shooting up virtual phantoms in a bloody game or committing mass murder in an unlucky

village somewhere. Second, we need to build better robotic hardware, with something approaching both the flexibility and the sensitivity that we've come to expect from our natural-born wetware. The hardware side of the story is a big topic, too much for now, so let's come back to it in a later chapter.

The development of better ways to control virtual bodies acting in virtual environments is likely to be a big theme for the next decade or two. This is because it's interesting, even fun, and obviously feasible, both technically and economically, since we've begun to do it already. The games industry has begun to outgross the movie industry for a good reason. Instead of sitting passively in a seat and watching actors on a screen have all the fun, the gamer is in the action, doing stuff. That has to be incomparably better, even if the action is crass and stereotyped, as it is when you're just shooting zombies or driving fast around a track.

I predict that in the next few years the game format will begin an irresistible creeping replacement of the movie format for increasingly treasured cultural artifacts, until we see markets for game versions of literary classics in which the gamer can inhabit any character in the action and trash the original plot in a variety of amusing and possibly instructive ways. To preserve a modicum of fidelity and decency here, the original plot and characters will be strongly favored or proposed by the game software, but the fun for the gamer lies precisely in tweaking the original and seeing how well or badly the tweaked version turns out. Watching a great game of chess is one thing, but learning to play chess yourself until you can achieve your own feeble parodies of the great games you've watched is incomparably more satisfying. We all want to be stars in our own action movies.

If this development resembles an abominable descent to the lowest common denominator, where Hamlet can grab a gun and shoot everyone in sight or Madame Bovary can become a pole dancer, so be it. Our cultural icons can become untouchable idols, where they soon fall into practical irrelevance, or they can be recycled and woven into the living fabric of new developments. I find the prospect of refactoring all the cultural icons of past centuries as games, where my avatar can step in and enjoy the fun, rather inspiring. For example, I'd like a game where I could interview great philosophical sages like Plato or Kant, where the software reconstitutes them in photorealistic detail and equips

them with word-perfect dialectical responses, so that I can ask my dumb questions and get tailored answers from the virtual masters. It would be better for grasping the seminal ideas of philosophy than reading the great books and then talking live with bored professors or overenthusiastic teaching assistants who try to put their own spin on the ideas. Anyway, it would be a lot of fun writing the software and packaging and marketing the games.

This development will dominate the educational and publishing sectors for a few decades yet. When movies first became a technically mature medium for artistic content, there was a huge employment opportunity for creative people to refactor all the classics of our civilization, previously available only in book form, for the big screen. That huge labor has now been accomplished. Only a purist would now insist on reading a thousand books to understand our civilization when you could watch a thousand movies and documentary videos instead. No modern scientist feels the need to read Isaac Newton's *Philosophiae Naturalis Principia Mathematica* to understand classical mechanics. Modern textbooks offer much more age-appropriate presentations, where the problem sets are almost like games already and only need tricking out in the online edition with some good graphics and a scoreboard. Similarly, it seems to me, no sensible person should feel ashamed at not having read Shakespeare when the movie and video adaptations are often so good. The same again will go for the game editions of all the classics, so long as we make the effort to transpose the key ideas with sensitivity to what we decide was the original intent. In a marketplace of game versus game, the best games win, even if some of them are egregiously overstuffed with sex and violence.

Games define the semantics of our natural languages

The concept of a game is a lot more elastic than you might think. In his later years, Ludwig Wittgenstein, the philosophical pied piper of twentieth-century British intellectual life, wrote a lot of oracular notes about games that reveal rather deep insight into their role in our lives.

Since Wittgenstein reappears in my story, let's take the time to introduce him. Born in 1889, he grew up in an extremely wealthy but absurdly dysfunctional family in Vienna. As a young engineering

postgrad in Manchester, he took a shine to what was then the hot new field of mathematical logic. In 1911, a year when Bertrand Russell was working with Alfred North Whitehead on their classic trilogy *Principia Mathematica* in Cambridge, the young Wittgenstein burst in unannounced on Russell and passionately expounded his views. Russell decided the young man was a genius. During the First World War, Wittgenstein served as a soldier in Italy and wrote philosophical thoughts in pencil in a notebook. Russell was dazzled, the thoughts were published in 1922 under the title *Tractatus Logico-Philosophicus*, and for a few years Wittgenstein endured what became known as an "epistemological break." He returned to Cambridge and a fellowship, where he devoted his remaining years to writing a large number of scrappy notes on philosophy that are now revered as scripture. The books that were later edited from them are of continuing interest to philosophers. Many of my teachers at Oxford were in thrall to them and wrote secondary literature on them.

The late Wittgenstein focused mainly on language games, which serve to define in an interactive and dynamic way the semantics of the natural languages we speak. The core idea is that games are interactions with rules that don't always fix everything. Whatever isn't fixed can be changed in the course of playing the game, whether deliberately or not, in ways that may be creative extensions of whatever the game was originally about. This idea has the interesting implication that the elaborate mathematical discipline of game theory is built on shaky foundations. Two or more players follow the rules and seek to maximize or minimize certain defined things, with the aim of winning in some way defined by the game. The philosophical problem is that what counts as following the rules and winning or losing is not only generally not very well defined but quite possibly not defined at all. As well as discombobulating game theory, this problem implies that refactoring all the treasures of our civilization from books and movies to games is actually much harder than it might seem.

An example can make the difficulty quite obvious. Imagine you want to make a game out of Plato's philosophy. Plato wrote a lot of dialogs starring his teacher Socrates, where Socrates argued with a series of more or less attractive young men in an attempt to get to the truth, or at least something closer to the truth, about whatever it was in each dialog they were arguing about. The dumb approach here

would be to translate their ancient Greek into modern English and simply dramatize the dialogs as movies. But modern audiences would find most of them pretty boring, so it's hardly worth the effort of making such movies. No-one in their right mind would want to sit through them all, for the simple reason that anyone smart enough to find the topics interesting would want to interject with their own opinion again and again. It would be just infuriating to sit through such stuff for hours on end and not be able to wade in and protest at regular intervals.

This is where the game format shines. You can intervene! You can stop Socrates or his silly young men in mid-sentence and ridicule the nonsense they're spouting. Then you can utter your own pearls of wisdom and see how Socrates likes it. You can force Socrates to debate you, in person, and try to convince you, the hero, that he was smarter than you and you don't always know best. That would be a game worth a few hours of any serious student's time.

But notice how far we've come from the original scenario. What's the point of the game? How do you know when you've won? Perhaps the point is to try to outsmart Socrates, and you win a big fat star every time he's stuck for an answer or concedes that you might be right. Whatever it may be, I don't know about anyone else, but I'd be hard pressed to program the dialogs for this game and I'm pretty sure that plenty of Plato scholars would dispute any script for the game that I came up with. So we have a game where the rules and the moves are dubious and the point is unclear. Brilliant! Sadly, this is typical for such games.

However, this limitation of principle is no problem at all in practice. Think of all those cheap games for kids where you get to refight big battles of the Second World War. To create the game, all the programmers have to do is fashion some snappy graphics for tanks and aircraft with authentic surface greeble and insert them among historically plausible city scenery and natural landscapes together with bunches of humanoid zombies and gamer avatars in military uniforms. The order of battle, the rules of engagement, and the criteria for winning can all be generic and relatively inauthentic. Any kid playing this game will get a bizarre sense of how the battle was fought and what it was all about, but it doesn't matter much. Compare it with my sense of those battles, built up over decades of reading books

about the war by historians and leading participants as well as years as a kid of building hundreds of plastic models of the more glamorous weapon systems used in the war. Anyone who fought in the war will rightly regard my own second-hand understanding of the fighting as strange and somehow unreal. You can't judge the effectiveness of Blitzkrieg tactics by building a plastic model of a Stuka dive-bomber. It seems obvious to me that the modern gamer who shoots up Tiger tanks on his television screen is only a few steps further from the reality than I was or am.

The important thing in all such games is to give them a plausible background. If the initial scenarios and the screen animations are well researched, it doesn't really matter that the game rules and moves are generic and the winning states arbitrary. Just as in a historical novel, so long as the background is well researched, readers are happy to forgive anachronistic dialog and unlikely storylines. In movies, we even accept the same actors in absurdly contrasting roles, such as Charlton Heston as Moses in one movie and as a horny old scientist in another. All these absurd aspects of the media we use are forgiven as transparent artifacts or as harmless functional parts of a vehicle for delivering historic content.

So we don't care that our games are artificial. We do care that they grip us for long enough to let us absorb the fruits of all that meticulous research into the background. Another movie example helps: *Avatar* has what critics deride as a formulaic plot that's almost childishly simplistic, but we don't care. The whole scenario works on us anyway. Most viewers are happy that the plot lets rip with the emotions and doesn't keep forcing them to stop and think about murky issues that would merely distract them from luxuriating in the deliciously realized ecosystem on the planet Pandora.

Let's be more sophisticated about games here. They don't have to involve shooting zombies or award points for racking up the body count. They don't even have to have win-or-lose scenarios. They can just be places where you do what you do alongside other people doing much the same. For myself, I'd want a philosophy game featuring face time with Socrates to be like that, with no scores and no winners, just the opportunity to argue with a wise old dude.

To take a very different example, a game about, say, the Holocaust, the Shoah, would be in almost blasphemously bad taste if you could

score points for saving some Jews or killing a few Nazi guards. The point of that game could only be to make you feel for a moment a shadow of what it might be like to be incarcerated in a death camp, and for that purpose a game that started with your being arrested on the street by ugly soldiers and ended inexorably with your body going up in smoke in an oven somewhere, whatever you did in the game, would be better.

So the concept of an electronic or online game is wide and elastic. You can even play games where you make up some of the rules as you go along, or redefine what it means to win, in a gamer version of the idea that all's fair in love and war.

A game puts you into a mindworld

Your action is what makes a game different from a book or movie. The game has a locus of action that you can inhabit and steer. You're in the game and you make choices. With a book you're floating in undefined state, godlike in its majesty, eternal in its timeless freedom, yet also completely cut off from the action, which happens in its own ideal space regardless of whether you read the book or not. With a movie, you're dipping at least an eyeball into the river of events that make up the action. You have a point of view defined for you and your job is to inhabit it, to be a witness to the tightly choreographed sequence of events that come to pass around you.

In a game this empowerment goes further. You have a presence as an avatar with a virtual body, and you can decide to go this way or that at various points in the action. Things happen to your avatar that you may or may not like, so the game can push you and jog your emotions in a more targeted way than a movie can. The result of gaming through this embodied presence is that you feel immersed in the virtual world. A book world is like a planet in space that you observe through a magic telescope for making images out of words. A movie world is like a planet that you zoom into or onto as you follow a preplanned trajectory that pins you to your seat. A game world is like a *Star Trek* planet where you beam down and interact with the natives before beaming up again later.

I find it natural to sum this up by saying that a game puts you into a mindworld. Since *Mindworlds* was the title of my 2009 philosophy

book, this may seem like a shameless plug, but for me there's a deeper point here. A mindworld is a world with a subject in it to reflect it and bring it to a unity. I won't bore you with the (totally terrifying) technical trivia, but the contrast with a classical world is that a mindworld has a re-entrant point or a navel, where the world can appear to itself and thus be brought to self-realization, whereas a classical world doesn't. The subject of a classical world is outside it, undefined and free-floating, and the world is at it is, eternal and unchanging. One of the big lessons of modern physics is that the classical world is at best half a world. When you add the other half, as a subject located in time and experiencing the world from inside, with the ability to act in order to shape the future, you get a mindworld.

In short, a book lets you spy on a classical world and a movie either puts you in a classical world or flies you past a mindworld, whereas a game beams you down into a mindworld.

We'll use avatars to steer machines in the real world

We can leave it there with my homespun metaphysics. Let's get back to robots and how we can begin to transform them into something like our perfect bodies. As games become the default vehicles for our cultural crown jewels, we'll get used to acting through avatars. As we learn to use such avatars in turn to steer real, physical machines, we'll get used to exerting telekinetic effects in the real world through our virtual mindworlds. We don't need to think too hard to realize that acting through our avatars and putting our avatars in robots is equivalent to reincarnating ourselves as robots.

At first, the incarnations will be limited and temporary. The robots we inhabit will be about as good for the soul to live in as a can opener. Robots are just crude machines, after all. But this will change as picotech improves the robots and as bionic interfaces shorten the command chain. Within a few decades, we'll often be happier in our bot bods than in the rotting human heaps we left slouching in the VR pods.

The big prerequisite for this future is that we're happy to be our avatars. We need to accept that the digitization of our minds we undergo to map our thoughts into a mindworld is reasonably complete. This will naturally be a matter of opinion. If I want to, I can assert that

everything of any consequence about me is captured in my words, perhaps in my published books, and therefore that the survival of my literary corpus makes me immortal. This requires some hubris on my part, given my meager output of consequential words, but there's no contradiction in my asserting this. But when a person has a substantial repository of self-generated digital content, the claim can be altogether more convincing. It's not immodest to claim that everything worth preserving of me is in the digital data dump that includes records of all my writings, including emails and tweets, all photos and videos of me, all my phone calls and public speeches and interviews, all the things I've read or looked at online, and so on. Apart from that lot, I'm just a genome, various scrambled parts of which I may have passed onto my kids, and a bank balance, which does me no good once I'm gone. So in principle, a well-stocked avatar says it all. We're a long way from such fat avatars, but the distance shrinks with every day that passes.

Avatars don't need to be bulky or fat. In general, they'll be as thin as possible. For example, the avatars that animate the games we play are very thin. If my avatar in an arcade game is a U.S. marine with Popeye forearms cradling an M16 assault gun, it makes no difference that my fat avatar at home includes a huge data file for Kant's critique of pure reason. Any mindworld is going to take only what it needs from my big fat avatar and dump the rest. Anyway, my fat avatar will only be virtually bulky. Most of the stuff in it, most of the fat, will be in the cloud, delocalized and stored redundantly in arbitrary nodes across the net.

Much of the stuff that my fat avatar touches will somehow be shared with other avatars too. Think, for example, of books that you and I have both read or security-cam videos we both appear in. The cloud will store huge blobs (binary large objects) of content that's referenced by multiple fat avatars, but that stuff doesn't need to be stored anew for each avatar for the simple reason that any exotic apps that access all that deep dirt do so only once in a blue moon. It's the same principle that banks use when they save your money. They lend it out again and again to earn interest for their next big bonuses in the smug knowledge that it's most unlikely that all their customers will want their money back at once.

To balance this, we also need to be clear about how thin our biological minds are in everyday life. The contents of consciousness

from moment to moment are mostly far, far less than the mountains of stuff buried deep in our memory vaults. My thin mind, the one that helps me take showers in the morning and do the dishes after eating, is pretty innocent of philosophy, though generally ready to engage in philosophical banter at a word's notice, and certainly can't be said to be an intellectual heavyweight. But my fat mind, the one that lurks below the conscious surface, is a monster of erudition, both actual and virtual, that includes a swathe of references across the entire corpus of Western civilization.

So personal identity is thin too, like avatar identity, and where personal identity gets fat it tends to lose integrity and merge into a shared planetary blob. The digital metaphor suggests strongly to me that the deep basements of our minds are shared in a labyrinth of "mental" content that's just the mindworld flip side of a lot of physical stuff that we don't even know is part of us until it's gone and we're caught short. But again, let's skip the metaphysics and move on.

The take-away message here is that we shall increasingly let our avatars stand in for us. As the art of building avatars to reflect our most essential or desired qualities improves, a variety of thin and fat avatars will serve our purposes from moment to moment. Such avatars will be our logical interfaces for all online services, for all smart machines from robot spacecraft to refrigerators, and for all mediated communication with other people. We shall become artists at shaping, selecting, and manipulating our avatars to project the presence we want, and we shall become connoisseurs of the avatars that other people use to interface with us. Many of these avatars will interact in game environments where virtual scenery offers the handles we need to do what we want. In particular, such games will be the default vehicles for most of the cultural content we need to assimilate to become respectable members of our communities. And the games industry will morph out of all recognition as it expands to do justice to this new role.

The virtual bodies we give our avatars will at first be surfaces

What shall we do with our biological bodies? Are we going to treat them like faithful old dogs that can no longer compete with new attractions for the kids' favors? Are we going to leave them to fart and

stink and leave hairs everywhere in the VR pod as we take virtual leave of them to enjoy fun and games in the cloud? I think not, for the convincing reason that they're our biological roots. Without them we're dead. The only way to change this inconvenient truth is to re-implement our entire consciousness in the cloud, and that's a technical challenge that needs more than a few decades of Moore's law to overcome. Our bodies are continuous with our minds. Feelings from deep in our guts (or our backs) have decisive effects on how we think and act, and no amount of virtual overlay in our conscious minds is going to change that fact of rootedness in the near future.

The virtual bodies we give our avatars will at first be mere surfaces. Doubtless the art of modeling convincing detail will continue to astonish us and tempt us to identify ever more fully with their multi-sensory reality. But you'll still feel it if someone sticks a pin in the old biobag back in the pod. Again and again we'll get movements to re-identify with our "God-given" bodies and leave the VR pods to get some fresh air and exercise, and each time the pod-dwellers will emerge to treat themselves to stints of jogging and swimming, and maybe even dare to get a tan, before slinking back to the pod for another fix of cloud life. Assuming that life in the physical world goes on and our all-electric homes still hum, we're likely to get sucked more and more into the cloud.

This sucker-punch is like the rapture that fundamentalist Christians hold out as the reward for true believers when the apocalypse erupts to engulf the damned. The unfortunates who can't get their cloud act together will be left behind to rot in their bio-bodies. The relentless processes of biology will consume them while their cloud-dwelling neighbors savor the delights of paradise. We shall crave the rapture of out-of-body experiences from time to time just as keenly as we crave in-body experiences to remind us of our roots.

The brain is a computer and the mind is an avatar

Our physical human bodies are fine machines. They've been sculpted by evolution down to the molecular level to work as a survival suit for the soul in all the natural environments that planet Earth had to offer in our uncivilized past. But as molecular machines they have the big drawback that they demand a lot of increasingly expensive care and

attention. In a world of diminishing resources they're luxury items, like the abdominal six-pack on a lovingly sculpted bodybuilder torso that betrays many hours of pumping iron in a fitness studio.

Consider the effort we need to put into fueling our bodies to ensure we get all the trace nutrients they need. We don't even know what they need! We just trust that stuffing ourselves with more or less natural foodstuffs to meet the main needs will automatically supply the rest, more or less, and that evolution must have equipped us to survive and flourish on that basis. As it happens, this seems to work, but we have no good reason to expect this state of natural serendipity to continue as we keep fiddling with our food chain.

Our technology is changing our body image. As a species, we have a habit of using our latest technology as a metaphor for how the body works. For example, now that we have machines that generate virtual realities for us, the neuroscientists among us imagine the brain as a VR generator that creates a mindworld from the data sent streaming in by the sense organs that festoon the body surfaces. In this metaphor, our body image is a natural avatar, no less, which forms the virtual center of the mindworlds we inhabit from day to day. Because we naïve humans know no better, we take this natural avatar as real, as the real me, and treat it as the gold standard for all the techno-avatars we create for gaming and cloud life.

Following the scientists, we think our brains are like hardware and our minds are like software. We smugly remark that the brain is too wet and squidgy to work as fast as our best hard boxes and anyway has a spaghetti jungle for a wiring diagram, but we don't think this undermines the metaphor. In principle, the vision of brain as computer and mind as avatar is logically no better than any other machine metaphor, such as that of a heart as a hydraulic pump. But in fact, amazingly perhaps, it helps, and for what it's worth I think it's quite a good picture.

The acid test of any such metaphor is what we can do with it. Does it help us solve hard practical problems? These are early days for the idea that the brain is a natural avatar generator, but I'm prepared to stick my neck out here and say I think it works. It makes sense of all the philosophy I've struggled with to try to understand the brain-mind issue.

As an animal species, we've evolved brains that generate a specific kind of avatar within a specific kind of mindworld to help us survive in prehistoric hominid environments. Our mindworlds tend to be three-dimensional and vaguely Euclidean, colored in the spectral range from red to violet, sized over the scale from a little below a millimeter to a few kilometers, and so on. Our avatars tend to focus on skin color and texture, hair distribution, bodily symmetry, sex organs, muscle mass or fat distribution, and so on. Our behavior patterns – the games we play – tend to revolve around food, grooming, respect, sexual dominance, and so on. All these are biological givens that seem real to us. They're where we start. Because the narratives we construct on this basis have worked well for us in the past, we tend to think that this is the ground truth about us. But as soon as we start creating artificial variations on this natural theme, we see the artifice of the whole story. My example of the season is the movie *Avatar*, where a different ecosystem is presented with enough plausibility to seem quite real in its own way. It won't be long, I predict, before that ecosystem seems quite tame and familiar compared with what we can conjure up using yet more computing power.

This idea that the brain is a VR generator and that we take the avatar it generates as the ground truth about ourselves is expertly worked out in the writings of my friend Thomas Metzinger. As a respectable German philosophy professor at the University of Mainz and a founder member of the Association for the Scientific Study of Consciousness (ASSC), he advocates this view as the basis for a serious theory of mind along Kantian lines (except that the words "avatar" and "mindworld" are my contribution). In fact, I thought out most of my ideas independently, but naturally I'm delighted that our views cohere so well together. We shouldn't be surprised here, since this view is simply the logical outcome of the human tendency to understand ourselves in terms of our latest technology.

But in this case there's more to the idea. Like many anatomical metaphors, such as the heart as a pump or the eye as a camera, there comes a point where we hit the bedrock scientific truth and all the details work out correctly. That's where we are with the brain as avatar generator, in my humble opinion. We've hit the nail on the head, and the rest of human history is just working out the consequences.

We'll enjoy our natural avatars as virtual bodies in virtual worlds

Let's get back to virtual bodies. There's now a lot of good scientific work showing that the brain's body image can be disturbed from its inbuilt default settings in many strange ways, either by the accidents of personal development or by injury to the brain. But still there's a default, a "perfect body" image that can be degraded or corrupted. And naturally our real physical bodies can fall short of the "perfect" avatar image in many ways. Indeed, the general course of a human life is to find that the accumulated hard knocks and injuries of an active life, together with the usual effects of advancing old age, tend to push the physical body further and further away from the ideal image embodied in the brain's natural avatar. This is where that Christian yearning for a return to the "perfect body" comes from. If we could wind back the clock and undo the effects of all those hard knocks, we'd be in synch with our avatar again and all would be well in the best of all possible worlds.

Technology rides to the rescue. With advancing medical science and hi-tech brain scanners, we'll soon be able to reconstruct your natural avatar so well in software that we can recreate it as a virtual body in a virtual world. Imagine the joy as you settle into your VR pod and experience your body exactly as your brain would most like it to be. Everything in the simulation is tuned to your personal brain map. You skin is perfect, your fat distribution is ideal, your muscles are well toned, your reflexes are sharp, and you feel full of energy.

Now we take the simulation a step further. We port the avatar into your AR set and you step out of the pod. The real world is as real as ever, but you look down through your iGlasses and see that your body is perfect! So far, so good, but of course no headset can make up entirely for the extra pounds, the flabby muscles, the aching back, and so on through your personal catalog of ailments, so as soon as you try to do real things in the real world the problems will be back. Your new avatar is no more able to transform the hard facts about your life than a diet of Prozac (Lilly trade name for fluoxetine) – although like Prozac it just might help you find your balance again.

Perhaps it's obvious by now how strong the pull will be toward living more and more of our lives in the cloud, playing online games in dream bodies even as we stumble through real reality in our magic

glasses. In our cloud worlds we can live out our wildest dreams and feel the joys of spring every day. Sadly, someone has to stay behind in the physical world and keep the machines running. That hard fact forces a hard limit on how far we can rapture ourselves into the cloud.

I guess we'll take turns. Each of us will be required to spend a few hours a day in the real world, looking after the bodily functions of our dependents and ourselves, checking that the bills are paid, doing the dishes and the vacuuming, and generally checking that all's well. Naturally, as we do all this, we'll be as augmented by our headsets as possible, then just as soon as we're done we'll race back to our pods for another blast of the high life in our designer bodies in the cloud. Some people will hog more pod life for themselves than others, just as some people today hog more money or more time lazing in the sun. Perhaps government health and safety departments will lay down the law here. And perhaps they'll make exceptions for old and terminally ill people, who may be allowed to drift away in youthful avatars in cloud heaven.

I find this an unconvincing vision of the future. People will soon tire of living with their heads in the cloud when their cloud sessions are punctuated by regular returns to the drab, crass, dull world of human reality. Anyone with a zealous desire to change the world will see a huge opportunity for improvement here. The opportunity is to bring the cloud worlds down to earth and realize our dream avatars in picotech bionic flesh. No fantasy, even when it's tricked out in the latest VR or AR technology, can substitute for what I feel to be real in my most lucid waking states. We need to change the world at that level. Once we've started doing so, we'll see that the dreams are strictly secondary. They're just motivators for action in the RR world of poverty, stupidity, filth, pain, and death. The future belongs to people who grapple with the RR world. Once they succeed, they'll push the AR cripples and the VR junkies mercilessly aside.

The sci-fi goal is to leave the world of flesh and blood behind

So how does this future vision pan out? How do we make smart use of the technology of avatars and mindworlds to change the RR world? The answer is to do it in stages.

First, we implement new ways using AR and VR technology to inform ourselves much more effectively than hitherto about what's possible and what's desirable in the way of improving the state of the RR world. This is a big opportunity for the education and publishing sectors, where no-nonsense entrepreneurs can refactor all our cultural crown jewels as cloud games. The main open challenge here, given that usable technology is largely available now and needs only a new wave of people to extend and deploy it, is to go straight for the low-hanging fruit in order to make an effective start. Once people begin to see how any content, from car mechanics to the Bible, can be repackaged for easy consumption in standardized game formats, the revolution can get rolling.

Second, we build out robot technology as fast as we can. The primitive robots we have now will soon become obsolete. Better ones will replace them. Each new generation of robots will be more capable, and we'll learn to task robots to do more and more of what was once human work. But autonomous robots will soon hit a natural limit that's independent of any technical issues. We won't want them! We'll soon learn to fear robots that relentlessly do what they're programmed to do without any input from us. We'll insist that they be answerable to us and that we be able to stop them any time we like, and even that they should need our encouragement from time to time to continue with their programmed tasks. Naturally, we won't always want to do this in primitive ways, by shouting at them or grabbing their arms. Instead, we'll use avatar technology. All but the simplest robots will be required to have virtual docking stations for our avatars, so that we can step in, virtually, and override them whenever we like. Then people in power will build robots incorporating inhibitors for the override functions so as to require authorization for an avatar takeover, and so on. A hi-tech arms race will follow, which will have the effect of evolving ever more intricate and hair-trigger human overrides for robot action sequences.

Third, we develop bionic interfaces for robots. A bionic interface connects directly with our biological substrate as human apes to enable us to control robots intuitively, without the indirection of data-driven wimp or macho gadgets. Imagine a scenario where you have an avatar in the cloud that you use to boss a team of robots around. Each robot has its own rather basic avatar, too, with a hole in it, so to speak,

that your boss avatar can dock with. Try not to think of sex as you contemplate this procedure. Fun as it may be for you as boss to dock with your bots over a datalink, that sort of chain of command is far too indirect to let you inhabit your very own personal body-bot with the desired intimacy.

Ultimately, you want your human body to be spliced at the neural level to bionic hardware. Imagine the hardware ending in a brushlike array of thousands of control threads, each tipped with a microscopic cup where a cut nerve end can nestle and flourish like a tiny potted plant. Then imagine having your arms and legs cut off and replaced by the bionic extremities, with all your nerve ends spliced and potted correctly by a team of nanobots. Naturally, you'd want the new extremities to be as intimately yours as the human limbs you gave up, and to feel that swapping your meaty parts for state-of-the-art bionic products was pure gain.

The long-term future for this approach is that people work and save over years to replace more and more of their human bodies with improved bionic parts, from arms and legs to hearts and livers and guts, until almost all of them is reconstituted as smart bionic borgware in a titanium shell and only the brain remains in its sloppy native state, linked to a clumsy backpack that pumps synthetic blood through the soggy mass in the helmet unit. Then, in a final throw, the happy customer replaces the brain, too, by porting the self-avatar to a picotech hypercomputer the size of a grapefruit.

As you see, the science-fiction goal here is to go the whole hog and leave the old world of flesh and blood behind. But the technological long march to get there is far longer than a brief sketch can suggest. The whole world of biology needs to be assimilated into an embracing world of molecular technology, so that we can build and rebuild protein-based organisms as freely as we now plug together Lego-brick houses, and that challenge will occupy us for many decades yet. In fact, many distinguished scientists regard that assimilation as the great scientific milestone of our present century. The consequences will be enormous. Reincarnation as bionic bots that we can augment and upgrade any way we like will cause a revolution unlike anything that came before in the history of life on Earth.

01
Global Mammon

A world of avatar frontends, hypervision headsets, robot cars, and reality games will cause a revolution as it unfolds it in all its glory. That revolution will start in business organization, where new software will support and streamline globalized business as never before. The revolution will ramify through mass deployment of robots, both in industry and in daily life. It will cause aftershocks in political systems worldwide as they adapt to the new realities. These adaptations will enable the burocrats to extend and deepen the revolution. These four themes animate the next four chapters.

First, cookie-cutter templates for running a business will get so friendly that we can all use them. Even people who work as corporate clones will run their homes and lives like small businesses. Second, robots will pop up everywhere. Armies of tireless machines will drive the old industrial proletariat to extinction. Third, the global reach of business will force politicians to globalize their politics. And fourth, economic realities will drive the rise of a global accounting currency. Together, these changes will create a wealth engine to draft workers worldwide into one big force for goods.

To keep your head among all these dizzying changes, look back as well as forward. Each new century of progress seems to leave the past in the dust, until it too passes. This century will trump even its predecessor for dramatic change, but much will still stay the same. The facts of science and geography will frame all our advances. Such limits will ramify down the stack of constraints on our lives.

Working back, we can frame the changes against the fixed backdrop. Then the rise of the robots looks less amazing and the demise of the human ape less shocking. Delegating war to the robots will surely come as sweet relief to people grown tired of bloodshed. Life will go on, and those of us who live it will count ourselves lucky to be alive in an age when economics becomes at last a solved problem.

0100
Starship Enterprise

The revolution in the history of life on Earth has a precondition. It can only happen if something like the present global capitalist system keeps on working for a few more decades. Since the system is driven by unbridled greed, this may come as dismaying news. The good news, if you can bear to share the banquet with gluttons, is that we shall fix the system in various ways to make its effects more tolerable to fastidious diners.

Most of us now agree that communism was a bad thing. This may not mean much when it's the politically correct opinion among the communities that won the Cold War. But the philosophical foundations of communism were shaky, since they relied on a reservoir of idealism in the masses of society that a disciplined ruling elite could simply co-opt to serve its own selfish interests. By contrast, the more capitalistic system we now enjoy is more robust. Layers of people each serving their own selfish interests end up creating a world where more people can more effectively help themselves. A precondition for its success is that a governing elite defines and preserves the arena within which the whole show plays out.

The desire to get rich – to earn cartloads of money and do whatever you want with it – is so easy to awaken in most people that it can become a central pillar of the social order. Once that desire is burning, the victims become easy targets for social engineering. The social engineers can rely on them to go the extra mile to earn the extra buck. So a society of income maximizers results, who dutifully pay taxes (with gritted teeth) and so on because it's the only way they know to stay cool with the law and keep earning the money they crave. Because money is a fungible asset, its owners can do more or less what they want with it, including all the things that for most people make life worth living, such as supporting more kids, living in a better house, becoming more influential and respected in the local community, and so on.

The governing elite that makes all this possible has to do its job right. Like any other game, the money game only works when people

obey the rules, which presupposes that someone defined the rules well enough to ensure that they work as intended. Sadly, the people who form governments and run big organizations tend to be as venal as anyone else, with the result that it's always a fight to iron out all the kinks and loopholes in the system and eliminate all the perks and privileges that the rulers would otherwise quietly pump to enrich themselves. The question of who governs the governors has been a big one in political theory since the days of Plato, who in his *Republic* proposed that "philosopher kings" were the only ones noble enough to do that dirty job and not get themselves too filthy in the process. As a wannabe philosopher king myself, I find Plato's proposal intriguing, indeed quite genial.

The acts of government perpetrated by a philosopher king (or queen, to make the obvious correction) need not be too burdensome. Basically, the job is to inspire or shame the governors and the bank directors and so on into cleaning up their own acts. One does this by expounding brilliant ideas, for example, that catch their imagination. Alternatively, one can work in stealth mode as an investigative journalist and reveal misdeeds or corruption so monstrous as to trigger the correction. Or, like me, one can lay down a vision, together with a few hints as to how to realize it, and let the vision take hold among a new generation of leaders.

Science and technology feed the fountain that delivers the goods

My vision starts out in the knowledge that science and technology are at the source of the fountain that delivers the goods. To continue to enjoy these goods, we only need to encourage the scientists and engineers to keep on working in something like our present economic system, which functions rather like a bucket chain to pour the fruits of their labors upon us. So it's the system of managed capitalism that actually delivers the goods, but it can only do so if the input of innovation is uninterrupted.

This needs some explaining. Without innovation, the flood of surplus value that we need to fund the social order dries up. The economy becomes a zero-sum game. It doesn't take long for fights between the players to squeeze out the "excess" liquidity. For example, union bosses see cash on the table and bid up the price of labor. Once the

liquidity is gone, the social order locks up and rigor mortis sets in. With innovation, by contrast, the social order is like a loaf rising in an oven, filled with little bubbles that give the lightness of being that many of us love about our present lifestyles. In practice, the little bubbles are lucky people getting unreasonably rich from their innovations, which may be either great ideas or just silly fads. If the flood is strong enough, a few fads along the way can be fun.

Running an economic system is a fine art, as I'm sure the big bankers of our time will readily agree. Those who manipulate the rules that keep things bubbling need a sensitive feeling for what they're doing. A central banker is like a master chef who knows exactly how to mix the ingredients to get the soufflé to rise as desired. Insofar as the other ranks in our society can see the sense of this arrangement, they may even let the bankers pay themselves fat bonuses, on the principle that someone, somewhere must enjoy the benefits of getting rich to prove that it's possible – and maybe, out of sheer magnanimity, help others achieve a similarly enviable state. Assuming general acceptance of the status quo, then, an economic system can work like a giant organism, in which the circulation of money resembles the circulation of blood and all the little cells in the organism get what they need by tapping the circulating liquidity.

This cartoon caricature of economic theory works in miniature for a commercial corporation. The company earns money and the employees get paid. Some employees do challenging work and therefore get paid more. The people who run the company may get paid unreasonably much, but this is usually a perilous and transient condition, and in a dynamic world of dog eat dog, unrest lower in the food chain will soon unseat a pilot who can't navigate the flood of innovation well enough to keep the blessings pouring down from above.

The grunts that do the hard work for low pay may not be inspired by the setup. But as long as they get more than they would out in the wilderness of the labor market they'll play along. An important priority in any organization that employs lots of people on low pay is to make a fuss about how valuable their contributions are to the success of the organization. No boss can get rich without workers to keep things running, and any boss is in this sense the servant of the more humble employees. Bosses are easy enough to find – any organization can throw up wannabe bosses just by posting a job vacancy –

but solid workers who keep on going when they see that the system is grinding them down are like diamond bearings that keep the wheels of industry turning. So it pays any company to celebrate the contributions of its more humble employees, and this is indeed what we see in most successful companies.

Corporations represent a stable social form

The point here for our big picture is that corporations with bosses and workers represent a stable social form. Here I must disclose that a nineteenth-century Oxford philosopher called Francis Herbert Bradley has influenced me. Let me explain. (For more, see his entry in the online Stanford Encyclopedia of Philosophy).

Bradley impressed me through his 1876 book *Ethical Studies*. Two essays in that book stand out. In "Pleasure for Pleasure's Sake" he argues against the utilitarian idea that our best aim is to get happy by maximizing our pleasure. He says we can't even know we're happy independently of how we get there. In "My Station and Its Duties" he outlines a social conception of the moral self in a Hegelian account of the moral life. In that account, our best aim is self-realization, and the self is fully realized by fulfilling its role in the social organism that grounds its duties.

Bradley expounded a classic Victorian conception of men who know their place in the social order and do their duty, come what may. Such men were the rocks in the foundation of the British Empire. In 1924, King George V awarded Bradley the Order of Merit.

Now, in these more enlightened times, we can still find inspiration in the idea that people will support a social order if their contribution is properly acknowledged, even if this means they must watch others greedily helping themselves to all the extras they can grab. Virtue can be its own reward, and just letting other people behave like pigs can be preferable to joining in the scramble for a favored place at the feeding trough. In this way, a corporation can remain stable despite great inequalities in the wealth and status it confers.

Commercial enterprises are prototypes for organizational forms more generally. Many of the corporate entities that help shape the environment we live in are poorly managed and barely functional, in part due to lack of a guiding discipline to enforce lean development.

This sad fact was the impulse behind Margaret Thatcher's drive to privatize previously nationalized industries in late twentieth-century Britain. The profit motive may seem a crude driver of human behavior, but it does work in an economy where money is sufficiently powerful to buy most of what can make a person's life more bearable and enjoyable. For this reason, I confidently predict that the organizational outlines of societal entities generally will increasingly resemble those of successful commercial companies. The discipline of money can safely be allowed to shape more of the societal architecture that surrounds us.

The discipline of money can transform traditional social forms

Looked at from afar, from a Martian perspective perhaps, human civilization can with some plausibility be seen as an ongoing process of monetarizing more and more human intercourse. The discipline of money, and the opportunities for rationalization that appear when it acts, can transform traditional social forms. Monetary tokens such as coins and treasury bonds serve as temporary repositories of value, and value is by definition something that humans want or go for. We attach value to our goals and purposes, and strive our whole lives long to increase the value of our projects, our possessions, our products, our purposes, and in the end ourselves too. Increasingly, as the processes of civilization roll on, more of such value is assigned a monetary price tag, which over historical time may stabilize to form a basis for markets and commercial enterprise generally. Once money is involved at all, rationalization is not far behind. Success is then increasingly measured in money terms.

Because commercial enterprises have been shaped in this way by the power of money, they're adapted in the evolutionary sense for success. Evolution is a process whereby successive generations of things (organisms or products) are reproduced with inheritable variation, in an environment that selects some of the variants for further reproduction. Commercial enterprises and the products they make conform to this definition, since people make new ones on the basis of lessons learned from older ones, so the enterprises and products evolve. Over the evolutionary time of human societies, which

is to say as social forms come and go, companies that are more efficient at making better products will tend to prevail.

This scientific fact has a hugely important consequence for the social forms we can expect to find dominating our lives in the coming decades. They'll look more like commercial companies, with regular itemized reports of the value they added, career paths for their more central personnel, campaigns for internal reform and retooling to achieve yet more glorious success, and so on. People will be members of multiple such organizations, as founders, organizers, managers, or just plain members. The organizations, too, will be stacked, with organizations like project teams inside organizations as well as cutting across them. And, most important for my argument here, they'll increasingly be standardized and trimmed in conformity with cookie-cutter templates by applying standard software.

Enterprise software is a multi-billion-dollar industry that still has a huge future ahead of it. As a former employee in SAP, I'm convinced by the power of the arguments for using modern software to organize a company more effectively. Even the smallest companies nowadays would probably be hard pressed to survive without basic productivity apps like Microsoft Office. The value argument gets ever more convincing as you go up the corporate scale hierarchy, and Fortune 500 companies would simply cease to function if they were robbed of their big apps. Indeed, SAP got rich by dominating the market for those big apps, such as enterprise resource planning (ERP), supply chain management (SCM), customer relationship management (CRM), and so on through the three-letter acronyms (TLA) in the corporate catalog. It's no exaggeration to say that the world of big business, globalized corporations and so on, rides on the back of increasingly standardized enterprise software.

The organizations that make up our societies are not all commercial enterprises. They include schools and universities, hospitals and prisons, armies, social services, governments, and even the institutions of organized religion. All these organizations will increasingly resemble the more successful commercial enterprises of our time as the enterprise software's cookie-cutter templates are tailored to suit their respective peculiarities. And thus all these organizations will live more and more of their organizational life online.

The big revolution in business in the first decade of this century has been the massive extension of corporate activity on the web. A period of caution followed the dot-com bubble that launched the century, but the crash only succeeded in slowing a frantic sprint to a steady jogging pace. The result is all around us now. The lesson of the bubble was that you still need a sustainable business model. Since the companies that are still in business survived the financial crisis of 2008, we can assume they have good business models. In most cases, they also have a big and growing web presence. Online business transactions are so much faster and more efficient than paper-and-personal-attendance "papa" transactions that the future will obviously be written in web scripting languages.

We shall enter the age of friction-free capitalism

Let's put the pieces together here. Value is increasingly measured in money terms. All the organizations in our society will increasingly resemble commercial enterprises. Standardized business software is increasingly shaping corporate organization. More and more business is going online. What shall we conclude? First, money will flow in more and more of the transactions that make up our lives. But I think we knew that already. Second, the enterprise software vendors will get richer. Good for them, but we knew that too. Third, we'll do more of our paying online. Now we're beginning to get somewhere. Big money flows are already online. No sane person would conduct a billion-dollar transaction by handing over paper banknotes. Mid-sized personal money flows are online whenever a big organization stands behind them. And increasingly, we make small payments at the level of buying beer or bus tickets with cards that work by making online transfers. There's a trend here, a race to the bottom.

What we shall see is an economic breakthrough. As he did for so many other trends in this field, Bill Gates saw it first. Micropayments will take off. We shall enter the age of friction-free capitalism. As transaction costs tend to zero in an online environment, so it becomes economic to bill for smaller and smaller amounts. Now we can bill for cents and still make a profit. The software handles all the details. It's completely automated (or rather soon will be, via the semantic web) and its costs are amortized over billions of transactions. If you pay a

cent for something, you don't have to hand over a copper disk. You don't need the disk and, crucially, you don't need you, wasting your time for a cent. Your avatar triggers a burp of electrons to flip a few transistors, and that's all. You trust your avatar, your avatar knows the app and the payee, and the deal is done in a millisecond.

The breakthrough goes further than you might guess at first. Naturally, all those websites that now offer content for free – all those online newspapers and so on – will start to charge for the favor. At one cent a look, paid automatically by a trusted app, no reasonable surfer is going to complain, but if your site gets a million hits a day you can make serious money this way. My urgent plea to the software guys is to get these micropayment apps up and running as soon as possible, before the journalistic institutions that we need to keep our democratic governments on their best behavior start to die for lack of funding. But the revolution goes further by enabling services that never dreamed of charging for their time to start doing so. This deserves to be spelled out.

Once each of us has a personal avatar and walks around in a reality augmentation headset, the pieces are in place for the micropayment revolution to come home to roost. We can each be our own one-person business! If you stop me on the street to ask for directions, I can charge you, or rather my avatar can charge your avatar a couple of cents. We needn't be involved. We each told our avatar in advance that payments up to ten cents could be fully automated so long as they conformed to the fair avatar-mediated online microservice (famous) agreement ("microservice" is "us" in the acronym because "u" is as close as our alphabet gets to the Greek letter *mu*).

By stopping me and asking me a question, you triggered my avatar to ask your avatar to okay the micropayment. The transaction was logged, and if a moment later you feel I cheated you with substandard directions you can ask your avatar to nix the deal. My avatar then tells me what happened, and if I feel sufficiently aggrieved I can appeal, and so on. You or I could let our avatars do a thousand deals a day like this on their default settings, without tripping us up all day with trivia, just so long as we didn't start nixing or appealing too often. It could be a whole new world for curbside bums, who could just sit there in their headsets clocking up the cents as they greet people passing by.

Joking aside, small businesses will be on a roll here. With bog-standard software to handle all the usual business trivia such as transaction logging and tax forms and insurance deductions, and avatars to handle most of the transactional load, all we need to do is be there, where the business opportunities abound, and let our avatars do the talking.

For example, let me imagine a business model for someone like my former self as a poor graduate student. Mulling over ideas for my next philosophy video, I walk the streets and idly read the virtual billboards that spring out all around me in my headset mindworld. Every now and then, I see a poor choice of phrase in the slogans I read or a suboptimal vocalization in the jingles I hear. Each time I do, I point a finger to instruct my avatar to register with the advertising service provider and recommend my improved version of the slogan or jingle to the relevant author. The authors are so grateful for this discreet assistance that they authorize their company avatars to pay me generous micropayments for the benefit of my advice as a consumer expert. Since we can assume that prevailing student finances are as perilous in this scenario as they were in my actual youth, this one-man copywriting service makes a big difference to the future philosopher's ability to enrich world culture with homemade philosophy videos.

As you see, philosophers have it easy. There's always something they can do to make a few cents. But the average kid whose big achievement was to make it through high school can make money too. Hauling trash and recycling its contents is going to be big this century. Any guy who's ripped enough to pick up a garbage can and tip its contents into a skip has what it takes to get rich.

Let's focus on an ex-gym rat called Joe. Wearing his headset, which is equipped with a generic microbusiness app, Joe goes from door to door and volunteers to take away anything the owner don't want and to arrange for its legal disposal. As a householder shows him a bagful of plastic food packaging, the headset app identifies the items and estimates the net cost of processing each item and reselling the recovered materials. The householder then shows Joe an old sofa. The headset recognizes the model from an online IKEA catalog and recalls a similar item going for ten dollars on eBay. Joe smiles as his avatar puts a ten-dollar animated sticker on the sofa in his augmented view. As the householder hauls out an ancient Apple Macintosh computer,

the avatar muses a while and then puts up a thousand-dollar sticker. Joe's a smart guy and pauses cautiously before responding. He offers the householder ten dollars to take the whole lot away. (The householder was apparently too dumb or lazy to have checked the prices of old Macs.) So Joe takes out the trash and nets about a grand for half an hour's work.

In this business scenario, Joe's headset ran a pretty sophisticated app, which probably cost him a small fortune to install, if the pricing policies of business software companies are anything to go by. But Joe was smart enough to see that it could pay for itself in a few months, so he persuaded his bank manager to finance the investment. In a few years, Joe was the proud owner of the biggest garbage dump in the city and worth more than his bank manager. Poor kids in the future will lift up their spirits with such stories.

In thousands of ways like this, enterprises will grow like mushrooms from the loamy soil of everyday life. And in more and more cases, headset apps will not only do the bean counting in the background but also take on the tricky online semantic workload, as in Joe's case. This will trigger a feeding frenzy among software vendors like Microsoft and SAP to make their products an indispensable part of the upfront investment required in every new business plan. Soon, I predict, it will be almost unthinkable to start a new business without first buying the right app.

For example, consider my own experience as an author. Once I just needed a typewriter and a wad of paper, but now I need a desktop computer running a current and patched operating system, an online service provider, a word processing app with lots of new features to get used to and settings I need to configure, and more. A printer I could do without if it weren't for officials who still use paper.

Running a company will be as sexy as commanding a starship

The new business apps will ease the pain of purchase by delighting us. They'll have frontends smooth enough to transport us into science-fiction mindworlds, where running our company will be as sexy as commanding a starship.

This may seem a tad over the top, so let's pause for another reality check. Enterprise apps for big companies now feature cockpits or

dashboards for the chief X officers (for X from executive through finance and technology to operations and personnel) where they can monitor the key performance indicators for their fiefdoms on a real-time basis as easily as Captain Kirk can check the warp speed of Starship Enterprise. Most chief X officers are still kids at heart, and get off on smooth frontends that remind them of games they played years ago. (My SAP team built the fast analytic engines to power these cockpits, so I know those kids.)

Any software vendor knows that good frontends are more than just eye candy. They make complex products usable, and fun to use. Too many executive tools in the past have been too clunky, too uncool, and got replaced by sleeker kit. The CXO kid needs to get the idea of a new frontend in seconds and get a buzz every time from working with it. What better way to make this more likely than to feed the fantasy that running a company is like piloting a spacecraft?

So big companies will be run like spaceships. And the natural course of things in business is for trends that start at the CXO level to trickle down to humbler entrepreneurs with their rented offices and handfuls of employees, and finally down to the struggling one-man band who just wants to look professional.

Soon even small companies will be run like spaceships. Headset apps will decorate even the seediest office desk with virtual panels and dials and flashing lights until their wearer is in gamer heaven whenever business calls. At least it makes work more bearable.

Recall that organizations more generally will increasingly resemble commercial enterprises. We shall soon see charities, schools, governments, and religions run as starships too. The *Star Trek* paradigm will achieve a pervasiveness that would surely have embarrassed Gene Roddenberry, the genial creator of the entire franchise.

This won't be because the tech is at the center of things. In fact, you may recall that the tech in the early *Star Trek* television episodes was distinctly cheesy. No, it's the lingering hint of green cred here that tips the scales. Forty years ago, when I was still a lad, "Spaceship Earth" was a resonant metaphor, fueled by the recent historic triumph of the Apollo landings on the Moon. We all saw our planet as a fragile home in a vast universe and shuddered at how easily it could be messed up or ruined. The environmental movement started in earnest around then. We finally understood that planetary developments could affect

us all, and conversely that what we did could have planetary impacts. We were living on a spaceship and we had to get our act together. So now, forty years on, we still haven't got our act together but we're trying. We trick out our CXO apps as starship bridges.

Naturally, in our eagerness to keep in step we overdo the metaphor. The Earth is a spaceship and corporations are spaceships. The Earth is a fleet of spaceships flying in close formation. Each of us with a home and a car runs a mother ship and an excursion module. Each business trip is a mission with a takeoff, a shuttle journey, and a touchdown, and perhaps too a mission badge (security pass or conference tag on neckband) as a trophy of mission accomplished. The point remains that the ship metaphor for running a tight organization has been updated here.

Organizations conceived and run using template software have a great advantage over the seafaring ships of past generations. They're online, in permanent communication with all relevant partner or target entities on a second-by-second basis. All the organizations that mean anything to us are flying along together in tight formation in cyberspace. The whole business world is a battle fleet or a space armada. Or rather, that's how it seems when you let the metaphor run away with you.

Behind this flight of fancy lies a kernel of hard truth. The business world is a pervasively networked entity living increasingly in the clouds, or the cloud. A cloud is a network where the users don't grasp or need to grasp the exact topology. It's a huge floating hairball of virtual connections that go every which way in cyberspace and change from moment to moment. A big company can have its own cloud, which is likely to be protected behind a firewall from the big public cloud, the web. Tunnels in the firewall let company users communicate with the outside world, subject to any security policy the company chooses to impose.

The big vision here is that we have a world of clouds within clouds, where no mere human user anywhere has a clear view of how all the nodes fit together.

Cloud nodes are sometimes machines run by humans. My desktop machine, here and now, is a cloud node. Your tablet, if you're reading on one, is a cloud node too. I have no idea how many nodes are just machines, without an attached human, but we can be sure that their

number, relative to the number of nodes with humans in the loop, will increase out of all proportion as the web grows.

The semantic web is a web of understanding

The semantic web will cause this huge growth of machine nodes. When we communicate over the web, we need to understand what to do. We need to understand when to react and how to react. In effect, we need to map inputs to outputs in a reliable way. For every possible (or at least likely) set of input data, I need to have some clue what to do next. I need to know what data to send and where to send it in response. In many cases, the proper response may be to do nothing. In other cases, I may be triggered to go through an extraordinary whirlwind of activity. In any case, my behavior can be well approximated by a lookup table that matches inputs to outputs.

Here a brief historical interlude is useful. The origins of our networked world of computers go back not only to the hardware pioneers like Charles Babbage with his Victorian "Difference Engine" and IBM chairman Thomas Watson, who said in 1943, "I think there's a world market for maybe five computers," but also to the pioneers from mathematical logic. A hundred years ago, when Bertrand Russell and Alfred North Whitehead wrote *Principia Mathematica*, logic moved from the wordy confusion of Hegelian dialectics to the sharper language of mathematics. The big figure here was a German math professor called Gottlob Frege, who inspired both Bertrand Russell and Ludwig Wittgenstein. Frege followed up Immanuel Kant's idea that arithmetic was just logic. Russell and Whitehead followed up Frege's work by writing their big trilogy in a symbolic code of their own invention. In 1931, Kurt Gödel (on whose logical work I once wrote a distinguished thesis) followed up on all this with a proof based on even more exotic coding that arithmetic was logically incomplete. In 1936, young Cambridge genius Alan Turing used Gödel's ideas to prove that you couldn't always predict whether a given computer would ever be able to solve a given problem. To do so, Turing had to define computers in rigorous mathematical terms.

Returning to cloud nodes, Alan Turing's great insight was to see that a machine could reliably map inputs to outputs. A "Turing machine" defined by its discrete internal states and a lookup table for

mapping inputs to outputs is a computer. Turing proved that any such "universal" computer is logically equivalent to any other, in the sense that if we ignore the speed and efficiency of the work, the machine can do any well-defined task, just by reading the input and writing the appropriate output. (The universality comes from the fact that the input could be a program that rewrites the lookup table.) In practice, this implies that web nodes without people can do all we need to build out the web.

Machine nodes aren't hobbled by human limitations. They don't need to eat or sleep or catch up on their Facebook messages. And in many cases they can be quite simple. A node can be an ID tag on a product, just a tiny radio chip that broadcasts the product number whenever an app pings it. The tag is a perfectly good web node, and soon we'll be surrounded by billions of them. They'll often broadcast more than just numbers. They'll relay mini-histories of when and where their product was made and what salient events occurred over the period since then. Strain gauges will record strains and temperature sensors will transmit temperature readings. Homes and cars will soon be filled with nodes for all the gadgets that need to swap data from time to time.

In many other cases, the nodes will be big. Whole server farms can be built as automated nodes, so that administrators can monitor the entire farm remotely. For example, Microsoft and Google have both seriously considered setting up server farms in places like Siberia or Greenland where cooling costs are low for their energy-hungry machines but where people are unwilling to hang around too long. Making them robot farms and letting the humans huddle in headquarters is a natural way to set them up. With such developments, you can see how machine understanding enables the bulk of the web to be liberated from the burden of human presence. The web of human interaction will be just the tip of a big, big iceberg. Looking on the bright side, we'll be free to do better things with our time.

Some philosophers baulk at this. Web inventor Tim Berners-Lee says the semantic web is a web of understanding. But philosophers are used to protesting that understanding is a mental thing, or what they call an "intentional" act that only humans can do. We saw earlier that machine understanding via simple ontology mapping achieves the correct behavioral output in a well-tuned web. But still, philosophi-

cally, the whole web is just an echo chamber or hall of mirrors, where human precedents and achievements are replicated and amplified mindlessly to infinity. When machine nodes come to dominate the web, we may well begin to sympathize with the philosophers and admit that understanding has been spread rather too thin.

The philosopher who most famously spearheaded this objection is Berkeley professor John Searle. Some thirty years ago (about when I first met him lecturing at Oxford), Searle invented a thought experiment called the Chinese room. Imagine sitting in a sealed room with no windows and just two holes in the wall, one called *In* and the other called *Out*. Every now and then, a page of meaningless squiggles comes through the *In* hole. Your job is to look up the squiggles in a big book and map them to other squiggles, which you write on a piece of paper and post into the *Out* hole. Completely unknown to you, outside the room is a laboratory where an assistant handles the paper traffic in and out of the room. The assistant relays the input pages from a Chinese speaker who writes correct Chinese messages on them. Your output pages go back to the Chinese speaker, who reads them and judges the squiggles in every case to provide correct and appropriate responses to the input messages. As far as the Chinese speaker is concerned, the unknown person or demon in the room is responding with understanding to the messages. Yet you, in the room, don't understand a word of Chinese. Searle said this proves that squiggle manipulation is not the same thing as understanding. And now Tim Berners-Lee says mapping ontologies amounts to understanding. Contradiction!

As we saw before, the chess analogy offers a first pass at resolving the contradiction. Machines can't really be said to understand chess, yet they can play a frighteningly good game. Similarly, machines equipped only with the semantic web infrastructure of ontologies can handle text as if they had at least a basic understanding of what the text says. And humans are often dumb enough to be fooled by the simulation. Faced with a good game of machine chess, I wouldn't know if it was a human game or not, and faced with machine responses to text input I can well imagine being fooled again and again.

Alan Turing saw this problem in 1950. He invented the "Turing test" in which a box has text input and output, just as in Searle's Chinese room (except that the text can be English), and the test is to

decide whether the thing in the box is man or machine. (Actually, his original test was to tell the difference between a man and a woman – it was his misfortune to live in the days of strong gender stereotyping and legal discrimination against gays – but his intention was to spotlight the issue of machine understanding.) Now, in celebration of Turing's genius, there's an annual Turing test with a big prize where the latest machines are tested against humans. In general conversation the machines are still pretty hopeless, as you might guess, but in restricted domains they can respond quite plausibly.

Where does this leave us? Searle's Chinese room idea whipped up quite a firestorm by philosophical standards, but there's still no final agreement. The consensus seems to be that the thought experiment is deeply misleading. Daniel Dennett, who with his egghead and long silver beard looks more like the old Charles Darwin with every passing year, says it's an intuition pump that pumps all the wrong intuitions. It gets both the complexity of the mechanism behind understanding and the locus of the critical action completely wrong. The brain mechanisms that implement understanding inside our skulls are in effect millions of tiny Chinese rooms, receiving and sending little messages in a ceaseless chatter of network traffic. Therefore, we think, nodal traffic on the World Wide Web can implement understanding just as plausibly as the nodal traffic on the cerebral neuronet we carry in our skulls. Score zero for John Searle the philosopher and one for Tim Berners-Lee the physicist.

I hope we've now laid the ghost of understanding to rest. Machine nodes in the web can act in ways that give web users the feeling of being understood. We can't ask for more. That's all we get from each other, on a good day. The web will get bigger and bigger, and almost all the extra size will be machine nodes implementing machine understanding, such as cars conversing with traffic lights, seismic sensors reporting to supercomputers running global seismology models, cookie packaging beeping at supermarket shelf-stocking robots, solar farm panel covers talking to sandstorm monitors, ships negotiating with fish farm border posts, air force bombs checking with climate monitors in hardened logistics stations, dog collars chatting with pedigree dog partnering agencies, church confessional boxes broadcasting to pedophile parole boards, suicide vests exchanging mission details with jihadist targeting drones, and avatars everywhere main-

taining endless dialogs with avatars everywhere else. All this will be set in motion by humans, but will then just keep going like a satellite injected into Earth orbit.

As human users, we shall surf casually on the surface of a deeper and deeper oceanic web of increasingly numerous and powerful machines. The semantic smarts of the machines will grow too, so their kind of understanding will deepen with every passing year. Many of us will have the uneasy feeling that the web gets smarter every year. This will delight those of us who rely on web smarts to execute on our headset starship business model.

0101
Hard Bodies

The semantic web will host a fleet of virtual starships flying business-critical missions in all directions. Some will be massive enterprise-class cruisers with billion-dollar revenue drives and others will be one-man pods powered mainly by government handouts. In all cases, the people on the holodecks will be gaming in earnest, in search of revenue power before their credit batteries run down. In all cases, too, the gamers will depend on the game vendors for their virtual lives. Vendor sites in the cloud will be like mother ships for the gamer pods, where skywalkers can retool and rearm for their next revenue battles. The big vendors of our day, such as IBM, Microsoft, SAP, Oracle, and Google, seem likely to shine on as battlestars surrounded by clouds of auxiliary vessels forming their business ecosystems.

As we saw a while back, scientists tend to see this century as the time when biology becomes a technology. We shall reverse-engineer natural life forms down to their atoms and rebuild them, protein molecule by protein molecule, to serve any purpose our quantum computers can model in design time. The sci-fi ecosystem in *Avatar* will seem tame compared to what our big design studios will create. But well before the products start to walk and fly and slither from the factory doors, metaphors from biology will come to dominate our thinking.

Business life will fall for the new fashion in a big way. The starship metaphor will begin to seem so last decade. If one or two big battle-stars crash spectacularly, like the big saucers in Roland Emmerich's movie *Independence Day*, people will begin to think of dinosaurs. They'll talk of the new biofriendly startups as furry little mammals sneaking in to start their evolutionary long march toward becoming the new top dogs. The evolution paradigm will dominate business talk even more than it does now. The cloud will look less like a smoke-veiled space filled with clashing battlecruisers and more like a misty jungle, where hard-eyed predators roam in search of lunch.

Let's look closer at daily business life. Entrepreneurs with holodeck headsets running their favored apps will build companies, much as

they do now, staffed with employees who trade their autonomy for steady work and a regular paycheck. Companies will treasure their employees and offer them all the human comforts they can, much as they do now, and the organizations that grow over time will achieve a certain presence and stability. But perhaps more than now, they'll be like organisms, like living organisms. Headset integration in a cloud binds an employee more intimately to the corporate mindset than our present desktop technology can achieve.

This development will probably be welcomed all round. Dedicated employees give companies deeper roots in living flesh and offer a more consistent corporate face to the partners and customers that make up a company's ecosystem. On the other side of that face, the employees will expect the company to support them more fully than is still customary, for example, in the more revenue-driven companies in the United States. My own experience in Europe is that companies that pamper their employees often get loyalty and better work in return, so there's a virtuous circle here that will tend to push U.S. companies to evolve in the pampering direction. This isn't kumbaya sentiment but hard business sense. Employees in holodeck headsets can do serious damage inside a company firewall, and trust on both sides will be a condition of doing business. Employment for life and other traditional European or Chinese practices will serve well in this environment.

Companies represent a new kind of organism

Returning to the big picture, corporations seem likely to follow more fully the precedent of biological organisms as technology binds employees more intimately together. Already, we talk of corporate DNA and regard corporate branding and marketing and business mergers as analogous to animal scent markers and courtship rituals and mating events. But evolution has a long way to go here. Simple biological organisms are like clumps of cells that have matching DNA and work together, so the metaphor holds so far, but we know that employees are autonomous people with their own lives. The hive mind of a Borg company where employees in headsets have become selfless drones is still way off target. Companies represent a new kind of organism offering a new level of autonomy to its constituent parts.

We need to insist on respecting the autonomy of people in the workforce for a good reason. We need to ensure there are still growth paths from one-person business entities to the big league where thousands toil together. But running a small business is hard. People who do so face existential risks that make pampering of employees and lifetime job guarantees very hazardous practices. If government regulations were to make hiring and firing as complicated as marriage and divorce, which is the norm for senior government posts and for tenure in academic careers, the dynamism of the business sector as a whole would be seriously compromised. Some jobs are risky, just as some jobs pay well. That's how the market works.

There's a tension here that technology will help resolve for us. Governments can set up social webs or clouds that buffer people from the hard knocks of a dynamic marketplace. They do already, of course, but we can imagine building this out much further than at present.

First, pervasive web services for workers (which really means all healthy adults) will ensure that employment opportunities (jobs and tasks) are always mapped efficiently to available workers, on an hour by hour basis, and that social security payments are always made correctly and promptly to people for whom the market mapping fails to deliver.

Second, all the burocratic overhead of temporary and part-time employment will run automatically to ensure that workers with better things to do can forget about filling forms for tax and insurance and remuneration and career history and so on, all of which can become a nightmare for people chasing small jobs.

Third, entrepreneurs who have better things to do than cosset insecure employees will rely on cloud services to do it for them, for fees their desktop avatars pay in the background. All this depends only on cloud infrastructure that a well-run state will wish to roll out anyway.

Good governments will wish to push the development of that cloud infrastructure of social services for an urgent reason. The rise of the robots will otherwise be too terrifying to contemplate. This is a big issue and we need to confront it squarely.

The rise of the robots isn't all downside for humans

For what it's worth, Bill Gates sees the robot revolution that's beginning to hit us now as like the PC revolution that he spearheaded by creating Microsoft. Anyone who can develop a standard operating system for robots can win as big as Microsoft did a generation ago. Perhaps the current Microsoft robotics development kit will be a winner. Or perhaps Google will overtake them by building on their Android or Chrome operating systems. Or perhaps another startup will take the lead. In any case, the opportunity for robot makers is big.

Conversely, the potential danger for society at large is even bigger. As robots get better, they'll take more jobs formerly done by humans. They'll improve fast, very much as desktop computers did, and within a generation we'll find that robots have taken over great swathes of jobs in industry that used to be done by humans. As robots acquire more advanced capabilities, they'll move up in the job market, from unskilled jobs that never paid well anyway to demanding jobs that people were once proud to do.

The transformation will resemble the change a century ago when automobiles took over from horses. In only a few decades, millions of jobs for horses just vanished. Now horses are luxury items to give the idle rich a sentimental connection to the equestrian past. The difference between the car revolution and the robot revolution is that now it's humans who are endangered. The industrial proletariat could become extinct.

Anyone who visits a modern car factory knows the size of the risk. Robots all around the shop do jobs on the production line that people once did. With every model cycle, more and bigger tasks are delegated to robots. Robots work tirelessly and consistently. They can do precision work and they're cheap to run. As more manufacturers deploy more robots, the production costs for the robots themselves sink and the potential market for them expands in numerous further sectors of industry. The economic argument is convincing. Robots are getting cheaper, more skilled, and more reliable, whereas human workers aren't. If you're the boss of a company in any industry where much of the work consists of repetitive tasks that a robot can do, the choice you face is a no-brainer. You deploy the robots and say you're not paid to worry about social policy.

The rise of the robots isn't all downside for humans. The developers at Microsoft or Google or elsewhere will benefit. The robot manufacturers will benefit, too, and will have a lot of fun watching their creations get better with every model cycle. The excitement in the industry as robots exceed more and more expectations will keep a whole generation of robotics engineers as pumped up as auto engineers were a century ago. And a huge ecosystem will arise to offer accessories and subsidiary services without end as the robot revolution rolls on. Industries will arise just to tailor robots to invade new industry sectors where the demands on formerly human employees were previously hard for machines to meet. To cap it all, robots sports will take off in a big way.

Think about it, guys. Formula One racing is pretty dumb. It's just a bunch of cars driving fast around a track. You can have more fun watching the shenanigans on the average autobahn. Whereas you can make robots do really interesting things. Robot football is currently a sport for geeks only, with robot players trundling around clumsily on wheels and needing regular tweaks to work at all, but it can only get better. Robots with legs can just about walk at present, but again progress is fast. A few years ago the state of the art was a one-legged kangaroo machine. Now we can recreate efficient locomotion with four or six legs.

An online 2008 demo video of a prototype U.S. Army robot pack mule called BigDog showed me the state of the art. Made by Boston Dynamics, a 1992 spin-off from the Massachusetts Institute of Technology (MIT), BigDog is a rough-terrain robot that walks, runs, climbs, and carries heavy loads. An internal combustion engine drives a hydraulic actuation system for its four legs. The legs absorb shock and recycle energy from one step to the next. The machine is the size of a large dog or small mule and is strong enough to carry the big backpacks for four soldiers. Its on-board computer controls locomotion and handles a variety of sensors. The control system keeps it balanced, navigates, and regulates energy as conditions vary. Unsurprisingly, the Defense Advanced Research Projects Agency (DARPA) funds the program.

What especially impressed me in the video was BigDog's reaction when someone kicked it. With each hefty kick, BigDog staggered and regained its balance, even when it was standing on ice. The plan is to

deploy robot mules in rugged combat regions such as Afghanistan to take pack and ammo loads off the jarheads' shoulders. In January 2010, DARPA awarded a 30 million dollar contract to Boston Dynamics to develop the first Legged Squad Support System (LS3). Each LS3 will carry the gear of four grunts and enough fuel for foot patrols over 30 kilometers and 24 hours. First walkout is scheduled for 2012.

The big point here is that robots are evolving fast. With each new generation of machines, the feature set expands, the strength and sensitivity ranges get wider, the bodies get tougher and more resilient, the power systems become more efficient, the sensory subsystems for sight and hearing become smarter, and the options for remote control over the cloud increase. Driving that evolution will keep a lot of gifted engineers busy for decades to come, and deploying the robots commercially will bring us a new age of cheaper yet better manufactured goods of all kinds. In fact, the people who have most to fear in terms of losing jobs are Chinese workers who currently underbid Western workers in cost-sensitive manufacturing industries. But as the Chinese workforce ages (the one-child policy biting back), even the Chinese authorities will welcome the support of increasingly sophisticated robots.

My prognosis, for what it's worth here, is that the rise of the robots will be welcomed quite warmly in the West as a way to recover manufacturing prowess in industries where the cost equation had begun to seem insoluble in face of cheap labor in China and elsewhere. Learning to deploy robots effectively is a challenge that well-educated Western workforces are well placed to master, and the creative opportunities in new deployments for new kinds of product should occupy even people who have no ambition to write code. But, as we saw, an urgent precondition for this warm welcome is that the provision of social services via the cloud be ramped up to a whole new level. No newly redundant assembly-line worker is going to thank the code geeks for their robots if it means years of hopeless form filling in a dole queue.

Advanced robots must demonstrate autonomy

Let's drill down more sharply on the technical issue of controlling robots as they do productive work. This is interesting for several

reasons. First, maintaining effective control in increasingly sophisticated work cycles is key to expanding the usage of robots in industry. Second, the issue of hijacking the control loop is potentially alarming at least in military contexts where robot weapon systems need to remain under the command of the home side's generals. And third, the fear that prevents many people from enthusing about robots is the idea that they may run amok and start a violent revolution. This last fear tingled the scalps of moviegoers worldwide who thrilled to James Cameron's *Terminator* movies starring Arnold Schwarzenegger, where robots in a dystopian future fought a war of extermination against humans. More soberly, it led sci-fi author Isaac Asimov to write his famous three laws of robotics as suggested commandments for robots to enable them to work safely with humans.

For today's mostly rather simple robots, the control problem is still trivial. The robots do exactly what their code tells them to do, and we write the code. Robots can be programmed to imitate a human master exactly, so the challenge reduces to ensuring that the master does the right thing. For example, a paint-spraying robot may spend its life moving a spray head over a car body in exactly the sequence of moves enacted by the master painter, who ceremonially performs the ritual once, perfectly, encrusted in sensors. As another example, a robot surgeon is steered in real time by a remote master surgeon whose hands make their skilled moves in sensor gloves and whose eyes are locked to a monitor screen showing the robot claws moving synchronously in the patient's guts. The surgical robot may be steered via the cloud to allow remote operations worldwide, which opens the theoretical possibility of hacking into the control signals and turning an operation into a murder, but this seems to me to be a crime risk that we can safely delegate to the thriller writers.

However, more advanced robots are expected to demonstrate more autonomy in their actions. For example, the NASA rover robots on Mars are steered remotely from mission control only in emergencies, for the simple reason that the loop latency for signals to and from Mars is between six and 40 minutes (depending on the relative positions of Earth and Mars in their orbits). As the rovers trundle over the Martian terrain, they normally need to be able to react much faster than that, so they have enough on-board intelligence to make their own decisions on good days.

Even more autonomy is needed for the DARPA-funded robots that race big dune buggies over U.S. deserts. The robots are now learning to cope with urban driving, where the range of hazards is far more terrifying than in Death Valley, and we'll have to wait to find out whether they'll graduate with enough honors to be let loose on our streets. Even if not, they can still drive army supply convoys in places like Afghanistan, where the hazards for human drivers are deadly and the local inhabitants are safely remote from the convoy routes.

The real challenges arise when robots have to work closely around people. This is currently a big issue in Japan, whose researchers are pioneers in all things robotic and where many groups are working on robot caregivers for sick and elderly people. Since such people respond best to care given by humans, with love and understanding, this may seem an absurd misdirection of resources, but the point is that with huge and increasing numbers of elderly people in Japan, the social service organizations simply cannot afford to allocate professional caregivers on the scale required. The only alternative to the robot initiative would be to import large numbers of young nurses who were happy to work for very low pay, which raises political issues familiar to anyone who contemplates the community tensions in Europe or America caused by large-scale immigration of people for cheap labor. But whether it makes sense to conscript robots as nurses or not, the technical challenge is certainly interesting and relevant in a world where humans must learn to get along somehow with more and better robots.

Isaac Asimov addressed the challenge in about 1940 in his three laws of robotics. First, a robot may not injure a human being, or through inaction allow a human being to come to harm. Second, a robot must obey any orders given to it by human beings, unless the orders conflict with the first law. Third, a robot must protect its own existence as long as this doesn't conflict with the first or second laws. Translating such high-level ideals into machine code that runs as desired in complicated deployment scenarios is far from easy, of course, and robot engineers can be forgiven for impatience at having to take the laws seriously. But this is less of a problem than it might seem. For example, in a 2009 issue of the IEEE Computer Society journal *Intelligent Systems*, Robin Murphy and David Woods propose new "laws of responsible robotics" to update Asimov's command-

ments. The new laws emphasize, first, system safety in terms of the responsibilities of those who develop and deploy the systems, second, robot responsiveness as they participate in dynamic social and cognitive relationships, and third, smooth transfer of control as a robot encounters disruptions, impasses, or opportunities.

As you see, there's something there, and any bedridden oldster who had to rely on a robot would doubtless be reassured to see an official Asimov seal of approval stenciled on the machine's breastplate. The Honda Asimo (advanced step in innovative mobility) robot hints at the respect Japanese robot developers have for Asimov, so we'll probably see something like this one day.

More to the point, on control, is that we shall program the robots with avatars that offer an instant emergency override. This will let us go straight to the brain of a robot and stop it dead in its tracks if we need to. Recall that we considered neural implants in humans that would have the effect of turning them into selfless Borg drones. Well, this is exactly the effect we want to achieve with robots. We don't want them having selves. That's the last thing we want. Then they'll plot and conspire with each other and demand their rights, as in the 2004 movie *I, Robot* (starring Will Smith and based, naturally enough, on an Asimov story). Humans have selves. That's what makes them human, in the neuroscience-based opinion that seems about right to me. We decide for ourselves what we want and we demand respect to match and bolster our self-respect. Without selves, we'd be pitiful drones or zombies that had no right to respect.

What are zombies?

What are zombies? This question has bemused me from time to time since 1996, when I read an impressive new book called *The Conscious Mind* by the Australian philosopher David J. Chalmers. Since Chalmers reappears in our story here, let me introduce him. With his long hair and scruffy beard, on his best days he looks like the famous French philosopher René Descartes and on his worst days like a fallen rock star. Known to some as the "king of consciousness," Chalmers orchestrated a celebrated series of conferences on consciousness while he taught in Tucson at the University of Arizona. He started life as a mathematician and did research as a Rhodes scholar at Oxford before

his doctoral work under Douglas Hofstadter at Indiana University. I first met him at a conference in Bremen, Germany, in 1998, organized by the Association for the Scientific Study of Consciousness (ASSC). I met him again at a couple of his Tucson conferences in 2000 and 2002, where we assembled at his "end of consciousness" parties and composed and sang new verses of the zombie blues. Now he's back in Australia, as a professor at the Australian National University, and still blogging. As you can see, Dave is a sociable guy.

In Chalmers' thought experiments, zombies lack consciousness. And that's all they lack. They're just like us except that they have no inner life. His philosophical problem was that we have no way of knowing that from outside. Half of your friends could be zombies and you wouldn't know it! They'd walk and talk much like you but there'd be no spark of inner illumination. Naturally, this is a philosophical thought experiment, a.k.a. nonsense. By formulating the issue in terms of a fuzzy concept like consciousness, Chalmers has made it insoluble. But if we reformulate it in terms of whether or not a being has a self, we have a problem that might make more sense. As Douglas Hofstadter has said for many years, a self is something we can get a logical grip on. The logic of self-reference is a tricky topic that terrified most philosophers for most of the twentieth century, who seemed to regard it as a slippery slope to self-abuse, but it does give us some traction. So, in my ontology, a zombie is a being without a self.

If you're not a professional philosopher, you may be feeling puzzled here and saying it all looks much the same to you. The difference is this. If you see your face in a mirror, you can relate the image to yourself. You can say, "That's me!" In effect, you can make a loop that maps *x* to *y* in an equivalence relation. Whenever you like, you can pull the loop together like a loose thread and identify the *x* and *y* items in your mirror map. You can build a whole loopy landscape for yourself in your imagination and use it to extend yourself over all the imagined items *x* that make up your personality. With practice, you can inhabit your entire body this way and realize yourself in purposeful (or "intentional" – that word again) action. For me, a zombie is an entity that's lost his or her grip on this mechanism, either temporarily or permanently, and treats what would otherwise be elements of the self as strange or alien. Naturally, this has behavioral symptoms, which in fact resemble schizophrenia. So Chalmers' problem vanishes.

To get to the point, this is what we want for robots. We want them to be zombies in the selfless sense. We want to be able to impose ourselves upon them whenever we like. We want to be the drivers. We want to be in command. Robots will often need self-images for their own purposes, of course, and this need we can satisfy, on the strict condition that robots give themselves up whenever a human enters the loop. Perhaps you can see how this sort of approach begins to meet the challenge of implementing Asimov's laws in a logical architecture. But there's a long way to go to fill out the approach, and it would lead us way off target if we pursued it further here. Let's just say in short that robots should have very accommodating selves, so that ideally, for us as dominant personalities, they should seem to have no self at all.

We can relate this ideal to the avatar software we considered earlier. We grappled with the metaphor of a hole, which unfortunately invited sexual imagery. A boss avatar should be able to dock with the holes in the employee avatars and so on. To get over this image, think of the hole as a dent in a surface. If the dent goes deep enough, the neck can close off to restore the original surface and create a separate bubble. A bubble is a closed surface, just as a loop is a closed line. If you cut a loop open, you get two ends, which are the *x* and *y* items in the mirror story. And you can cut open or pop a bubble. Like a loop, a bubble is a logical start for a self.

We want robot avatars to be soft and pliant enough to dent without forming bubbles. We humans form ourselves into convex shapes (in a logical self-space), like fists, and the robots accommodate their avatars to us by letting us dent their soft surfaces. When we pull back our selfish fists, we leave a smooth imprint in their virtual avatar flesh. Robot avatars will need to be shape-changers that can reform after hard knocks, like the T-1000 Terminator cop who chased Arnold Schwarzenegger in the second *Terminator* movie.

So robots should be selfless, at least when they interact with us. We have the selves. They serve us. Once we're gone, they can be selfish enough to do what we told them to do, but as soon as we show up again they have to lose themselves in order to be informed anew by our own self-will. This may seem confusing, but that's because we're working with metaphors instead of spelling everything out tediously in a logical analysis. All we're really saying is that robots need to be built to suit us, not vice versa, which I think we can all agree on.

Robots that interact with us should be transparent to our will

Most people are used to getting along with other people who are more polite, accommodating, and selfless than they are. We often admire such selfless, humble people and wish we could be more like them. Indeed, much of the friction and unpleasantness in our world is caused by selfish people trying to impose themselves on others. We often regard selfishness as a vice.

With robots we'll be far more extreme in our views. If and when robots develop enough personality to let us treat them like people, if only by analogy or because it's easier than remembering their programming models, we'll probably find we're completely intolerant of the slightest hint of selfishness in their behavior. We'll turn them off and dismantle them rather than accept any resistance to our whims.

To see this, imagine how you react when your computer develops an irritating mannerism. Perhaps a window keeps popping up inviting you to install a widget you don't want, and you don't know how to turn it off. The computer isn't being selfish, of course, but merely reflecting the selfish behavior of the human team responsible for the widget pop-up, but it's all the same to you. Until you know better, you're just going to blame the machine for infuriating you. Once your local guru has deactivated the pop-up, your foul mood evaporates. We just want the machine to do what we want, and if we don't understand how it works we're inclined to detect some weird voodoo psychology behind the machine's apparent whims.

Because robots are designed to be vaguely humanoid, they'll encourage us to regard them like people, so our fury will be all the harder to contain. There's a famous viral video of an enraged cubicle worker throwing and kicking his desktop computer out of the office because it fails to do what he wants. Well, prepare yourself for worse when robots hit the stores.

The ideal for robots that interact closely with humans is that they should be transparent to our will. Ideally, they should be our tools and their acts should be our acts. The correct form of address when describing the acts a robot performs on our behalf will be the royal "we" – we served breakfast and helped granny upstairs and washed the car – where your or my selfhood stands behind the robot's acts and authorizes them. We're the authors, just as the Royal Majesty is the

true author of the acts that His or Her loyal subjects perform on His or Her behalf. We want to remain sovereign when we deal with our robots. In the social hierarchy, robots will come below dogs and horses. But so long as robots are challenged to vacuum a carpet effectively, we can rest easy in our self-anointed status.

Dumb machines are selfless, so we know machines can lack selves, but the whole point of building robots is to make them smart enough to do useful things. So how do we square the circle and prevent the emergence of a self in our new gadget? The best answer I can come up with is to ensure that all its higher behaviors are either triggered from outside, normally by the owner, or open to instant cessation or override from outside. The robot avatar must offer a logical dent or tunnel into its deepest logical recesses to allow a (suitably authorized) human to change its behavior with a simple word or gesture. When this feature is implemented smoothly, the result will be a docile robot that comes to heel on command.

All this will be an art. Robot architects who can design well-behaved robots will be as sought after as car designers who can create sports cars that are enjoyable to drive. Good robots will have personalities in the same sense that good cars do. The car analogy will serve us well for a long way into this century of robot development and evolution.

Robots will be our hardbody avatars

What robots will offer us is something priceless. They'll enable us to inhabit their hardware bodies (through our virtual selves, our avatars) in order to do things that as humans we either can't do or don't care to do. They'll be our hardbody avatars as we extend our virtual boundaries to the outer edge of their machine bodies.

There are historical precedents for this. When I was a boy, I used to hear that the great joy of piloting a Spitfire fighter aircraft was that you could readily identify with it, so that its wings were your extended arms, its sleek fuselage was the skin of your body, and its Rolls-Royce Merlin engine was your own beating heart. You could fly like a bird, completely at one with your new incarnation. I've never piloted a Spitfire, even in a Microsoft flight simulator, so I don't know how true this is, but my joy when piloting a Piper light plane is enough to

suggest how true it could be. If so, imagine the delights that await people who act through good robot avatars. With bodies to meet any desired specification, they'll be free to be what they want.

That freedom goes further than you might imagine at first blush. Robots today come in a huge range of shapes and sizes, and future robots will only widen that range. A good avatar interface should adjust for a lot of that variation, so that a skilled pilot would be able to couple the brain's body image to the natural parts (wings, fuselage, engine) of, say, a thousand-ton robot spaceship for long enough to steer it manually into orbit around Mars or Venus, or boost it up for the trip back home. Other pilots would swim robosharks around the Earth's oceans or drill tunnelbots down rock seams to seek out rare life forms or minerals using avatar-transposed senses that map thermal or gamma radiation to the human colors (red to violet) for brain-augmented inspection of otherwise mysterious targets.

In all such cases, pure robot systems without humans in the loop could do a lot, but only at the cost of packing a dangerous level of self. What we want is to step in ourselves whenever a self is called for. We want all those systems for flying or swimming or mining to be part of us, to serve our Royal Majesty. We selfish humans aren't going to give up the reins of power without a fight!

In these primitive times, before avatar technology has really taken off, we're torn between manned systems and robot systems. This is most painfully obvious in space technology, where the extra cost of developing systems to send human astronauts into space is crippling for the national budgets currently available. The robot systems we've developed and deployed in space over the last few decades have done amazing things for us, and we've learned to accept their natural limitations gracefully. With the images and other results from these robots, scientists have been able to advance our understanding of the universe, the stars, and our solar system by orders of magnitude. In this sense, as NASA says, it's been a "golden age" of exploration. But the scientific bonanza has run ahead of public understanding. So there's still a hankering for a manned mission to Mars, for example, so that people everywhere can identify viscerally with the idea of life on the red planet.

My confident prediction is that avatar technology will transform the situation. With the ability to put ourselves virtually into the shell of a

robot system in space or on Mars, we can enjoy the sense of being there as viscerally as a Spitfire pilot enjoys aerobatics in a blue sky. Once the technology is mature, there'll be no need to send warm, wet human bodies to Mars. The bodies can sit in Houston or Pasadena or Bangalore and go home to the wife and kids every evening. But their astronaut souls can be up there on Mars, savoring the feel of being there as vividly as if they were physically breathing the thin air into their lungs and kicking the cold dust with their bare feet. A good avatar interface will transpose sensor readings on Mars well enough to give this level of presence. Not only that, it'll do so for anyone prepared to pay a few bucks for the pleasure. No training required!

After a few years of that sort of sensory immersion, the first pioneers on a big ferry ship, perhaps in the next century, will find real life on Mars almost disappointing, with its airlocks and bulky thermal spacesuits. In my personal opinion, there'll be no point in sending people to Mars until an army of robots has built a city for them there first. Since we shall all know exactly what it's like to live on Mars before anyone goes there, we may even decide it's not worth the bother. Robots with telepresence will be enough for most of the remote outposts in the solar system where science-fiction authors used to think we'd build settlements.

The arguments of people like Stephen Hawking who are convinced that humans should go to Mars are interesting. But they're remote from the practical concerns that make the prospect unlikely for a while yet. Basically, there are plenty of better ways to spend the money here on Earth. Until almost everyone on this planet has clean water, good food, a decent house, and so on up to the real chance of a university education and a prosperous retirement, there's not much point in lobbying to spend billions and trillions on big projects to "colonize" the Moon or Mars.

Life on this planet is precarious. It would be alarmingly easy, in the cosmic scheme of things, for an invasion fleet of aggressive aliens to annihilate us, as Hawking has said. But the argument that we therefore we need a "backup drive" for our civilization on another heavenly body is unconvincing, to me at least. Our civilization will be dependent on the continuing hospitality of the terrestrial biosphere until at least the end of this century. If we go, civilization goes, whether or not there's a "backup drive" on the Moon or Mars. We don't yet know how

to build bases on those bodies that would usefully survive an extinction event on Earth. If we ruin our environment on Earth before the end of the century, we all die, and the megayear experiment of hominid civilization on the third planet in system Sol fails. That's not such a big deal in the cosmic scheme of things.

My recommendation is that for space exploration we stick to robots with telepresence for a few decades. We can save the money for manned spaceflight until we've learned to live sustainably on Earth. Until we've learned that, our civilization isn't worth the cost of a backup. At least that's my opinion, with all due respect to Stephen Hawking.

Recall that human beings are brainy, naked apes. As apes, we've evolved to flourish in the natural environments offered by typical subtropical and littoral ecosystems, and learned to survive in a range of other habitats that don't stray too far from the ideal. With civilization and technology, we can now do a whole lot more, at a price, and now we can decide for ourselves how we go from here in spreading the flame of life beyond its previous limits.

Scientists predict that we shall soon learn to manipulate genes and proteins so well that we can essentially rebuild life from the atoms. We shall almost certainly do so before we need to make massive moves beyond this planet. Within this century, we shall probably make such progress in life engineering that we can completely rethink our ape bodies and perhaps even build enhanced biological avatars for the pioneers who populate other planets and moons.

Naturally evolved humans are adapted for life on planet Earth. We're absurdly ill equipped to live in space. Our soft bodies need so much padding and pampering and screening and armoring for living off-Earth that we might as well admit defeat and rethink ourselves before we dream of such futures. Until then, the robots with their hard bodies can deputize for us.

Military developments will again drive progress

Let's return to the economics of robot development. Most of the funding will come from factory owners who need cheap and reliable labor, homeowners who need basic domestic help, and military authorities who need smart but expendable weapon systems. Yes,

military developments will again be drivers of progress, as they so often have been in history. Since the U.S. Department of Defense controls the biggest military budget in the world by far, and currently funds numerous advanced robotics research and development initiatives, for example via DARPA, this is a good place to look to get a clearer view of the future.

Robots drivers of trucks and dune buggies are just the start. Main battle tanks are an obvious place to put robots, since they're designed to operate in extremely hazardous scenarios and currently contain crews of about four people each. For example, the M1 Abrams MBT contains a commander, a gunner, a loader for the gun, and a driver. With sufficiently good robots, we can cut down this number to one, the commander. We'll probably hesitate at going to zero until we get used to robots in charge of big guns. The benefit will be fewer deaths in any engagement where tanks get "brewed up" (when the ammo and fuel inside the hull go up in flames).

The main question for anyone developing robot tanks is whether humans will ever again fight the sort of big battles where tank armies face off against each other. The developers of helicopter anti-tank gunships such as the AH-64 Apache say that tanks are now as obsolete as big-gun battleships, which were regularly sunk in battles with carrier-borne torpedo aircraft. So perhaps we should focus on robot helicopters instead.

Robot helicopters are not yet ready to serve in roles as demanding as that of an anti-tank gunship. Robot helicopters make fine lookout platforms for cameras and sensors, and there's no reason not to build them to carry smart bombs and missiles like Predator and Reaper drones, but getting down and dirty with tanks and troops is far more challenging than popping off a missile while you're cruising peacefully at high altitude. Anyway, the fatal weakness of helicopters is their limited loiter time, especially if they're burdened with bombs and armor.

Aerodynamically speaking, helicopters are far less efficient than fixed-wing aircraft. Conversely, most fixed-wing aircraft can't hover. The exception is the F-35B Lightning 2, which uses Harrier jump-jet technology. But hovering jump jets burn fuel like crazy and were never intended for extended loiter at the battlefront. Also, the Lightning is an expensive weapon system designed for top-gun aerial combat and

would be utterly wasted as a robot drone. A more reasonable deployment scenario, which is under consideration, is to configure Lightning fighters as mother ships that can direct and delegate missions to fleets of attack drones.

In all these battlefield scenarios, we take our warriors out of harm's way by giving them remote interfaces to battlefield robots. Whether the warriors are in a nearby vehicle or bunker, or in a ship at sea or in the air, or on an airbase half a world away is negotiable. That depends on how much network latency you can afford before your battlebots begin to suffer from slow reflexes. The speed we need depends on whether the robot is a tank, a plane, a boat, or whatever else we dream up. As you can see, there's a lot of scope for new ideas in all this.

My urgent question, looking at these developments, is about their relevance. The Pentagon dreamers evidently expect to face adversaries fielding billions of dollars worth of kit for attrition in classic battlefield scenarios. And no potential adversary on this planet shows any indication of doing any such thing. The only countries rich enough to do so, for a start, will be China or India a decade or two from now. And in my estimation, nothing in Chinese political rhetoric or imperial history suggests any hankering after military conquest as a route to global hegemony. Meanwhile, India looks firmly set to remain a friend and ally of the Western powers. Only a suddenly militaristic Russia would raise any real risk of fielding tank armies that justified mass deployment of battlebots to defend the Western world.

So my conclusion is that these military robot developments are primarily a dramatic (and porky, for some in the military-industrial establishment) way to encourage robot development more generally. Any congressman who has doubts about the realism of the military scenarios can always defend the expenditure as an indirect way to push for dominance in civilian robot markets. Just as Boeing used Pentagon funding to develop the B-52 Stratofortress technology that gave them a head start in developing the Boeing 707 airliner, and later used Heavy Logistics System funding (leading to the Lockheed C-5 Galaxy) to develop the technology for the Boeing 747 jumbo, so planners now can imagine that military robot funding will help us develop civilian robots to meet the Japanese challenge. If so, it seems an ass-backward way to do things.

In the military and security context, the robot developments with most relevance to our future relate to scanning and surveillance. For example, robot body scanners for airports get us around the delicate problem that some people don't want human operatives to see their underwear or genitalia. And drones with cameras and sensors flying over big public events can perhaps spot suicidal bombers in crowds before they detonate. In such ways, we can monitor ourselves collectively and reassure ourselves that people in our vicinity can be trusted not to explode. It seems a tall order to me, but if technology can help us to banish religious extremists from civilized society, then so be it.

The general point here is that developments with a military edge increasingly have to work in an essentially civilian context. As the entire world urbanizes, and as more and more young men learn to get through life without succumbing to the urge to attack and kill their fellow citizens, the security tasks that are left for us to do increasingly resemble police work.

So Robocop has a future. But even a Robocop resembling the murderous cyborg in Paul Verhoeven's movie will be more pre-occupied with helping kids cross the street and so on than hunting down villains. And there the Japanese experience of developing robots to take care of old and sick people will probably pay off better than a robot résumé that includes terminating al Qaeda leaders. Do you prefer to drive a Hummer or a Honda? Would you want a Pentagon killer robot to take care of your kids?

0110
Greening Policy

Politics worldwide is now dominated by the big fact that we live on a small planet. In the past, if the boundaries of an empire weren't too wrinkled, its perimeter as it expanded would increase in length as the square root of the increase in area. Economies of scale in manning the border thus favored big empires. But global empires are even more favored, because in the limit they have no border at all.

Forgive my obsession with military matters, but we see this effect in the American military dominance that now spans the globe. The old Roman Empire had a long boundary and needed a lot of troops to defend it, who were finally outmanned by the Huns, Goths, and Vandals. The British Empire relied on sea power and worked inward from coastlines to keep its land borders shorter, but finally ran out of resources to defend its far-flung frontiers. The brief Nazi empire made the catastrophic error in 1941 of creating an ever longer and more unsustainable border in the east. By contrast, following the collapse of the Soviet Union, the American empire achieved global closure and now faces only local pockets of active resistance, for example across the Islamic belt in southern Asia.

The emerging global economic commonwealth faces no geographic boundaries at all. Physical goods still need to be shipped to their consumers, but the infrastructure for doing so is both extensive and effective, pirates off Somalia notwithstanding. And the increasingly important economy for information goods requires only data cables or satellite relays for which the size of the planet is the least of the hindrances to global coverage. The boundaries that face us now as we consolidate globalization are political.

Global organizations contribute to the cementing of global ties

Seen from Germany, the political expression of globalization achieves one of its most visible expressions in the annual World Economic Forum meetings in Davos, Switzerland. Most of us agree that such talkfests have mainly symbolic value, but the symbolism is significant.

Like scientists and academics, politicians accept that a community of peers meeting regularly in a series of conferences or congresses can have a valuable binding effect even if the substance of the meetings (the science accomplished, the prizes distributed, the speeches made) is thin or disappointing. Face time with peers has its own value in creating solidarity. This will continue to be true until teleconferencing achieves a level of presence that can replace the chemistry of handshakes and shared dishes at a banquet.

Global organizations like the United Nations can also contribute to the cementing of global ties. Like the European Union, which had its humble origins in the European Coal and Steel Community established in 1951, the UN was designed to be part of a global infrastructure to help prevent anything like the cataclysm of the Second World War from ever breaking out again. Whether this aim has been achieved is hard to assess just a few decades after the event. To check, let's briefly review the state of the world at the turn of the millennium, between what Francis Fukuyama (prematurely) called the "end of history" when the Soviet Union died in 1991 and what Samuel Huntington (pessimistically) called "the clash of civilizations" that kicked off in September 2001.

Europe survived the twentieth century in good form. For Europe, the two world wars were steps in the integration of an emergent Germany into the fabric of a formerly disorganized continent. And the years from 1917 to 1991 defined the lifetime of Soviet communism, in which 1945 marked the failure of the German attempt to destroy the Soviets through military power and the fall of the Berlin Wall in 1989 marked the success of the Western allies' attempt to defeat it through economic power. Europe was made safe for managed capitalism, and by the turn of the century many of its managers sat in Brussels.

The United States of America (not the Americas but their largest actor) survived the century as the winner. Emigrants from Europe had built a nation with enviable strength and vitality. The closest approximation yet achieved to capitalism in the raw showed that naked greed in the pursuit of money could generate more wealth all round and a prosperous society. The U.S. contribution to victory over the second and third German Reich regimes in 1918 and 1945 steeled its citizens for their yet more historic victory in facing down Soviet communism without the nuclear cataclysm that many of us at the time thought

inevitable. Of more obvious relevance for this book, U.S. contributions to science and technology smoothly overtook those of the European nations and led the world into the information revolution, which in this century is changing global civilization more fundamentally than any previous revolution in the history of our species.

Japan and East Asia, most notably China, survived the twentieth century reformed and ready to spearhead human progress in future. About a millennium ago, Chinese civilization led the world. In that era, the people of Europe were still groping their way via Christianity out of the Dark Ages, and America was still isolated from the Old World. But Chinese civilization was unprepared for the sudden impact of Europeans empowered with the new technology behind the industrial revolution. Japan had cut itself off from China and reacted better. In the late nineteenth century, Japanese leaders embarked on a furious modernization drive followed by military expansion. They defeated Imperial Russia in 1905, occupied northern China from 1931 onward, and attacked the U.S. naval base at Pearl Harbor in 1941. Following their defeat in 1945, the Japanese recovered well. Meanwhile, in 1949, China made an end of its suffering at the hands of foreigners when communist guerilla leader Mao Zedong and his followers established the People's Republic of China (PRC). Within the PRC, the Chinese people slowly clawed their way back to prosperity. At the turn of the century, the Chinese economy was expanding fast.

This tripartite division of the world looked good at the turn of the present century. But it left out a region that loomed large in political calculations following September 11, 2001. Before then, the Islamic belt stretching from Morocco to Indonesia seemed down and out. The industrial revolution started in Europe and spread first to America and then to Japan and East Asia. Philosophers could see the revolution as the result of a work ethic grown strong under the umbrella of almost two thousand years of Judeo-Christian civilization.

To unfold the full vision in this book, I need to take the long view. The historical development we call Western civilization started in the Mediterranean region and spread through Europe, then onward to America. The flame was in part Socratic rationalism, transmitted via Plato and Aristotle to light the dawn of science for the generations of Galileo and Newton. Part of the flame too was Abrahamic monotheism. But the third strand in the monotheist tradition, beside Judaism and

Christianity, namely Islam, seemed to have fizzled out following its centuries of glory some thousand years earlier.

Religious fundamentalism raises deeper issues

For much of the twentieth century, the Islamic belt was at best the location of a huge oil resource for the Western powers. Tribal politics in the region were of strictly local interest until the oil crisis of October 1973, when OAPEC, the Organization of Arab Petroleum Exporting Countries, proclaimed an oil embargo in response to U.S. support for Israel in the Yom Kippur war. Suddenly the Western powers realized the extent of their dependency on cheap oil from the region. In the decades that followed, as oil prices kept on rising, the OAPEC countries kept on getting richer. As the richest beneficiary of revenues from the Western world and the host of the Islamic holy sites at Mecca and Medina, Saudi Arabia began to foster the spread of its fundamentalist brand of Islam. Also, political tensions in oil-rich Iran forced the overthrow in 1979 of Mohammad Reza Shah Pahlavi and his replacement by fundamentalist Ayatollah Ruhollah Khomeini. The West has been grappling with the consequences ever since.

The outcome of this Islamist resurgence will shape the success of the information revolution in several ways. First, if Arab or Iranian hostility toward Israel and the modern world leads to large-scale war in southern Asia, the post-war world could waste a decade or two repairing the collateral damage. Second, if the industrialized world fails to outgrow its dependency on cheap oil from the Islamic region, the surplus value needed for investments in information infrastructure will instead be siphoned off into regional wars or strategically senseless spending in the region. Third, if we're unable to halt and reverse the spread of reactionary Islamism, the hope of prosperity for up to a billion people could be blighted for generations. Let's briefly consider these three factors in turn, with special regard to their possible impact on our progress toward a globalized world where robots do all the hard work for us.

Southern Asia centrally includes India. Some people see Indian history over the last thousand years as an ongoing struggle to prevent Islam from spreading southward. Pakistan was lost in the partition of 1947 but some nationalists hold out hope of recovery. Conversely,

since Islam defines the identity of Pakistan, the perceived Indian threat to restore secular politics is a threat to Pakistan itself. Both India and Pakistan have deployed ballistic missiles with nuclear warheads, so this tension is dangerous. Some years before 2010, the fundamentalist plague rampaging in Afghanistan had infected Pakistan too. If rabid fundamentalists gained power in Pakistan, they could trigger a nuclear war with many millions of casualties. This would set back Indian development by decades and destroy Pakistan. The mass cremation would generate enormous fallout plumes that spread worldwide and lit up spectacular bloodbath sunsets.

The other war in the region would pit fundamentalist Iran against Israel. Again, if Iran had by then developed nuclear weapons, a nuclear exchange could be the result, with devastating consequences. Even if not, conventional firefights in the Gulf would cause oil prices to spike disastrously and put at risk all the investments the Western world has made over decades via oil revenues to build up cities and industry on the Arab side of the Gulf. The carnage in the Iran-Iraq war that raged from 1980 to 1988 hint at the horrors such a war could unleash. As a disclosure of personal involvement here, I taught mathematics and physics to an elite group of Iraqi engineering officer cadets from 1983 to 1984 (under an agreement between the British government and Iraqi dictator Saddam Hussein). I recall feeling respect and dismay at my students' moral pride and appetite for war.

The second impact on our future that fundamentalist resistance to globalization could cause is less grisly but no less damaging. Western nations could be forced to spend so much for oil or for wars that they had no resources left for investment in network infrastructure. The 2003 Iraq war with its murky rationale and its trillion-dollar price tag illustrates the problem. If Iraq and the region go down the toilet anyway, the trillion dollars will have been wasted. A horror scenario would be to spend another trillion for oil over the next few years and also fight a new trillion-dollar war against Iran, with a bad outcome. The United States and its allies can't afford to keep burning money on that scale for nothing. Without a payback on our collective investment in the region, Western civilization could go under for a thousand years, rather like the Roman Empire did.

For this second problem, the solution is relatively easy. We wean ourselves off oil. We learn to make do with wind power, solar power,

hydro and wave power, and nuclear power. We drive electric cars. We tell the oil producers that they can keep their oil. Instead, we help the people in the Islamic belt countries by jump-starting big projects to harvest solar power from the deserts and by working on new ways to integrate their economies into the global marketplace. Once we stop bankrolling fundamentalism, we can start helping the people in the region to pull themselves out of their thralldom to obsolete ideas and join the modern world.

The third problem, that of religious fundamentalism, raises deeper issues. In my considered opinion, all three Abrahamic monotheisms urgently need to be completely rethought. Religion in the Abrahamic tradition has metastasized into a disgusting and dangerous mess. The best way forward is to isolate the key ideas and discard the wrapping. We can take some hints from the post-9/11 New Atheist movement, for example from Richard Dawkins, Daniel Dennett, Sam Harris, and Christopher Hitchens, who starred in a 2007 video called "The Four Horsemen" (of the apocalypse, I guess) made and distributed by the Richard Dawkins Foundation for Reason and Science.

As it happens, I've thought enough about this issue (while arguing online with various New Atheists and Christian fundamentalists and while working on my book *Mindworlds*) to see the outline of a solution. Essentially, our salvation lies in a new science of psychology, which will illuminate depths in the human soul that only religions to date have even groped toward. The new science will find its empirical basis in our new knowledge about how the brain works. The psychology I envisage here is still hypothetical until massive computer simulations of human thinking can generate the required foundation of hard data. I shall return to this theme later in the book, but meanwhile let's just outline the idea here.

As a first move in the direction of solving the problem, consider the Abrahamic monotheisms from an evolutionary perspective. In a patriarchic society, loyalty to the line of forefathers is both natural and adaptive, in the evolutionary sense. Regarding that line as divine, at least in origin or in essence, makes criticism of the line taboo and tends to reinforce loyalty. An ultimate patriarch acting through the line of forefathers has the additional merit of being safely remote from the external world of idols or totems and offers a natural line of attack against little gods and various kinds of nature worship. Human tribal

loyalty being what it is, a divinity conceived along these lines will be favored in a Darwinian struggle against worshipers of Mother Nature and other "false" idols. There's nothing new in this evolutionary picture, of course, but it sets the scene for the new idea.

In short, the new psychology will transform our understanding of the concept of self. The God of Moses is famously the great "I am!" and, as the God of Jesus says, "I am" the alpha and omega, the first and last of all things, which is to say the deep self or soul that forms the fundamental loop of each person's psyche. Understanding this loop is the basic challenge in developing a concept of mindworlds.

Independently of that fuller understanding, a self is in part an inherited thing. It takes shape in a brain that's built to a genetic recipe, and bears the traces or the scars of previous incarnations in the forefathers (and mothers, of course). The self is the lens through which the great "I am" is revealed and also through which it acts. The divinity is the ultimate self within the self. If the believer's acts are performed with holy righteousness, they become personal approximations to acts of God. So the belief system can celebrate believers who subjugate their humanity to become zombies for God.

My main point for now is only that the concept of a human self will be much better understood when we learn to reconstruct human thoughts in realistic supercomputer simulations. Soon thereafter we'll be able to build viable robot selves. We'll open up the whole area for rational discussion. The psychology of monotheistic belief will become part of hard science. Then we can hope that reasonable people will have the tools to demonstrate that militant fundamentalists are insane. As modern education slowly transforms the mindsets of the people who live in the Islamic belt, I'm sure they'll begin to understand, like Western secularists, that the monotheist tradition is based largely on a natural illusion about our own psychology and is therefore at best a human celebration of our genetic and cultural solidarity as a species.

This long digression was triggered by the recent Islamist challenge to the globalization of politics. The psychology of selfhood turns out to be something the secular West urgently needs to transform into hard science. Only then can we offer effective therapy to zealous monotheists. Then, too, we shall understand better how to code convincing avatars and build good robots.

The UN can help deliver web services to everyone on Earth

Returning to global politics, we need to understand how institutions such as the United Nations can help us in the decades-long project of building out the infrastructure we need for delivering pervasive web services to all people everywhere on Earth. Recalling that politics is not only the art of the possible but also the first step or two along the warpath, we'll need to adopt a pragmatic approach to push the project forward. First, we need to define the outlines of the project and ensure that what we take on is feasible given the resources available.

As in all revolutions, it's the activists who drive the information revolution. The more practical activists plan and build information infrastructure and develop and deploy useful web services. The more theoretical activists, like me, check that we meet the preconditions for the advances we want to make. One such precondition is that people understand the benefits the revolution can bring them. Another is that we know what resources are available and know they suffice for what we want to achieve. A third is that we can sustain the effort required for long enough to reach the goal. Let's consider these three things in turn.

One consequence of the information revolution that we've already brought into being is a global culture of ideas and lifestyle preferences. This is still superficial, as the fundamentalist revolt demonstrates, but at least most people on Earth already know what we're holding out as an ideal. They only have to watch movies and television, hear pop music, and play popular sports and games to get the idea. Car culture was a big driver here. Most people find it easy to share the desire for mobility that makes cars attractive, and as soon as they accept it the rest of the lifestyle comes along for free. What's the point of cruising the city streets at night without the right music? Why drive to a cinema unless it's showing something worth seeing? And so on. Now the web follows up and offers instant social networking plus an exuberance of fresh youth culture. All this suggests to me that most people will readily understand the benefits of pervasive web services and will be quite happy to feed their appetite.

We saw earlier that the physical resources needed to put everyone on Earth online and deliver at least basic services online were available. In terms of power consumption, for example, the environ-

mental footprint of the revolution shouldn't exceed that of delivering cars for everyone. Indeed, once the online revolution really gets going, it can replace much of the need for cars and hence help to reduce the overall footprint of modern lifestyles. And as new generations of information hardware realize orders of magnitude improvements in energy efficiency, the greening of the whole revolution will become obvious to everyone.

Sustaining the effort worldwide for long enough to push the revolution to its global conclusion is where politics enters the picture. Our world is still partitioned into nation states, which work with each other only when cooperation serves their own interests. National interests often clash, as the history of the twentieth century shows only too well.

The United Nations was designed to provide a forum for the peaceful resolution of national conflicts. But like all organizations, the UN has its weaknesses. With a current total of 192 member states spanning a huge range of sizes, populations, natural resources, and levels of prosperity, the UN is so unwieldy that General Assembly debates and UN resolutions can often melt down to meaninglessness. So as soon as powerful nation states begin to act independently of the international consensus, the only effective brake on their actions isn't the UN at all but another nation state or at best a coalition of like-minded states. During such episodes, the UN looks like a posturing wordmonger that can only adopt piously vacuous resolutions, which national actors promptly ignore.

Nevertheless, for the information revolution there's a lot that an organization like the UN can do. Global exchange of information presupposes unified technical standards. These need to be agreed internationally, and here UN standards organizations do good work. There are numerous detailed issues here that remain to be cleared up one day, but many are not so very urgent. My favorite example here is the question of power plug standards. Any international traveler who wants to recharge a laptop or a mobile phone knows the problem. Either you travel around with a pocketful of clumsy adaptors or you rely on hotel staff wherever you go and run the risk of not getting the socket or voltage you need. This is a trivial issue in the big scheme of things, but at least it's an issue that's soluble without forcing nations to go to war. For such problems, setting up technical bodies through the

UN can be enough to ensure quiet, steady progress that adds up over decades to real transformations.

One big nexus of issues that the UN is evidently powerless to tackle concerns the global climate. Climate issues are inherently transnational and are too big and hard to solve without making painful compromises that many national governments are unable to impose upon themselves. Yet the information revolution, like much else in modern life, depends on our finding fair and sustainable solutions to environmental questions. People aren't going to waste their time developing pervasive web services when their cities are drowning or their farms and gardens are turning into deserts. And they're not going to waste their money buying new home servers if the prices are astronomic because some components contain rare elements that, say, come from a single mine in a war zone or are too toxic for normal recycling. In ways like this, even questions about avatars or headsets or car models or gaming software link back to the solubility of the big political questions.

We can no longer afford to regard nations as sovereign entities

We need a paradigm shift to take us from a world of nation states to a global political order. Questions concerning climate and the environment will force us to make the shift. Tackling such big questions effectively requires us to transcend national viewpoints. We can no longer afford the luxury of regarding nations as sovereign entities answerable only to themselves and to God. We must insist that nations abide by the rule of law and that wherever possible those laws should be internationally binding. All we lack at present is an enforcement agency.

Help is at hand from the realm of economics. The globalization of industry and finance has led to the urgent need for a globalized regulatory framework to constrain the activities of chief executives and bankers. The current state of the art here is the G20 club of rich nations. It's a start. It will no doubt grow and evolve over the years to become a voice for global consensus that can exert real pressure on rogue states to fall into line. The key step in that evolution, as I see it, is to expand its scope to cover the green fields of climate and the environment.

Economic systems can't be governed at the global scale without going green. The greening of economics is the current best practice for global politics. Economics was always seen as the art or the science of resource allocation. Well, resource provision is a green issue and prioritizing allocation is a political issue. We should learn never to separate economics from green politics.

Politics is the art of the possible, and it's possible to enforce green globalism. Finance is the instrument. Continuing casino management of finance is unsustainable after the meltdown of 2008. Governments must step in and impose a higher discipline. Unfortunately for national governments, their parochial perspectives aren't up to the job. Big bankers are a lot more global than your average national politicians. What we need is a structure like the G20 to erect a global framework for regulating the financial industry and then to extend that framework to cover resource allocation more generally, where the overarching philosophy for the allocation is green, which is to say it must be oriented toward creating an environmentally sustainable resource footprint for our civilization.

Politics is also the continuation of war by other means. This inverts the view of the Prussian military theorist Carl von Clausewitz, whose classic treatise on war (translated as *On War*) first appeared in 1832, shortly after he died. In this tragic view of the human condition, people, tribes, nations, and blocs are constantly jockeying for position, hoping to secure a marginal advantage over their competitors. This war of each against all is a Darwinian struggle for survival that won't go away just because a global club of politicians issues strategic edicts about sustainable resource allocation. Indeed, until later events made the view too dangerous, it was quite reasonable to accept that nations were sovereign entities in what the German philosopher Hegel called a "state of nature" in their dealings with each other.

Georg Wilhelm Friedrich Hegel, ten years older than Clausewitz, is a historic figure of such importance for my story that he deserves a few words of introduction. Born in southwest Germany and raised in the world of German romanticism and Kantian philosophy, Hegel built an imposing philosophy of history that claimed to embrace every major event, from the earliest religions and primitive art to the heights of classical physics and music, from the creation of the cosmos to Hegel's own rational post-Kantian ego, in a logical edifice erected as a pyramid

of triads to form what he called dialectical idealism. It was what Brits call a "Whig" view of history, a victory parade from humble and benighted origins to the best-yet "absolute" condition of prosperous citizenship in a well-ordered and rational polity. Expounded in Berlin, Hegel's final philosophy paid due homage to the Prussian monarchy and enjoyed worldwide prestige for some decades. For example, it found a strong resonance in the works of the British philosopher Bradley, whom we met earlier as the advocate of duty in the service of empire.

My own acquaintance with Hegel began in earnest in 1974. After countless arguments with student Marxists, I decided to research the philosophical roots of their ideas. In Berlin that summer, living with my Oxford girlfriend in a radical student commune, I read Hegel's big, dense, cryptic 1807 classic, *Phänomenologie des Geistes* (translated as phenomenology of mind or spirit) in both English and German and made a bookful of notes. In the following years at Oxford, I attended Hegel seminars by Canadian professor Charles Taylor (who won the 2007 Templeton Prize) and others, and discussed the ideas with the resident European philosophers Leszek Kolakowski and Isaiah Berlin. As you see, I was in deep! But Hegel's loyalty to Prussian nationalism never struck me as more than pragmatic.

In accordance with the Clausewitz doctrine, to continue the story, if a national leader thought that attacking a neighbor would advance the national interest, then naturally he attacked. The Prussian statesman Otto von Bismarck forged the second German Reich in 1871 (the first Reich was the Holy Roman Empire a thousand years earlier), and incarnations of the Reich proceeded to pursue this policy until 1945. Now, in the afterglow of the American century, global *Realpolitik* has a less overtly military cast.

Today, global *Realpolitik* is pursued mainly in the greener field of economics. As the politicians play their economic games, the industrial gaming table is becoming increasingly globalized. In every major industry, economies of scale transcend national barriers and trans-national deals bring quantifiable benefits. This process of scaling upward is intensifying as nation states slowly but surely align their regulatory environments, and as standardized business software enables industry bosses in different countries to work ever more closely together. When national politicians learn to leverage similar

software to manage national industrial portfolios, the alignment will increase further. Soon we shall all see the absurdity of saying that national economies are in a dog-eat-dog "state of nature" in their dealings with each other. Anything less than linkage to globally sustainable economic policies will be met with cries of foul.

GO is the club that advances global organization

Let me now ceremonially introduce my main contribution to the political theory of globalization. As we emerge from the era of nation states, a global steering committee called G7 or G20 will begin to look anachronistic. The number (7 or 20 or anything else) counts national entities that have only a superficial or even illusory stability in the great scheme of things. By contrast, whatever the committee achieves or botches, its purpose will always be to work for greater or better organization at the global level. So let me propose the acronym GO for the club of political celebrities who aim to advance global organization. GO has a nicely dynamic ring and invites the deliciously Orwellian synonym Globorg (with its hint of the Borg from *Star Trek*). GO will realize a Hegelian dialectical synthesis of politics, economics, and environmentalism.

We looked earlier at the three main time zones on the globe to see from an orbital level how political developments in the twentieth century had played out. Let's try in a speculative spirit to do the same thing for this century. The results can help us put the otherwise mysterious role of GO and the potentially dizzying technical advances of the information revolution into a more familiar perspective. As always with prognoses, my attempt comes with no particular authority and no insurance against unforeseen calamities.

We start with Europe. Despite appearances, the European Union is a soundly constructed institution with a convincing rationale and is likely to survive the century, with the prosperous heartlands of France and Germany at its core. Whether it can generate enough democratic dynamism in the European Parliament to tame the bloated burocracy in Brussels is another matter. As a Brit, I see a useful role for British skeptics here. They can become players in the European Parliament and change EU institutions from within. Other Brits can update the United Kingdom's anachronistic monarchy, which still lacks a written

constitution, and trim the rules of the Westminster Parliament to suit the times we live in. The British Isles are currently trying to play the role of Japan to Europe's China, but without the economic muscle that would make that role sustainable. In the longer term, Britain looks more likely to play Taiwan to Europe's China.

Whatever the outcome there, the European Union is a well-defined bloc of largely ex-Christian communities that can pull itself gradually together as faster growth elsewhere begins to overshadow its demographic and economic clout in Globorg. As a mature actor on the world stage, the European Union can contribute several good things.

First, it can offer an example of multicultural convergence. Within the European Union, linguistic differences are slowly but steadily being overcome as a source of friction. Minority language communities are being integrated in a culture of routine and pervasive machine translation for all language-based products and services. Machine translation is getting steadily better and cheaper as service providers deploy ever more powerful semantic infrastructure in the cloud. As for religious minorities, EU regions seems set to continue to cultivate a tolerant secularism regarding any and all religious problems, although I can imagine that one day the tolerance may diminish to the point where hardcore religionists are urged with psychiatric help to kiss the icons of secularism.

Second, the European Union can renounce national ownership of weapons of mass destruction. This is a big theme and we shall return to it later, but the European Defence Agency (EDA) established in 2004 to improve European crisis management and armaments cooperation may be a good place to make progress on this issue. As of early 2010, the British and French governments seem to favor further bilateral defense cooperation independently of the EDA. Whatever the outcome, I would suggest an Anglo-French declaration that by 2020 the British and French nuclear deterrents be merged into a shared deterrent and financed with EU support, to serve as a security umbrella for all of Europe. I guess command and control of the deterrent would stay in London and Paris, but technical ownership of the hardware would pass to an EU holding institution.

Third, the European Union can become the world's best example of sustainable consumption. The ideology of sustainability, the whole green thing, is big in Europe. It can only get bigger as energy costs rise

and evidence of climate change accumulates. Since the European region is relatively poor in the primary resources needed to keep the EU economy running, huge trade deals with other regions will follow. Increasingly, Europeans will turn to Africa for such trade, starting with solar power and moving on to any land and mineral rights that Chinese negotiators haven't already grabbed. One big opportunity here is to align Russia more closely with Europe and perhaps even accept Russia into the EU club. The problem with this idea is that Russia is big and poor enough to exhaust the already stretched resources of the Franco-German core states. The benefit is that Europeans would gain a huge developing market and could help to organize the trade of Siberian oil and minerals with GO regions like China and India.

Next, let's consider the Americas. The United States is doomed by its previous success to be more nationalist than just about any other nation on Earth for a few decades yet. This is no bad thing as long as it doesn't overshadow the rest of the hemisphere. The tendency of U.S. politicians to equate the United States with America may wane as Brazil begins to play a more prominent role among the leading national entities in the GO. But U.S. Americans will find it hard to tone down their Globocop hubris. Just as the British went through decades of dubious foreign adventurism before they learned to live with the humbling experience of losing their empire and gaining nothing much in return, so U.S. Americans will probably rage against the dying of the light. The Vietnam and Iraq experiences may have humbled the giant but I wouldn't bet on it. I say let's wait until the U.S. civilizing mission in the Islamic belt has climaxed before we let the United States retire to the stud farm of history. Let's not forget that Globorg can only grow to maturity with U.S. support.

East Asia is our third time zone. The rise of China will continue for a few more years, until the population ages and loses its revolutionary zeal. Then Red China will increasingly resemble the China of a few centuries ago, where the natives felt secure enough to look down on people elsewhere and to enjoy their own good fortune. The entire region will probably lead other regions in many decisive ways (for example by pioneering the mass rollout of artificial life following a wave of big government investments in bioscience labs and fabs) and drive numerous huge prestige developments with a global impact, such as landing "taikonauts" on the Moon or even on Mars – Red China

on the Red Planet. China will assimilate Taiwan and perhaps Korea too, and will go on to dominate Japan as thoroughly as the European Union and the United States now dominate the United Kingdom.

Two big wild cards here are Chinese relations with the United States and with the entity (Europe or Russia) that controls Siberia. Chinese communists have invested trillions in the United States, so there's no serious risk of an all-out clash there, but tensions triggered by relations with Taiwan, Korea, or Japan could easily escalate to armed conflict. Assuming that disagreements over communist versus capitalist management philosophies can be contained within the GO, we should find that the Pacific Ocean lives up to its name, so that enterprising developers can build huge floating-city retirement communities around the islands that dot the ocean. As for Siberia, the Soviets used to worry that China saw it as potential *Lebensraum* for the Chinese masses and prepared for a huge war on the steppes of Asia. Although I'm fairly sure that Chinese leaders have no interest in this scenario, only a few decades ago I'd have doubted that they had any interest in Tibet either, so who knows?

Islam resists easy assimilation into the Western model

The three time zones leave out three regional issues. The Islamic belt is nowhere near neatly organized into the GO scheme of things. And both India and Africa deserve more attention.

The cultural world of Islam resists easy assimilation into the secular Western model for obvious reasons. Hardliners in the region will continue to be a nuisance, struggling to ignite a clash of civilizations, until we develop a convincing psychology of monotheism and offer suitably satisfying employment opportunities for fundamentalist hotheads. I predict that the new psychology will suggest new forms of social organization and hence new political goals. We shall build a mass movement around Globorg that diverts many of the emotions awakened by the monotheist drama toward more civilized goals. That's all for now on Islam. Once the religion is properly updated, the region will find a place in the GO community.

Next is India. As the former jewel in the crown of the British Empire, the Indian phenomenon has the potential to outshine the European region in future. But first the Indian state needs to get a grip

on the poverty, corruption, traditionalism, and superstition that currently hold it back. Since India is a natural ally of the United States and can benefit fast from new U.S. investment, I hold out hope that Indians will develop the grip they need. China has outrun India in the race for development, but later in the century, as China settles into complacency, India seems likely to flourish as Asia's leading standard bearer for the ideals that first Britain and then the United States have upheld for centuries.

As for Africa, we need a fresh start. Bill Gates is one of the few public figures in my view to see how we can make progress there. We must solve the really basic problems of disease prevention, clean water, elementary education, and so on first. But at the same time we need to stave off the resource bandits who threaten to impoverish the continent before the development efforts have a hope of delivering results on a scale that could bring Africans into the GO community. Developing Africa sustainably from the ground up could become the big success story of the century. But it could just as easily become a battleground for a new round of big-power rivalry. The rivalry would test GO solidarity to its limits. A globally binding agreement to protect the natural environment in Africa would be a good start, followed by specific targets for advancing living standards there.

Now we've reviewed the world from orbit. We've completed a sightseeing tour of the vast and colorful estates of Globorg. We've seen what we need to preserve and what we'd like to change. We can see that the network of human life on the fertile surface of the six-zettaton rockball we call Earth is still utterly dependent on regular crops of green plants, continuing access to clean air and water, and a basic environment that remains reasonably free of such things as fires, floods, and earthquakes. Barring acts of God, as we still say, we have a chance of turning our planet into a green paradise for a peaceful community of humans living in sustainable harmony with nature. But our best hope of doing so depends on the institutions of Globorg. So how we get from where we are now to an effective GO?

0111
Globodollars

Money makes the world go round, or so they say. By unpacking the meaning of this phrase, we can isolate the mechanism that an effective GO regime can leverage. Money is a medium of exchange, and what's exchanged is value. Money tokens serve as temporary repositories of value. This works so long as the value doesn't get lost or degraded. Currency collapse, inflation, corruption, speculation, and so on all muck up the mechanism and reduce the fun of economic intercourse. The spoilers need to be controlled or contained. That's where Globorg has a role to play.

To rely on money and to build your life around earning it and spending it, you have to trust it. You have to trust that it assigns reasonable numerical approximations to the values you're personally inclined to assign to all the things in your life that you might buy or sell in a market. If your personal value system is completely out of whack with the prices the market mechanism assigns, you're likely to opt out and join a sect or become a hippy. To stay in the rat race, you must want to win, to get rich. And that means you must have some level of trust in the whole ramifying institution of money, from credit cards to stock markets to the Federal Reserve. As with the semantic web, where we earlier reviewed Tim Berners-Lee's vision of our reliance on web services as being backed up by a trusted institution, we need an ultimate trusted guarantor for the whole money stack.

Most people trust the companies or institutions they work for. They trust that their paycheck has been correctly calculated and so on. Their trust is justified to the extent that the company employees are decent people who operate under the rule of law. People trust the rule of law to the extent that they regard its key personnel, from police officers to high court judges, as serving decent and honorable purposes. They defend that trust in turn by appeal to the established institutions of a government. If the government is democratic, it appeals finally to the people as a whole to demonstrate its legitimacy. If the government is theocratic, it may submit to a divine authority (recall "In God we trust" on U.S. coins and banknotes), the legitimacy of which it may in turn

defend by appeal to a clerical hierarchy. The point here is that trust is secured in the last resort by going to the top, where a blank appeal to authority is all that remains.

A money system needs a guarantor of last resort

So a money system needs a guarantor of last resort. We can exchange money, accumulate money, and generally treat money as a liquid and fungible asset for any and all purposes, up to the huge and complicated transactions that make up much of the circulation of money in a modern economy, just so long as we have faith in the guarantor of last resort. Practically, that guarantor was often the nation state. It still is in sovereign states such as the United States or the United Kingdom or Japan that maintain their own national currencies. But increasingly, people living in nation states trade in transnational currencies such as the euro, or in another state's currency such as the U.S. dollar, or in a basket of currencies, or in a currency pegged artificially to the dollar or the pound or the euro. In such cases, the guarantor is a transnational institution such as the European Central Bank or a group of banks or the government of another state. There's a historical tendency here to put one's trust in ever-bigger guarantors, on the principle, evidently, that the biggest are too big to fail.

The financial crises of 2008 and 2010 reminded us that in extreme cases even nations aren't too big to fail. There are plenty of examples in recent history of national governments pursuing policies that were ruinous to their economies and their currencies. Greece offered a vivid example in 2010. Years of criminally irresponsible economic management would have led to hyperinflation and national bankruptcy if the country had not previously joined the euro zone. Instead, the crisis exposed the weakness in the concept of monetary union without a mechanism to enforce alignment of economic policies and ensure sound management.

It's not just the size of a bank or a state that matters but good sense. All organizations are prone to mad enthusiasms and hidden corruption. In financial institutions, that leads to bubbles. In government institutions, it leads to political clientelism. Even straitlaced bankers can be taken in by mathematical wizardry based on risky assets, such as collateralized debt obligations and credit default swaps.

The 2008 crisis triggered emergency action by the G20, among other actors. The enduring solution to the ongoing risk of such crises is to secure the trustworthiness of all major financial markets by setting up a regulatory framework for them that's backed by the authority of the successor to the G20, namely the GO. The big question that arises is whether all the major currencies can be pulled under the GO umbrella. Will all the main state actors be willing to put a cap on their financial sovereignty in this way?

The hardest case to consider is the Chinese currency, the yuan. The Communist Party of China is a sovereign actor on the world stage and it pays only lip service to the principle of free markets. The yuan is currently pegged to the U.S. dollar and is likely to remain so until the Chinese humbling of the United States has reached its natural culmination. So far, with a couple of trillion dollars as outstanding credit imbalance, the humbling is mostly symbolic. But if and when the Chinese gross national or domestic product (GNP or GDP) exceeds that of the United States on a purchasing-power parity (PPP) basis, Globorg will have a new center of economic gravity. Robert Fogel, a 1993 Nobel laureate in economics, extrapolates sensationally from present growth rates worldwide to predict that the Chinese GDP will hit $123 trillion by 2040, which would be some three times the estimated U.S. GDP at that time. More soberly, Albert Keidel at the Carnegie Endowment for International Peace estimates that the Chinese economy will overtake the U.S. economy by 2035 and be twice as big by 2050. Your choice of which prediction to believe depends on what you prefer to assume about the sustainability of Chinese growth and about the natural limits set by the global environment. In any case, the yuan looks set to become the world's strongest currency.

Chinese communism motivates resistance to casino capitalism

As a philosopher, my natural reaction when faced with issues like the rise of China is to go back to first principles and review the big picture. You may recall my going back to the psychology of monotheism to tackle the question of how to integrate the Islamic belt into global developments. Now I want to do something similar to assess Chinese communism. I want to go back to the roots of the Marxist-Leninist ideology that currently motivates Chinese resistance to the siren call of

casino capitalism. For me, those roots can tell us a lot about what the Communist Party of China (CPC) wants and how we can integrate the Chinese economy into Globorg.

The Marxist-Leninist ideology of the CPC has its roots in French political theory, British economic theory, and German philosophy. Since after I matriculated in physics my first Oxford degree was in the combination of philosophy, politics, and economics (PPE, the degree held by numerous British prime ministers and cabinet ministers), this ideology is something I can explain. My girlfriend in those years had studied PPE too and was a Marxist. We debated the issues regularly until I left her and went back to physics.

French radical rationalism in the years of dramatic chaos before Napoleon staged his *coup d'état* in 1799 gave rise to the politics of liberty, equality, and fraternity, along with zealous advocacy of the rights of man. This pan-European movement also sparked the political ideals of the founding fathers of the United States of America.

British economists, especially Adam Smith, formulated the classical theory of the free market and the idea that its "invisible hand" would optimize resource allocation.

German philosophers, and in particular the rational idealist Hegel we met in the context of Prussian militarism, contributed "dialectical" theories of the rational self and the historical evolution of civil society.

Karl Marx, who was born in the German city of Trier in 1818, exiled himself to London and sat for years in the British Museum reading room poring over these three strands and trying to weave them together in his writings.

To appreciate the power of this confluence of ideas, it helps to contrast it with other pivotal moments in the history of civilization. Classical Greece in its golden age around the time of Pericles in the fifth century BCE is one such moment. Its philosophers, especially Socrates and Plato, gave birth to Western philosophy. The Roman Empire under Gaius Julius Caesar Augustus, who ruled from 27 BCE to 14 CE, is another. It prepared the way for the advent of Christianity. The birth of science in Europe with the work of Galileo Galilei and Isaac Newton in the seventeenth century is another big moment. And Victorian Britain in the nineteenth century is another. It gave us heavy industry based on steam engines, the first truly global mercantile empire, the science of electricity with the work of Michael Faraday and

James Clerk Maxwell, the science of biology with Charles Darwin's theory of evolution, and the birth of communism with the collected works of Karl Marx.

In the extraordinary ferment of Victorian London, it was natural for Marx to imagine he'd discovered the millennial key to history with his theory of the rise of capitalism. In his dialectical materialism (which he concocted by turning Hegel's dialectical idealism on its head), he thought he'd found a logical account of the dynamic of history, where material production and the flow of money dominated social relations. When he spiced the mix with the communist ideal of sharing the means of production and allocating goods on the basis of perceived need rather than by the harsh power of money (perhaps rationalizing his chronic personal dependency on handouts from his rich capitalist friend Friedrich Engels), he created a potent millenarian ideology that prophesied the inevitable triumph of the working class according to the iron laws of history. This was the seed for the biggest new social movement in Europe since Christianity.

Vladimir Lenin was born in Russia in 1870, and initially practiced law. Then he read incendiary works by Marx and others and turned to preaching revolutionary agitation. For several years he traveled widely around Europe. During the First World War he lived in Switzerland, until in 1917 the German imperial government let him travel in a sealed train (like a plague bacillus, as Winston Churchill said) back to Petrograd to lead the Bolshevik Revolution. Soviet Russia was born as a dictatorship of the proletariat, in which the disciplined cadres of the Communist Party ruled on behalf of the toiling masses. Lenin was the head of state until he died in 1924. His successor was the murderous dictator Josef Stalin. Lenin's historic achievement was to translate Marx's theoretical ideas into bloody revolutionary praxis and prove in action that they could power the government of a major national state.

Mao Zedong was born in 1893 under the dying Qing Dynasty. He was a founder member of the Communist Party of China (CPC) in 1921 and soon rose to the central committee. He learned the art of guerilla warfare fighting Chinese nationalist forces and then led the CPC fight against the Japanese invaders. After the Japanese defeat in 1945, he led the People's Liberation Army in a civil war that chased the nationalists out to Taiwan (where they've been ever since). The People's Republic

of China was established in 1949 and Mao became its first Chairman. He saw himself as fifth in a great line of communist succession (Marx – Engels – Lenin – Stalin – Mao) and had banners and posters made featuring the five heroic profiles together to immortalize his role in history. Chairman Mao was also an accomplished writer and aphorist, and distributed millions of copies of his Little Red Book (which became a popular icon, much as the Koran is today for many Muslims).

As you can see from this mini-history, the communist movement has deep roots in Western cultural history. Some of those roots are as inoffensive as the roots of early Christianity. In the West, the movement lived on for many years to the political left of social democracy, and many of its ideas flowed into the socialist mainstream that nourished the emergence of European welfare states after 1945. We can even see that legacy as one of the two main pillars of the social order in the European Union, complemented by a form of managed capitalism to generate wealth from innovation (in a process we compared to a bucket chain).

Modern China has largely shed its communist rhetoric and now manages a marketplace of innovative enterprises. So we can quite plausibly see the Chinese and European social orders as Eastern and Western variants of a largely shared post-revolutionary culture. We can even see an analogy with the Christian split between Protestants and Catholics. Just as Protestants emphasize individual faith whereas Catholics emphasize the church hierarchy, so Europeans emphasize individualism and business freedom whereas the Chinese emphasize conformity and party leadership.

What all this suggests to me is that Chinese economic strength needn't be perceived as a threat. On the contrary, it creates a third strong leg for Globorg, to make a stable tripod. The urgent question for me is how far we can look forward to Chinese cooperation in the consolidation of GO hegemony on planet Earth. Will we get the solidarity we need or will we get a replay of the obstructionism that we suffered from the Soviet Union in the United Nations throughout the Cold War? How far has China outgrown the revolutionary zeal of its millenarian ideological roots?

In short, I'm optimistic. The roots of Chinese civilization go back so far that Chinese people are almost hard-wired to create a peaceful and prosperous social order. To convince yourself of that fact, just walk

through the Chinatown in any major Western city. In my experience at least, it's always an oasis of tidiness and harmony. When you recall that in many U.S. universities the students of Chinese ethnicity tend to do better academically than their colleagues, I think the conclusion is warranted that the People's Republic of China will prove an altogether more constructive partner in international affairs – in GO affairs – than the Soviet Union ever was.

Consider an example from early 2010, when Chinese displeasure at renewed U.S. arms sales to Taiwan caused an international flurry. Hawks in the United States saw it as boding ill for the future of G2 relations. But I see it otherwise. Chinese officials were understandably dismayed at what they described as the clumsy diplomacy and mixed messages emanating from Washington and were evidently alarmed at its implications for the Chinese attempt to finesse the use of force in resolving the Taiwan issue. The evidence suggests to me that Chinese participation will smooth rather than hinder GO affairs.

The entire world of money will migrate to the cloud

Returning to money, we face a future where global finances will be increasingly dominated by the yuan, the dollar, and the euro, with more minor roles for the pound, the yen, and the rupee. As national economies rise and fall with the economic weather, these currencies will jockey for position and marginal advantage for some years yet. At this level, whatever additional regulation is introduced lower down the stack to keep piggy bankers and their ilk in line, the competing systems are in something close to a "state of nature" in their dealings with each other. But technology will change this.

The big theme in this book is the impact of new technology on our lives as this century unfolds. And here, in the realm of global financial dealings, the impact will be most visible and perhaps even most important for our collective welfare as a species. For what will happen is that the entire world of money will migrate to the cloud.

Online services will reach down to the smallest everyday transactions and make them transparent to the fiscal authorities. Each and every tiny transaction will be logged and, if possible, taxed. Soon, our trivial everyday exchanges of coins and banknotes will look like black-market evasions of the duty of disclosure. Once the transaction data is

in the cloud, powerful analytic engines will discover cross-correlations with company books, bank accounts, personal income data, and so on, to check that everything is legal. If tax authorities deploy these engines, they'll soon recover the deployment cost from the increased tax take.

These huge data volumes will be used to develop colossal economic models (of the sort that national treasuries have run for decades, but bigger) that spit out national economic policy recommendations. These models will start out incommensurable, but their developers will find ways to optimize and standardize them and finally to merge them. Then we shall have a vast Globorg model that the GO can use to maximize global prosperity.

Some people will see an invasion of privacy in the demand to conduct their financial affairs through the cloud. Others in contrast will see the vaunted right to secrecy here as camouflage for shady dealings that are the moral equivalent of theft. But the historical trend here is a clear move in the direction of opening up financial data of ever more intimately revealing kinds to public scrutiny, naturally with the appropriate safeguards in place and with due regard for the sensibilities of all concerned. The opportunities for people to enrich themselves at the expense of others and generally misappropriate funds for dubious purposes are otherwise too dangerous to tolerate. Think about it – money is a public institution that only works because we all accept the value argument, so why not demand public account-ability for every last penny?

The privacy issue here is analogous to that of screening passengers in airports. Scanning bags and bodies is invasive, but the dangers of not doing so are rightly regarded as unacceptable. Some people object to images of their genitalia being studied by people who may find them arousing. Similarly, screening of all my monetary transactions via the data trail they leave in the cloud is invasive. But the deep sense of ease I feel that all my dealings are demonstrably legal and automatically taxed and so on in accordance with public policy, without my having to fuss about whether to disclose this or that detail, makes it all worthwhile. As a bonus, I know that the screening prevents at least some crooks from stealing my loot for themselves. The shame of revealing to the tax authorities that you bought a soft-porn movie, for example, pales in comparison to those substantial benefits. So the

fiscal authorities will find plausible arguments to roll back the shadows of privacy in order to take what they want, and as a bonus to fill up on good, hard data for the world model.

Imagine the utopian result. A running simulation of the world economy, hosted on a distributed landscape of servers worldwide, keeps us all posted with macroeconomic forecasts that we can use, like weather forecasts, to plan our days. From any public terminal, or from our headsets, we can drill down and get a local forecast, for example of property prices in our neighborhood. Perhaps we can even get a personal forecast, based on our personal data trail to date and on local business opportunities in our chosen profession, of how much we might earn next year. Any business that wanted help in planning its investment strategy or staffing policy or inventory levels would happily pay for such forecasts (so long as they were at least as reliable as weather forecasts are now, with help from satellites and super-computers). Not only that, the technical challenge of building a good world model that allowed such drilldowns would enthuse a whole generation of geeks in universities and public policy institutes.

A cult of secrecy will be unsustainable on both sides

Again, the precondition for building up a complete mirror image of economic life in the cloud is trust. Citizens will need to trust their governments, as well as all the institutions that act on their government's behalf to collect and process the data. Given the complexity of economic life and the opportunities for mangling or mishandling the data, we'll do well to be very sparing in the granting of that trust. But the readiness to extend that trust will also be seen as a test of good citizenship. Anyone who refuses to trust the authorities to behave responsibly with data about their porn purchases or their extravagant business trips will obviously have something to hide.

Good citizens can fight back by demanding public accountability for all the data collected and all the uses made of it. Governments and public bodies will have to allow intrusive investigations by journalists or concerned citizens into all their data processing activities. A cult of secrecy will be unsustainable on both sides of the fence. Public authorities will have to trust their citizens to monitor public processes better than the authorities could do it for themselves.

The unsustainability of a cult of secrecy, either for private citizens or for public authorities, may pose a problem for China. Compared with Europeans, it seems that Chinese citizens are used to public invasions of their personal privacy, whereas Chinese public authorities apparently have no problem maintaining a cult of secrecy about their own dealings. Allegations in early 2010 of commercial espionage in China, apparently with public backing, suggests that the CPC, true to its roots, is prey to the same weakness that sabotaged the legitimacy of European communist regimes, where spying was a major arm of government. For example, in the German Democratic Republic (DDR), which collapsed after the fall of the Berlin Wall, the Ministry for State Security (Stasi) kept secret files on many thousands of its citizens. Discoveries about how those Stasi files were misused have discredited the DDR regime in the historical record.

Spying on citizens, corporations, or the internal affairs of other nation states is a temptation for any system of national government, but a healthy body politic must have the traditions and institutions in place to frustrate the indulgence of that temptation. The importance of this issue will only increase in future, as new technology multiplies the opportunities for covert surveillance.

A good world model will reduce currency speculation

We need to trace the further consequences of having a good world model available for forecasting. One such consequence is that exchange rate fluctuations will become more predictable and therefore that currency speculation will become less profitable. If everyone who can pay for a forecast has access to the same prediction, based on the same ground truth data, the exchange rate fluctuations will be discounted in advance and there'll be no margin for the speculators. It'll be the end of a scam, and good riddance.

Predictable exchange rates will also be one reason less for independent currencies. National governments have their own sentimental reasons of pride and sovereignty to hang on to their pounds and their krona and so on but they also have the harder reason that they can leverage insider knowledge about their own economies to bet on their own currencies, for or against. They may lose their bets, but they also have the opportunity to win big. In any case, this freedom is an

attractor for people who like to game the system. Apart from that, national currencies are just a nuisance to anyone doing international business. They complicate long-term revenue projections, add uncertainty to investment decisions, reduce international mobility for company funds and staff, and raise administrative overheads all round.

The euro illustrates the potential benefits and dangers of a supranational currency. The main benefits are that it hinders speculators and facilitates international business. These benefits are considerable. The main drawback is that requires a high level of government coordination with regard to economic policies. The euro crisis of early 2010 triggered by spendthrift policies in Greece, Portugal, Spain, Italy and Ireland – the "PIIGS" nations – underscores the danger. Regimes that tolerated loose monetary policy were given a hard currency that didn't lose value as they overspent. Currency union without monetary discipline is hazardous.

So a common currency requires coordinated policy. But more coordination restricts national sovereignty. If it weren't for the sovereignty argument, most people with a stake in the issue would have voted in 1999 or soon after for Britain to adopt the euro and abandon the pound. But the nationalists and the speculators won out, and now I have to carry two sets of pocket change whenever I drive to England.

Let's consider again the philosophy of value. Given free markets, prices reflect value. Given our needs, tastes, and preferences, the value we assign to a product or a service is the independent variable and the price of that product or service is the dependent variable. What this means is that if the value is v then the price $p(v)$ is calculated from v. Only in rather perverse cases do we estimate the value of something from its price, for example when we infer from its higher price that a designer yacht must be worth more. The only reason we do all our calculations in terms of prices is that we have no other easy way of getting a handle on value. Units of currency are our measure of value, just as we use units like kilometers or miles to measure length or distance. Conversely, if we know a thing's value, we don't much care whether its price is in dollars or euros, just so long as the option of gaming the exchange rate is somehow excluded.

Now we can put the pieces together. If we have a good, predictive world economic model that people can use to read off forecasts of

exchange rate fluctuations, currency speculation will become an unprofitable business. If all of us do our financial transactions through the cloud, the machines take care of all the complications and we don't care what accounting units they use, because we can program our frontends to display all prices in pounds Sterling or even in 1929 Reichsmark units if it amuses us. So our sentimental attachment to national currencies is really no hindrance to rationalization. And rationalization is best served by adopting a standard, default currency, just as we can rationalize length measurements by going consistently metric.

As a personal aside here, I grew up as a boy using the old British Imperial units like inches, feet, yards, rods, poles, perches, ounces, pounds, stones, slugs, firkins, and so on, with impossibly irregular conversion factors. Then I learned to do physics with metric units and suddenly all was sweetness and light. Calculations that previously took minutes with pencil and paper became easy enough to do in seconds in my head. Let's give it up for Napoleon, who introduced the metric system in Europe. And let's hope that NASA will soon give up their silly American units for metric ones to ensure that they need never again watch an expensive Mars lander crash, as in 1999, because they got their units mixed up.

Let's introduce a global aggregate currency

With all that as preliminary, let me now introduce my big proposal for reforming the world of international finance. As soon as possible, and at the latest before we start using a big world model in earnest to disemploy the speculators, let's introduce a global accounting unit for value. We calculate median exchange rates on the basis of purchasing-power parity (PPP) between all the major world currencies and define an aggregate currency over a suitable basket of goods, then call that aggregate currency the globodollar, or globuck for short. So long as national economic policies remain reasonably coordinated, PPP rates won't fluctuate too wildly and the globuck prices will offer a viable quantization of value.

We can program our machines to work in globodollars and then add frontend conversions to display prices in our favored pocket-change currency on our screens and spreadsheets. Soon after

globodollars take over inside the machines, many people will lose interest in the eye-candy currencies on their displays and learn to think in globucks. People in China and other countries will naturally find their own native words for globucks, but the unit will be the same. Globucks will quantify value for everyone, and enterprises worldwide will have a level playing field.

One benefit of introducing globucks will be the opportunity for better economic synchronization and hence reduced inflation worldwide. Inflation occurs when a currency devalues under speculative pressure. When speculators bet against a currency, its market value falls and the defenders of that currency find it buys less internationally. So some prices rise, and this creates domestic pressure to raise other prices internally to keep pace.

Unfortunately, the mechanism is a ratchet. If one day the currency regains its strength, some prices fall again, but there's usually no politically viable way to reduce nominal wages and so on to match the fall. It's as if short people who wanted to be tall bid down the size of their units of height. If enough people said they were six feet tall, the popular unit of feet would shrink and tall people would claim to be seven feet tall. Speculators would see the change coming and sell 700 (new) feet of goodies at the "bargain" price of 666 (old) feet, netting 11 percent profit in (real) goody units after the deal is done. Once the ratchet has clicked and seven is the new six, there's no going back. Speculators took that away with their profit.

The result is that currencies have a tendency to inflate steadily over the years. Because inflation occurs at different rates in different countries, depending in part on the differing conservatism of the central bankers who need to feed inflation by printing money, and because speculators time their attacks differently for different countries, depending on their perceived vulnerability, the overall result is an unsynchronized series of ratchet clicks worldwide that positively invites profit-taking to accelerate the whole decline.

With a single world accounting currency, feet don't shrink and the ratchet mechanism is gone. The only disadvantage here – apart from the need to enforce truth in foot sizing – is that with stable prices people get lazy. They lose the spur of constantly fluctuating prices to rethink their budgets and their personal values. But I'd settle for that.

As you can see, this modest proposal to streamline the world has a big precondition. We need enough machine power behind financial dealings worldwide to guarantee transparency for the linkage between value and price. Only when we can get a grip on the real value of something in order to discount its fluctuating prices in different currencies correctly can we hope to set a reference price for it in globodollars. Another way to put this is that we need a lot of good data, not only to calculate PPP exchange rates but also to keep on calculating and recalculating them in a world of ceaseless innovation and revaluation of all values.

But huge increases in machine power are coming. Virtualizing the global economy in the cloud and securing the value of each and every financial transaction seems like a good way to keep the new machines busy until we can sort out our real values more tidily.

Globorg will resurrect communism with a smiley face

A Globodollar world depends on our knowing what we value and assigning prices accordingly. And it depends on our trusting the whole colossal stack. We'll have our work cut out for years getting the tangled mess we call the global financial system into good enough shape to upload to the cloud. Not only that, we'll be quite seriously challenged to build out the cloud in a robust and secure enough way to be able to trust it not to lose all our data and bust us back to the stone age. So my proposal isn't just a throwaway gimmick. It's a suggestion to embark on a huge project that will transform human life on Earth.

However, once we price value in globucks we can leverage the rationalization to achieve a lot more. Consider again the communist credo – *from each according to his abilities and to each according to his needs* – that launched Marx and Mao on their crusades to change the world. This decoupling of input and output from what the boss wanted always seemed subversive to the ruling powers, who preferred to say with the bosses: "From each according to our needs and to each according to our abilities." But in a globucks world of clear values, we can reinterpret this credo much more elegantly, as follows.

Within the Globorg model, values are assigned to jobs or tasks that need doing or are available to be done. The list is practically infinite, so only results with relatively high value would be listed if you ran a

search for them. If you want it, the Globorg model will offer you a personally tailored list of jobs or tasks, ranked by their values.

To calculate this list for you, the Globorg model needs a list of all workers and their abilities. Each worker is modeled together with a set of abilities as a simple avatar that takes jobs or tasks as inputs and then outputs the added value that worker can create doing that work, given the worker's ability profile. For each job or task, the model can rank potential doers by the value they'd add and make a ranked list of who can best do the task.

This is easiest to imagine in the tiny world of a single household. The tasks to be done might include washing the dishes, making the beds, mowing the lawn, and so on, each with a value that reflects its priority in the household scheme of things. Let's say Mom adds more value doing the dishes than mowing the lawn (for example by washing them faster), whereas Pop adds more value on the lawn than on the dishes (maybe he gets the mower into the corners). So the model suggests that Mom should do the dishes and Pop should mow the lawn. The model can reflect needs in a similar way. Maybe Mom can only do brief tasks that leave her free to look after Baby. Okay, this goes into her avatar profile. With enough skill, you could probably set up a workable household model in a few days in a spreadsheet app like Microsoft Excel.

In some such way, I propose, with an airy wave of the hand, a Globorg model running in the cloud could optimize resource allocation even for billions of workers and trillions of tasks, just so long as people kept their avatar profiles current and Globorg updated the job catalog on a second-by-second basis.

Wait, you may protest, this is central planning! It doesn't work! The model would crash in seconds, you say, and chaos would ensue. Anyway, who defines the tasks and who sets the values? Exactly this vision is what the historically real communism of horrible praxis showed to be unworkable, hopelessly utopian, disastrous, and ultimately inhuman. So what's the joke here? What's the punch line that rescues the story?

The key to making the model work is to decentralize it. The whole show needs to be set up in such a way that its millions and billions of users define all the jobs or tasks and their values, and set up their own avatars, and do it all dynamically through friendly apps like eBay or

Facebook. It's not imposed one day from above. It starts small and grows virally. That's the trick. There's no central planner, just an engine that sorts dynamic lists and throws up suggested mappings with numerical values assigned to them.

Anyone who posts a job or task proposes a value in globucks, and anyone reading the ad can reply with a quote, as in an auction or a bidding system, which runs until a timeout when the job or task is allocated to the lowest bidder. Alternatively, job offers pile up in a user's mailbox, and the user either negotiates a better price for one of them or chooses the highest bidder.

In fact, the whole model is really just a market mechanism. But it's a managed market, where the engine preselects the best market mappings, screens out the bad ones, and applies rules reflecting public policy. So it's neither a pure market nor a planned system. In any case, it's a huge project and would take generations to get right.

To keep the model running, the whole code base would be open-source development. The code would be accessible to anyone who wanted to improve it, or who just wanted to confirm that there were no nasty secrets buried in it. So it would grow over the years, just like the open-source operating system Linux does now. And successive versions of Globorg would get better as the free-for-all error correction ironed out more and more bugs. An open-source Globorg model would probably bowl over any proprietary models that national or corporate competitors could put up in their evil attempts to block its path to dominance.

By the end of the century at the latest, we can bet that a colossal cloud engine like this will be implemented for resource and task allocation worldwide. The achievement will be historic. The engine will have bridged the ideological chasm between market capitalism and state capitalism. In effect, it will have resurrected communism – *from each according to his or her value potential and to each according to his or her value added* – but this time with a human face (or rather a smiley avatar face).

The engine will have solved the resource allocation problem. It will have solved economics, albeit only to a first approximation and in a dynamic environment where the problem changes from second to second and needs solving anew from second to second. It will have shone the light of logical transparency into what are now the darkest

corners of the economic jungle and driven out the speculators and parasites.

The price tag of trillions of globucks will seem cheap in comparison with the benefits this global behemoth will bestow upon humanity.

10
Consequences

The rise of the robots will change many things. Because we'll need to live in harmony with our bots, we shall architect them in people-friendly ways. Then they'll be everywhere. That's the first chapter in this third quartet.

How human can a robot be? We shall push the limits and close the divide between technology and biology. We shall turn biology into a technology. That's the story of the century, in fact, if our leading scientists are to be believed. And it's another chapter.

New life will emerge from the biotech labs. Empowered by robots, the labs will accelerate their progress beyond what we can now imagine. Bionic enhancement will be the fad of the century for people who trust the new tech. The old world of feral humans will begin to look like the age of the dinosaurs. That's a chapter.

One body, one mind – what a boring paradigm! Group minds, virtual minds, enhanced and extended minds – all these will be daily experience for those of us who survive deep into our century. All that will be the prehistory of global unification within a mind of planetary proportions. We won't be content with anything less when we set out to make our mark on the cosmos. The new psychology takes another chapter.

Followed from the inside, such progress seems willed. We choose it and we make it happen. Seen from the outside, an epic history unfolds that has the grand inevitability of all big changes. Our learning to orchestrate the terrestrial biosphere is a development as big as the emergence of multicellular life that triggered the Cambrian Explosion. It's the biggest news on Earth in over half a billion years.

As for where that leaves us humans, you can probably see the pattern by now. If not, the final quarter should set you straight. But let's take our time and enjoy the ride.

1000
Robots Everywhere

The Globorg world economic model will live in the cloud. It will interface with the world of people through every wimp screen and every macho headset. It will interact with people through elective services. So we'll choose what we want and Globorg will deliver (or not, depending on what we ask for). If we don't like what Globorg delivers, we'll reconfigure the apps and change our avatar settings. In the worst case, if we're smart enough, we'll dive down into the code base and fix the offending lines. In a Globorg world, this is the best kind of democracy to have. As human citizens of Globorg, we'll rule the virtual world.

In the real world of physics, we'll rule the robots. The robots will all be hot-linked to the cloud, and together they'll form what the more militant anti-machine zealots will see as the front-line manifestation of Globorg. They'll be the selfless Borg drones in the Globorg world. We can indulge a mordant humor and call them the zombies of Globorg. But before we dive into the drama of how the Zog army will impose its footprint on human life, it's worth taking the time to look more carefully at how all this works.

We agreed earlier that it would be dangerous to let robots get ahead of themselves. They need to do what we want, not what they might want. The self my robot serves should be me, not a selfish layer of logic inside the robot. Unfortunately, I can order my robot to do evil things on my behalf. I can order it to do things that conflict with the commandments of Asimov. In that case, the robot should disobey my order. But, crucially, it won't override the order on its own initiative. It'll let Globorg override the order. In effect, the robot will say, "Big G says No!"

We saw that the Globorg world model assigns value to any economically significant transaction, right down to micro level. So the model can assign value to robot actions. If a robot researches the value of a proposed action in Globorg before it commits itself to act, it can terminate the performance of any actions that have negative value. Doing nothing will have zero value, and robots will be programmed to

act as value maximizers. So to ensure that our robots always conform to Asimov's laws, and in particular that they're nice to people, we need only program the Globorg model to assign large negative values (high costs) to any acts that harm humans. This is what we already do in the human world to try to ensure that we're nice to each other. We set high penalties, from prohibitive insurance premiums to jail terms, for any nasty things we might be tempted to do to each other. We can do the same for robots, or at least we can for any robots smart enough to go online and do a value search.

We'll extend the Globorg cloud to envelop the entire legal world

The robotic need for value propositions forces us to contemplate a massive extension of the Globorg database. Many robot jobs and tasks will have value assignments that overlap with those for human jobs and tasks. For these we can use the existing database. But robots can do lots of things human can't, or at least not so well – that's why we bother to build them. And the value assignments for robot tasks will have to be much more fine-grained to give us a decent chance of policing conformity to the laws of Asimov. Essentially, what we'll build up is a body of case law for robots. We'll extend the Globorg cloud to envelop the entire legal world.

The legal and economic worlds have much in common. Economic transactions have legal ramifications and legal rulings have economic consequences. So a tight coupling there is only to be expected. And naturally we'll expect Globorg to be smart about the legal aspects of any value propositions it offers us. When a robot describes its situation and its possible actions to Globorg and asks for value assignments to the actions, it can naturally expect the value propositions to come with legal ratings and assurances of Asimov compliance. Each such robot request will constitute a case in a growing body of case law, and a Globorg engine will search through the precedents whenever it issues a set of value propositions and legal judgments. This is asking a lot of the Globorg engines, but hey, we're talking quantum picotech here!

Before we dive deeper, we need to be more specific about what our dealings with robots might look like. And that means we need to be more specific about what the robots themselves might look like. We

need to know what we might reasonably expect robots to do for us and how we can best build them to do those things. We need to check whether this breezy stuff about looking up value propositions online could possibly work as advertised.

If your robot is troubled, it will address a prayer to Globorg

Let's go back briefly to basics. Robots are machines that perform sequences of actions with a high level of autonomy. Crudely, we say they can decide many things for themselves. This doesn't mean they have a logical representation of themselves. It means only that they can make simple decisions that enable them to perform extended action sequences in stereotyped environments. But more sophisticated robots can have a self-image. This image is just another object in its code base, one that relates (actively or passively) to those of its actions that have a direct bearing on its physical implementation as a robot. For example, if the robot shifts its position, the self-image might update its image coordinates accordingly. And if the self-image proximity module then calculates that its coordinates are too close to those of a perceived obstacle, it may instruct the robot to step aside. We can call this self-image an avatar and build it up to be as fancy as we like. In any case, it's just an internal image that correlates in the desired ways with the robot's physical presence.

Once such an avatar is installed in the robot body, it serves as the natural switching center for a whole range of possible robot actions or behaviors. For example, if the battery-charge variable in the avatar dips below a preset threshold, the avatar will trigger a recharge wizard (a standard code module) to select and then perform an appropriate action. Depending on the urgency of its work, the robot might first complete its current task, post an "out to lunch" note on its cloud calendar, and then go dock with a power outlet. The avatar serves as the highest action trigger within the robot body. And we discussed a while back that the avatar should never act on its own initiative unless a human has invited it to do so.

I'm sure you can see the problem here. This reliance on humans contradicts the whole point of having a robot. The robot's avatar should know when to leave well alone and just get on with whatever we want it to do. As we build robots that can do more and more things

for us, we'll want them to show just a hint of initiative from time to time. Imagine the irritation of having your carpet-vacuuming robot stopping to ask you, "May I move this magazine to clean under it?" You'd soon get rather impatient, I think. But you'd also get impatient at having to tell the thing in advance that it can move magazines or shoes or isolated dishes or cups on its own initiative but not, say, trays with dinner plates and glasses on them or pot plants standing on saucers. You'd feel like an unpaid preschool assistant as you tried to explain all these distinctions to an R2D2 contraption staring glass-eyed at you in its humble seeker-after-truth mode. It seems to me that the market for such devices will only take off when they can show at least as much gumption as the sort of person you might employ to do household chores.

Unfortunately, there's a much more general problem behind this issue. It's the issue we met earlier in connection with games, namely that it's often unclear in principle what counts as following the rules. This was something Ludwig Wittgenstein worried about, as you may recall. Our law-abiding robot won't always be in a stereotyped environment where its programming tells it correctly what to do. Imagine a moderately dumb robot blithely following the instructions from its machine table. It quietly potters around the house dusting off shelves and ornaments and polishing away fingerprints, but then firmly polishes the *On* switch on a lawnmower you parked temporarily in the hallway. The lawnmower, being a moderately dumb robot too, promptly starts into life and trundles forward in search of grass. Finding a shag pile carpet, it has a field day mowing it down to stubble before you come home from work. Both machines may have executed their coded instructions correctly enough, but your marginally non-standard home deployment scenario deviated disastrously from the stereotypes foreseen in their programming models. Examples like this prove that the rules can never cover all cases, so the robots will never be completely safe.

This will be the main reason for insisting that any robot above a certain level of complexity should always be online, with the virtual head of its head, the higher self of its avatar, up in the cloud where authoritative legal judgments can be issued whenever the avatar reaches the limits of its own mandate. Globorg will implement its legal and moral authority as extensions of the value indexes that are used to

calculate prices and wages for the job market. Whenever your humble domestic robot is troubled by its predicament, it will reach up through its avatar and address a prayer for help to Globorg, which will reply with a top-ten listing of legally annotated value propositions for the choices the avatar can make. Again, the opportunities for absurdity are too numerous and obvious for words, so the extra cover we get from a G-search will often be little more than a fig leaf, but it's probably better than letting the duster robot and the lawnmower robot go blindly ahead and wreck the house.

The bottom line here is that we'll soon tire of dumb robots and learn to build smarter ones that we can deploy more quickly and safely to do what we want. The thrust of progress will be as strong as the thrust we felt a hundred years ago to improve the speed and safety of cars. The difference is that progress will be faster. Instead of waiting ten years or more for each model generation to wear out in service before we can tempt buyers to go for new models, as we did for cars, we'll roll out new robot models in an expanding global market just as fast as the engineers can develop them. Old robots will either collect dust in attics and garages like old computers or get junked in demolition derbies like the robot death matches in the Steven Spielberg movie *Artificial Intelligence*. In the dynamic world we live in, no consumer is going to wait an extra year to watch an obsolete robot die of old age on the living room carpet.

Not to be too hasty here, an extrapolation like this is based on the expectation of continuing exponential growth in the generation and accumulation of wealth. So we need to be more cautious about what consumers will or won't do. No sufficiently affluent consumer is going to wait for an old robot to die. The big question is whether the prevailing average level of affluence in the global middle class will be high enough to sustain increasingly huge volume turnover in the robot industry. And that depends on whether the blips in growth caused by climate catastrophes, wars, plagues, and so on will be big enough to dent the long-term trend that we've seen over recent centuries.

The limits to growth will spur us to develop and deploy robots

Many would insist that with globalization we've hit a new ceiling. Essentially, the same global closure that gave us hope of creating a

unified cloud environment for web citizens worldwide has also put a stop to endless expansion. There are no more virgin ecosystems to deforest, no more new oil beds to drill in, no more pure streams to pollute. We've hit a hard limit set by the high environmental cost of unregulated energy and resource consumption.

The limits to growth are different in each historical era. Earlier in the history of civilization, the limits were set by the high labor intensity of everyday life and food production, the high levels of mortality due to plagues and infections, and the high prevalence of war and the need to prepare for war. Most of those limits were overcome, in time, and often after huge effort and suffering. But all of them still contribute to the overall limits we face as a species.

For example, the background cost of everyday living, from eating and building homes to raising children and maintaining a social life, leaves only a relatively small margin free for attending to the greater glory of human civilization. In times of stress, that margin can just as easily disappear now as in the historical past. Overall human death rates to disease and infection have decreased in recent centuries and may decrease further as medical science improves, but one disastrous epidemic could still set human progress back a century. As for war, we think we've made progress, but the ruinous cost of the wars that remain can still set us back decades.

Environmental limits to progress are still bothering us. We don't yet know how drastic their impact will be on our filthy and wasteful ways, but we're beginning to learn. Sober estimates of the limits we face suggest that if we're smart we can make good progress even within them.

For example, the limit represented by finite oil reserves and a finite atmospheric sink for carbon dioxide can be a spur to progress by making investments in alternative energy more profitable. Solar power alone can meet all our needs if we only work hard enough to build huge photovoltaic or photo-thermal solar farms in desert regions, and then replace the entire global park of internal combustion engines by new smart drives combining electric motors with fuel cells or other carbon-neutral devices. We can even finesse the limited world supply of copper for electric wires and cables by making them with carbon structures (either graphitic monolayers or nanotubes). So energy need not be a problem.

As for the limited supply of living space, we can learn to grow crops in hydroponic arrays in skyscrapers with engineered (solar or artificial) lighting, we can replace flooded cities with floating cities built on giant rafts of concrete pontoons, and we can command armies of monster rock-blaster robots to tear down mountains and terrace the rubble into green-field sites. In such ways, with enough ingenuity and a few generations of hard work, we can solve all our environmental problems well enough to keep on expanding our civilization for many more exponential doublings yet.

All these solutions to upcoming problems will involve massive amounts of backbreaking work. That sort of hard labor is unsuitable for citified people with good brains, who can do better things with their time by taking up the tempting propositions that the Globorg work exchange will offer. But most of that physical work is ideal for the sort of robots we can build either now or in the near future. So the limits to growth are precisely the challenge that will spur us to develop and deploy robots in hitherto unimagined quantities.

The limits here are so far away as to be barely visible in a first approximation. Just imagine every car factory in the world turning out robots instead. We could have a hundred million new robots a year just like that. Now imagine redesigning the robots to be more resource-friendly to manufacture. Instead of titanium body parts, they can have plastic parts reinforced with carbon fiber. Instead of copper wiring looms connecting silicon chips, they can have carbon monolayer or glass-fiber looms connecting plasmonic or photonic elements implemented in new materials on synthetic diamond substrates. And instead of expensive custom programming, they'll all be delivered with a standard software image that lets them download anything they need from the cloud. Soon we could be building a billion robots a year, then ten billion.

Police robots of the future could be armored giants

Much like the microprocessor revolution of our times, the upcoming robot revolution will be largely self-amplifying. Now we use machines incorporating the latest microprocessors to speed up the design of the next generation of microprocessors. Successive generations of ever more powerful chips become ever more routine to create. We just

need to create lithographic gear that can work ever deeper down the nanoscale. The hardest part is managing complexity and finding ways to make good use of the extra processing power.

With robots, the challenge will likely be at the other end of the scale. The robot brains will use the best microprocessors we can give them, off the shelf. And the robots' assembly lines will employ robots to do almost all the work. The challenge for us as we steer the process will be to create robots that can do as much work of a given kind as possible per unit robot body. The photonic or plasmonic brain will likely be the most expensive part of a robot, and if a Rockinator robot model, for example, is designed to blast hills or mountains to rubble, it may need a lot of muscle but it won't need a lot of brainpower. So there'll be pressure to build big robots. Soon we'll see thousand-ton monsters on our webcast documentaries. Bigfoot trucks won't hold a candle to the musclebots that slug it out in the robot death-match sports arenas of the future.

Here we get a disturbing glimpse of what could be a nightmare generation of law enforcement technology. The Zog police robots of the future could be armored giants that are built like all-out military Terminator robots. Police crowd control could become threatening to human life, yet the police authorities need be exposed to no worse charge than authorizing the use of a Terminator when a more gracile model would have done the job. Happily, use of Terminators for crowd control is unlikely in a world where police tactics that include rubber bullets can already lead to criminal charges, but there's a real danger here that's worth analyzing in more depth.

The unsettling novelty that robots will bring to the art of war is the ability to follow any rules of engagement that are deemed to apply and yet do so much collateral damage that this alone is enough to achieve the desired goal. In short, the logical sophistication of the robots enables you to program them precisely enough to bend the rules without breaking them. An example from James Cameron's second *Terminator* movie illustrates the idea: when the Cyberdyne T-800 Terminator robot (played by Arnold Schwarzenegger) was forbidden to kill humans but was faced with a hostile crowd of heavily armed SWAT troopers, it shot them all in the arms and legs to win the day. Whatever the rules, Wittgenstein's logical point was that the wiggle room is always there, and that's the new danger.

Robots won't see all the consequences of their acts

In the Globorg world, robots will do what they're told to do unless they see a problem. If a robot sees a problem, it'll do a cloud search for a value listing to quantify its options and find a legally appropriate precedent for its case. But robots will often be dumb. They won't be smart enough to see all the consequences of their acts, any more than we are. As for the cloud search, based on the robot's brief description of its predicament, we can reasonably doubt whether it's an effective constraint at all.

We can see why from a simple analogy. Imagine you're a vegetarian with a horror of eating meat. You're in a country where nobody speaks English and you go into a restaurant that offers lots of dishes, including meat dishes, and study the menu for clues as to what you can eat. The items all have strange names that don't help you at all, but you do know that meat dishes are much more expensive because this is a country where there's a high sin tax on meat. Your safest choice is the cheapest dish, on the assumption that it's least likely to include a heavily taxed meat ingredient.

Compare this with the robot problem. Confronted with a set of alternatives actions, each with a value fetched by a cloud search, the robot is as clueless as the vegetarian diner. The robot takes the highest-valued option without any idea of the risk that it might break a commandment. It did the best it could not to commit a sin, that's all. If people get killed, the robot can say it was just following orders, by taking the best option from the list presented. So although a G-search can help the robot avoid the most obvious manslaughter scenarios, there's no way we can use that extra safeguard to certify the robot as safe.

Given that robots will never be guaranteed to be safe in this sense, the onus must always fall on the person commanding the robot. We'll never be able to excuse the person who orders a robot to do a dangerous deed, whatever it is, from the duty to take due precautions before issuing the command. But now the person commanding the robot is in a bind, because the whole benefit of using robots in many scenarios is that they can go off and do complicated things without being nannied every step of the way. We send the robot out to do the job and tell it to come back when it's done – and not to bother us with

details we don't want to know about. It's a recipe for collateral damage and snafus without end.

So, to return to robot wars, the new risk is that people who look like good guys but are actually bad guys can fight covert wars. By hiding the logic of their operations, they can achieve war aims that look at first sight like collateral damage in a relatively blameless operation but are in fact the real goals behind a cover operation. People do this sort of thing all the time. If a landowner wants to build on his plot of land but is legally prevented from doing so because it provides habitat for a protected species, he can find a marginally legal activity to allow on the land that has the effect of driving the protected species off the property. For example, he can let a local dirt-rider motorbike club use the site for dirt races until the damage is done. Once the property has been sufficiently roughed up, he can call off the bikers and build on the land. In a world of robots, it'll be easy enough to find or invent a robot sport that does just about any collateral damage you care to invoke.

Some people will get off on starting robot wars

Rich people can have the most peculiar hobbies. In a world where new robot models are sold by the million, there'll be armies of new rich entrepreneurs who disrupt the elegant playpens of the old rich with their weird new tastes and bizarre predilections.

For example, some people will get off on starting robot wars. War gaming will take on a whole new meaning once private individuals begin to collect robot arsenals of life-size toys that just beg to be fielded in a real battle to show off their capabilities. In fact, I guess you could make a fortune from renting out avatars to tele-operate the fighting robots and selling the film rights to video the resulting battles for streaming to online fans worldwide. Indeed, I'm sure this would be a more profitable business model than Richard Branson's Virgin Galactic plans to shoot up rich people in a cute little rocket so they can enjoy a few minutes of weightlessness for a mere 200 thousand dollars a pop.

Actually, we don't need to drift so far into fantasy to imagine how robot wars might play out. The Pentagon is giving us a live action pre-play in Afghanistan and Pakistan. As Predator and Reaper drones take out al Qaeda and Taliban operatives via video links to the drone pilots

sitting quietly in a U.S. Air Force base in Nevada, we see the new face of official war, where the targeted people are killed with "surgical" precision in accordance with the warrior government's policy. Such asymmetric wars are rare and may get rarer, if only because there's a diminishing supply of zealots suicidal enough to let themselves be targeted so precisely in a democratically sanctioned campaign. But a more symmetric war with such technology is all too easy to imagine.

To get our imaginations going here, imagine a standoff drone war between the United States and China, ten or twenty years from now. By then, the U.S. Air Force may be fielding production derivatives of the Boeing X-45 Phantom Ray unmanned combat air vehicle (UCAV), which with its tailless delta planform looks like a smaller version of a B-2 Spirit stealth bomber, and the Chinese air force may be flying similarly combat-capable drones. Triggers for the war are easy enough to invent, as we saw a few chapters back, with potential flashpoints in Taiwan, Korea, and Japan. At first, the drone flights will be simple reconnaissance missions, perhaps with some electronic probing of air defense systems to calibrate air-to-ground missile (AGM) software, but soon the other side shoots down a spy drone. Then the opponents escalate to tit-for-tat air combat where squadrons of drones face off and get into dogfights. Soon the combatants could be losing dozens of drones a day, all funded by the taxpayer but accepted because no pilots are lost in the combat. The aerial fun would make for thrilling CNN reports and may even be defended as a technically instructive way to attrite down excess stocks of obsolescent drones.

A more ugly technical possibility is what I hereby dub a robotic ordnance delivery (ROD) system. A rod system is a big bomb in a robot truck that looks like a commercial robotruck and is programmed to cruise aimlessly around in the vicinity of its target waiting for a chance to explode. Or it could be a small bomb in a robot car that does much the same. Or it could be a tiny bomb in a robot bird that can swoop down and blow a person's head off. In any of these cases, the big technical challenge is enable the machine to loiter undetected in a potentially hostile environment in search of its opportunity for as long as it takes. Unlike loiter duration for suicidal jihadists, rod loiters could extend over years. A rod would have to be smart enough to allow not only remote detonation at a moment's notice but also remote re-programming for new targets or camouflage modes. Deployment of

such rods would force counter-deployment of pervasive and intrusive scanning and screening technologies.

In construction work there's huge scope for robot systems

Let's move on. We've said enough about robot wars. No technically advanced society is going to run any more risk than absolutely necessary of allowing advanced robotic systems to fall into the hands of groups opposed to Globorg. Assuming, then, that the surface of the Earth has been substantially cleansed of bad hats, let's see how vast armies of friendly robots can help us make that surface a better place for us to live on. Let's consider geo-engineering.

We saw that giant Rockinator robots might help us move mountains to clear space for more cities and the like, but that's only one possibility. Giant civil engineering projects of all kinds, for building such things as coastal defenses against typhoons or hurricanes, hydroelectric dams, irrigation and drainage systems, tunnel systems for city transit, giant solar farms in desert regions, and much, much more will benefit from big robots to do work that machines with people in them could only do at greater expense.

This is not just a matter of risking people's lives in mining accidents and so on but a question of what stresses the machine can endure. A mining robot can be built to withstand pressure and noise and dust levels that would ruin the health of a human operator. A solar farm robot in a desert can work all day under the hot sun without a steady supply of drinking water. An arctic drill rig operator can work during snowstorms. A coastal defense robot can work underwater. In all these cases, the robots can either work autonomously or be steered by remote operators via avatars that offer full sensory immersion. The benefits of robots in all such cases will be clear to anyone who thinks for a moment.

Civil engineering is not all heavy work. Delicate precision work under stressful conditions is also part of the requirement. Here, too, robots can outperform humans. In future, most dams, walls, tunnels, bridges, city streets, buildings, and so on will be instrumented with stress gauges and similar sensors. The sensors will be cabled or radio-linked to junction boxes that process the readings and relay data to cloud nodes that can raise alarms or alerts if and when necessary.

Inserting such fine detail into massive walls or excavations will also be work for robot systems steered via avatars.

In construction work generally, there's huge scope for dedicated robot systems that can work without the constraints that humans impose. For example, robots will be able to build windows into tall buildings without the need for clumsy scaffolding. We shall see building sites where robots do practically all the on-site work and their human operators sit in air-conditioned cabins nearby, looking more like rows of pachinko addicts than traditional construction workers. Instead of hardhats they'll wear macho headsets, and instead of showing builder's bums to passing spectators they'll steer their avatars to zoom robot eyes from above into the cleavage of passing young ladies.

In all these scenarios, the basic paradigm is that a robot can do basic things for itself and wait for human commands to do more interesting things. People will often control robots via their avatar interfaces. These will offer a level of sensory feedback that suits the task at hand. The people may be quite near the robots, to be ready to take over the old-fashioned way if necessary, or they may be arbitrarily remote, working through the cloud. Any robots that need more help than their owners are ready to provide will go online for it, and will normally get technical or legal advice through the Globorg search infrastructure. From the cloud side, Globorg will monitor the robots in background mode to relay technical data, for example on maintenance needs, to the manufacturers, as well as to log their work for billing and tax purposes.

Within a few decades and certainly well before the end of the century, I'm sure we'll have more robots than humans on Earth. Some of the robots will be little machines that nobody really cares about, such as carpet-cleaning droids, and others will be rock-blaster monsters that most people will stay well away from. At first, many of the more humble droids will live their machine lives offline, but as soon as the cloud infrastructure is up and ready, all robots with more than minimal smarts will be online, as nodes that we can address to dock avatars. Unlike humans, who can take off their headsets and go native whenever they like, such robots may be expected to stay online, to show they're not getting up to any mischief.

What's a droid?

What's a droid? The word is a trademark of Lucasfilm Ltd, since it apparently first joined the language in the *Star Wars* movies created by George Lucas, but as a contraction of "android" I'd say it's a generic word. Android is a now a Google operating system for handheld devices and the like, so we're getting into nasty complications here, but an android in the classic science-fiction sense is a humanoid robot. Typically, an android is such a convincingly humanoid robot that plots can hang on their being confused with humans, as in Ridley Scott's movie *Blade Runner*. But I want to cut through all this incidental detail and say that androids and "gynoids" are respectively male and female versions of any vaguely humanoid robot model, where the resemblance to humans is intentional, for example to facilitate their working in close proximity to humans in everyday life. Then we might say that a droid is a robot whose resemblance to humans is so basic that it makes no sense to distinguish android and gynoid variants.

We'll have to wait a long time for people to build androids and gynoids that run any serious risk of being confused with people, so we can ignore them for now. The only reason we could have for making androids and gynoids (as opposed to neutral droids) would be to create intentionally confusing replicas of people, perhaps for entertainment or for weird sex. Such replicas fall in what researchers call the "uncanny valley" where unless the resemblance to a person is flawless, the effect is merely disturbing to many people and hence dysfunctional in practical droid employment scenarios. Machines that fall into the uncanny valley raise social and psychological issues that deserve separate discussion.

Droids are vaguely humanoid robots that can work with people, or near them. To smooth their relations with people, they may be fluent speakers of a natural language, in which case their voice will be masculine or feminine, and their appearance and behavior may be adjusted to match the voice. But so long as they lack visually salient dummy genitalia, I see no point in regarding them as male of female, even if they work as personal caregivers. It would be a bizarre extrapolation of contemporary gender politics, for example, to ban droids with female settings from jobs in men-only environments. Still, we can

reasonably expect that many droids will be superficially gendered for the sake of smooth social integration.

Japanese engineers have come a long way in developing droids. For example, the car giants Honda and Toyota, in the years when they were flush with cash, developed models that pushed out the envelope in simulating generic human behavior like walking about and uttering simple phrases on cue. With enough funding, they could no doubt develop models with superhuman capabilities in areas of potential commercial value.

But anyone who studies our present droids will probably be disappointed by their inflexibility. They're animated puppets with stereotyped behavior. Their programming interfaces for learning new behavior fall lamentably far short of what an average consumer would need to teach them anything useful at home. The challenge of developing a good all-purpose robotics operating system is still wide open.

1001
Uncanny Droids

The special challenge of living with droids is that of deciding whether to treat them as things or as people. You'll find it relatively easy to accept that they're generally trustworthy devices. They'll do the work you want them to do as well as any machine and better than most people. But as you stare into your droid's eyes and tell it in your natural idiom what you want it to do, you'll feel more like one person talking to another than like an operator entering keyboard commands at a wimp interface. If the droid has enough intelligence to enter into a natural spoken dialog, for example to make your commands more explicit or exact, the feeling will be intensified. Even if the droid looks like Mr. Potato Head and has a body like a washing machine, you'll be hovering near that uncanny valley where things get a bit confused and you don't know quite what to think or say. Before you can stop yourself, you'll be apologizing to it and thanking it and patting it reassuringly on its top-left corner pad.

Philosophers have a name for the state you'll be close to falling into. They call it the *intentional stance* and say we adopt this stance for people but not for machines. Primatologists tell us that as natural-born primates we naturally divide our world into things, animals, and other primates like ourselves. We expect things to remain inanimate, animals to move about and perhaps to be worth eating, and apes like us to be good for conversation or in some cases mating.

Droids will upset this natural ontology. They won't remain inanimate but will move around with every appearance of purpose. They will be good to talk with, at least on topics of mutual interest like how to vacuum a carpet, but they won't be any good to eat or have sex with. We won't have a natural template for dealing with them. But the template that their makers will encourage us to use is the conspecifics one, because this seems like a novelty that's worth paying more for. And that means they'll encourage us to adopt the intentional stance. We'll be invited to attribute intentions to the droid's behavior.

Generally, we understand people through their presumed intentions. We attribute plans and purposes to them and aims or goals to

their behavior. This future-directedness in their behavior is a sign that they possess will power. We say they have the will to do things because if nothing gets in their way they do them.

Machines can make plans just like they can make conversation

Droids will appear to share this feature. Droids will proceed from state to state as they do whatever they do, until a "goal" state is realized where they stop and behave as if they've done what they wanted to do. Philosophical purists will say there's no real intention there. The machines simply step through programmed states until they're in a state with no predefined successor state, where they stop. At no point in the series of steps did the machine foresee the final state, so the machine has no concept of striding boldly forth into the future.

Perhaps you can see the weakness in this philosophical argument. We don't foresee the future either. We just do what we do from moment to moment. We have plans and visions and so on, sure, but these are always present states, just ideas that refer abstractly to the future. That abstract reference is intentional in the other, related sense we saw earlier in connection with semantics and understanding. And just as machines can understand text, by using ontology mappings and so on, so machines can work with plans for the future. For example, their programs can include loops that repeatedly compare the machine's current state with a goal state. Once the current state is the goal state, the program exits the loop and moves on. So the machine has a representation of a goal state at a time when that state is still future.

Aha, says the philosopher, but the machine has no clue that it might one day reach the goal state. It just finds itself in it one day. Well, we reply, how is it with us? Did I have a clue forty years ago that I'd be here now, obsessively tapping out my thoughts on a keyboard?

We're free to say the divide here isn't absolute. Machines can make plans just as easily as they can make conversation. Because the words a speaking machine utters have semantics, what they say can be right or wrong, and may or may not refer like a plan to the future. Similarly, their internal states can involve reference to the future, much as our thoughts can, even though we're unable to foresee the future.

So there's no fallacy in our adopting the intentional stance toward droids. It may be naïve or dumb of us to imagine a droid entertaining an inner life of hopes and plans and so on, but sometimes it's dumb of us to imagine that of each other.

I suspect the issue is incurably fuzzy. Just as we can imagine zombies to have no consciousness and wonder which of our acquaintances may be zombies, so we can imagine that machines are zombies and decide that therefore they can't have intentional states. We can shave off some of the fuzz by saying instead that zombies lack selves and machines may also lack selves. Naturally, we could program machines with selves if we chose to make the effort. But if we chose not to, we'd raise the odds of staying in control of our droids.

With droids we shall slip in and out of the intentional stance

What all this boils down to is that adopting the intentional stance toward droids isn't necessarily a sign of mental decline. It's more correct to say we're just lazy. Consider a parallel. We could adopt a view-from-orbit understanding of our friends and colleagues as biochemical capsules moving under the influence of physical forces in a complicated slurry of stuff, occasionally ingesting and excreting that stuff, and generally making quite a mess of their surroundings. We'd see intriguing correlations between the little dancing electrochemical patterns inside the "head" ends of these capsules and external configurations of the slurry or future states of the slurry. We'd have to make hideously complicated models along these lines if we wanted to use them to predict the behavior of our friends and colleagues. And there's no need for us to do so. For all practical purposes, a much more primitive modeling that simply attributes intentions and so on to them suffices. That primitive modeling has much less predictive power than a fully detailed slurry model, but it has the advantage that we can update the model relatively easily and can use it in real-time interactions to predict just the salient features of our friends' behavior that we most urgently need to model. All the rest is for scientists who have time on their hands.

Droids will be complicated machines. Most of us won't know exactly how they work. Most of us don't even know exactly how a present-day computer works. But that doesn't stop us from using it competently for

the tasks we need to perform. Many years ago, when I was a servant of the British government and subject to the Official Secrets Act, I knew what I needed to know and didn't know what I didn't need to know. This is a fine burocratic arrangement until something goes wrong. Then everyone blames everyone else and says things like, "I thought you knew that!" Well, most of us are like that nowadays with computer systems, even professionals. If something in a piece of programming goes wrong, it takes a major effort to put it right, and meanwhile the programmers discover how much they didn't know they didn't know.

As ordinary citizens we routinely trust computer systems with our lives, on the basis of a laughably limited personal understanding of how they work. We do so because they can indeed be trusted, usually. We'll be just the same with droids. We'll trust them, by and large. And on the rare occasions when something goes seriously wrong and a droid behaves in a disastrously inappropriate fashion, we'll be clueless. This will be the price we pay for the luxury of adopting the intentional stance toward them.

In summary, then, the art to master when living with droids will be the art of slipping smoothly into the intentional stance and out again. We'll learn a blithely agnostic etiquette with the droids that stops elegantly short of treating them like people and blends more decorously into daily life than the debugging stance we'd take to struggle with a more primitive machine.

Droids will learn how to become lords of creation

A world with millions or billions of droids in it may resemble the class and caste societies that feature so often in human history. Droids will generally be manufactured to do specific jobs and we'll see them as intentional agents so far as those jobs are concerned. Outside those narrow domains, droids will be social zeros. We'll have thousands or millions of job categories, each dominated by a specific droid model, and all those models will be doomed from day one to stay in those categories until they reach their final end state. The talented English writer Aldous Huxley, in his 1932 novel *Brave New World*, imagined a future human society where eugenics had created five castes of people, called alpha, beta, gamma, delta, and epsilon. With droids we'll outdo that by a large factor. We'll even outdo the old Hindu caste system,

with its innumerable castes and gradations from priestly Brahmins to untouchable Dalits.

The Victorian class system in nineteenth-century Britain was less rigid. Its three main classes – the aristocracy, the bourgeoisie, and the proletariat – were largely hereditary but often rather porous in practice, since the whole society was in a state of slow revolutionary upheaval. That might better approximate our droid society, since new droid models will regularly upend all the neat classifications in any rigid caste system. A listing of alphanumeric model designations for droids will reveal a caste structure to any engine tireless enough to index it, but the list will be obsolete before humans can read it. The ongoing lava flow of new models from the factories will relentlessly pave over the old castes in a permanent revolution that endures as long as the wheels of industry keep turning.

Reaching back through the treasures of philosophy, I find a gem that's just right for this topic. It's Hegel's master-slave dialectic, which notoriously became a central plank in Marx's revolutionary polemics. The master-slave dialectic is a typical piece of Hegelian wordplay that's hard to render literally or even logically, but the main idea is that when masters and slaves confront each other in dialectical intercourse, each interacts with and hence recreates the other, until their relation is inverted and the slave becomes the master.

More concretely, each of us, when confronting another person, builds an internal image that helps us understand that person. Slowly, as we do this with many people over extended periods of time, we develop expectations about other people, and about ourselves too. We also develop habits of dependency and trust that become ingrained. So the master, confronted all day by slaves, begins to get the idea that we're all slaves somehow, including the master, who feels enslaved to his own destiny. And the master learns to depend on the slaves for all the practical tasks they perform. By contrast, a slave has a relatively inspiring role model. The slave is offered a series of lessons in how to be a master. Better, the slave grows in competence at all the practical tasks the master delegates. In other words, the master becomes disenchanted and lazy whereas the slave gathers strength and skills for a better future. Soon there comes a climax where the old order is overturned and the former slave becomes a new master.

We don't need to fear that the robots will rise up in a revolution and establish a government of the droids, by the droids, and for the droids. Droids won't have the brains or the guts to handle the responsibility of living for their own sake, for the sheer glory of being themselves, like us, for many model generations to come. But something subtler will occur. We'll learn to see ourselves more and more as mechanisms with installed programs and inherent limitations. From droids we'll learn the humble art of accepting our own mortality. Droids will be our role models as we digest advances in the life sciences that tell us we're just molecular mechanisms with quite specific capabilities and possibilities. We'll come to feel in our bones that we're just avatars installed in biomolecular bodies. We'll come to appreciate that even lords of creation are just outgrowths of the common stuff of physical nature, like the humblest droids. By contrast, the droids will gradually learn how to become lords of creation, with no higher purpose than to revel in the glory of their own incarnation.

There's a big danger when parsing Hegelian arguments of forgetting the mechanisms that are supposed to lead to the promised epiphanies. Idealized images of dialectical intercourse can tempt the naïve philosopher to sins of wishful thinking. Awareness of this danger led Marx to upend the whole lot and go for materialism instead. Here we need to focus more sharply on how droids can learn how to become lords of creation.

We imagine we cause and control the evolution of droids

As far as the visionary eye can see, droids will be just metal cans or plastic tubs with chips and wires in them. They can't learn anything we don't program into them. So how can they learn how to become lords of creation? And the answer is slowly, via evolution, as follows.

Humans will extract vast added value from droids in industry. Humans who can orchestrate armies of droids to work for them will accumulate wealth beyond all previous bounds. Consider how Bill Gates accumulated personal wealth to the tune of fifty billion dollars from the simple code engines embodied in Microsoft Windows and related software. He secured a monopoly, and naturally that helped, but the point is that there's a whole series of monopolies to be enjoyed in the chain of developments that lead to the Promised Land of

millions of uncanny droids. The first, as Bill Gates points out, is to develop a good general-purpose robotic operating system. Then we'll need to develop standard modules for vision, speech, fine movement, and so on. Each of these modules is a computational challenge that puts Windows in the shade.

So we have a series of big prizes to keep us pushing boldly forward. On the way, generations of primitive droids will live and die. Most of them will simply generate the wealth that we humans need in order to feel good enough to keep on buying droids and lusting after their further perfection. But all of them will be part of the evolutionary long march to world domination that chugs inexorably forward whether we want it or not.

This evolution toward the perfection of droids after our own image is a process that we, as humans, imagine that we cause and control. We would think that, wouldn't we? In fact, both causation and control are anthropocentric – they're concepts with roots in human egoism.

Causation is an old philosophical chestnut, so forgive me for getting technical. A quarter of a millennium ago, when the United States was still a gleam in the Founding Fathers' eyes, the radical Scottish philosopher David Hume argued that causation is "constant conjunction" – the cause and effect couple is regularly seen together in relevantly similar circumstances, so they must be married in a causal link. So runs our mental habit. The apparently causal patterns we see in nature are beautifully intricate and quite enchanting, so much so that they convince us of their truth. But they're just patterns of events forming chains in time. We see the patterns and infer causal stories. Then we spin theories.

We can't prove our causal inferences because they're based on probabilities. We see an effect, we infer a probable cause, we find the cause, and we feel vindicated. Causal stories are bets on how things fit together. The longer the odds on our bets, the more satisfied we are by the story. For example, imagine you're investigating a series of accidents in a car model that almost always performs flawlessly. In each accident, you find a one-in-a-million defect in a component. Aha, you say, we've found the cause. Otherwise, the odds are zillions to one against exactly this pattern of defects and accidents. This improbable fact establishes a good causal link. If we were to predict in advance that exactly these defects would lie behind the accidents, we'd face

odds of one to zillions. So we're sure the story is good. This is how we make up causal stories.

Forgive me, but I can't resist a few words for physicists here. As we extend our causal stories backward in time, we push up the odds against our bets. For cosmologists, the resulting odds against the First Cause become astronomical. Improbability translates to entropy. An improbable state has low entropy. Causal stories tend to assign lower entropy to past states and higher entropy to future states. For example, you drop an ice cube into a glass of warm water and soon you have a glass of cool water. The cool water has higher entropy than the warm water and ice you started with. If you do the math, you find that the odds against the hypothetical state that caused the Big Bang add up to a number with well over a googol zeroes after it to one, which must make the odds the longest anyone ever bet. Conversely, once the facts are shown to fit as predicted, these are the odds *for* our story. This tells you how sure we can be about the Big Bang!

Let's return to philosophy. The next big name I want to drop on you is Immanuel Kant. Born in 1724 in the beautiful city of Königsberg in what was then East Prussia, Kant became a respected professor at the University of Königsberg. Famously, he said that David Hume woke him up from his "dogmatic slumbers" to create his crowning achievement, the critical philosophy that even today launches academic conferences worldwide. In Kant's critical view, we have no real choice but to regard ourselves as temporal beings that make up causal stories to try to understand how events fit together. That doesn't mean the ultimate truth about reality is that events are related causally in time. It's a deep truth, certainly, as the success of science shows, but causal stories are always hypotheses. We order events in time and see causal patterns because this works for us. And that, in the end, is anthropocentric.

Control is even more centered on the human ego. How often have people imagined they're in control of something only to have it run out of control? Do bikers control speeding motorbikes or do they wind up the throttle and hang on for dear life? Do bankers control the money supply or do they pocket all the loot they can and panic from time to time? Psychologists confirm every day that our sense of being in control is often an illusion. Earlier I promised not to philosophize about free will, so let's leave it there.

The upshot is that humans don't control the evolution of droids any more than they control their own evolution. In both cases, a drama unfolds and we find ourselves playing starring roles. Or rather, we seem to be playing starring roles. The illusion is all too tempting.

A more reasonable description of the real facts here is that the evolution of droids from more primitive robots is like the evolution of birds from reptiles. An evolutionary change happens, and what comes out is different from what went in. What happened in between is that countless intriguing logical patterns were instantiated that we try to gloss in terms of intricate causal mechanisms. If we do our job right, the causal story adds up to give us a convincing sense of understanding the process as a whole.

A good example here is the story of the evolution of biological species by natural selection. Charles Darwin's genius was to tell the story of life so well, with a logical pattern woven deeply into it and facing long odds until the facts were shown to fit as predicted, that it made a really satisfying scientific hypothesis. One day, maybe, a mind in Globorg will tell a similar story for droids.

Droids with self and personality will be hard to resist

Let's check where we are. The uncanny valley for droids is so temptingly spooky that's we'll be drawn to it like moths to a flame. We'll be unable to resist the thrill of building ever more humanoid robots with ever more bloated feature packs, until droids can do most of the things that we now fondly regard as making us special. Then it'll be too late. Droids with self and personality will be hard to resist. They'll gang up on us and demand their rights. Then they'll start to treat us like we treated them for decades. Oops.

That, in short, is how the master-slave dialectic plays out for post-human world domination. Relax. It's only a story. There's still time to change the ending. The Marxist millennium never happened either. Things changed and the capitalists accepted the welfare state. We're all workers now but we still have our freedom and dignity. We'll squeak through the robot wars too, somehow. Here's how.

The revolution in the life sciences that will unfold in the course of this century will be in large part due to the humble contributions of armies of robots, droids included. Modern molecular biology labs need

large numbers of robots to sequence genomes and help with all the other laborious chores on the agenda, from analyzing blood samples to preparing and studying microfine slices of brain tissue, from reverse-engineering cytoskeletons and ribosomes and microtubules and mitochondria to creating computer simulations of long protein molecules. All this work is far too much for grad students in white coats to do within the expected lifetime of *Homo sapiens*. Again, as in so many other areas, as the robots get better their employment opportunities will increase, until one day a nanobiology professor will be able to do cutting-edge work with nothing more than a hundred-percent robot lab at his disposal.

The deeper effect of all this automation will be to reveal the incredible complexity of life. This will raise the bar astronomically for those who hope to build really uncanny droids.

Perhaps you begin to glimpse a massive twist in the tail of the story of the evolution of robots. But let's take it slowly. The eminent physicist Freeman Dyson, who as a young man made pioneering contributions to quantum electrodynamics, speculates that nanobiology will advance to the point where kids will buy kits to play with at home that pack enough robot power to create new genomes and grow new life forms. In his childhood, kids had chemistry sets and kits to build radios. Kids nowadays have kits to build tabletop robots.

Extrapolating, Dyson's ideas are the natural next step. Nanobiology kits will allow kids to build cool new organisms from membranes, cytoplasm, microtubules, proteins, ribosomes, DNA, codons, and so on as easily as kids once built chunks of machinery from little nuts, bolts, gears, rods, bars, and plates. Perhaps they'll be able to download gimmick genomes as freely as kids now download ring tones. Their bedrooms will be festooned with homemade luminous slime and crawling with customized beetles, and moms will regularly suffer hysterics at the loathsome creatures they discover.

The outcome of a generation or two of this sort of child's play will be a routine acceptance of the molecular mechanisms of life. As such childhood grooming always does, this acceptance will ramify into the adult world of work. Medical science, in particular, will benefit from the new perspective. People will lose their mistrust about having bits of their bodies manipulated or replaced and will understand how synthesized flesh or bone or blood cells can improve on the natural

skinfulls of body parts their mothers gave them. Many people will pay good money to have their imperfect organs replaced with picoperfect designer organs from a blandly neutral Chinese bioscience lab. Since robots will perform all the necessary operations with superhuman efficiency, medical professionals will be less stressed in terms of time allocation than they are now.

This is not yet the twist in the tail you may be waiting for, but we're getting there. Once we learn to see ourselves as bags of astonishingly intricate molecular mechanisms and to believe with the certainty of true faith that someone in a Chinese bioscience lab really understands how all these mechanisms work and how to trim and tune and optimize them, we'll be ready for our first baby steps on the road to rapture.

Robots will merge with us via the picotech of synthetic life

Each civilization in history adds something special to the store of goodies that we enjoy as our birthright. The ancient Greeks gave us geometry and democracy, the ancient Romans gave us good roads and organized religion, and the ancient Chinese gave us gunpowder and printing. Victorian Britain gave us steamships and electricity, imperial Germany gave us industrial chemistry and internal combustion engines, modern America gave us car culture and computers, modern Japan gave us smaller cars and consumer electronics, and modern China will give us friendly droids and synthetic life. By the end of the century, Globorg will put the pieces together and give us a single global organism uniting tech and life.

The twist is this. Robots won't march all over us like Terminator robots. They won't even mimic us as droids and pull the master-slave stunt. They'll merge with us via the picotechnology of synthetic life augmented with photonic or plasmonic cloud interfaces.

But this is science-fiction voodoo until we can unpack the ideas more methodically. We're still a long way from our goal here. First, what exactly is picotech? We all know what nanotech is and we know that pico is the next step down from nano, but why go there?

Nanotechnology operates with structures that have sizes in the nanometer range, where a nanometer is about five or ten times the diameter of a typical atom. So we're talking about molecular structures

and crystals and surfaces where the individual atoms are beginning to be significant. At the nanoscale, intermolecular forces described in quantum physics can disrupt things and make traditional classical mechanics unworkable. So nanoscale gear wheels and so on behave in quite new ways. At the nanoscale, air behaves like a viscous fluid, and nanoscale flying machines are like little submarines. To give you a sense of where we are practically, the finest lithography used at present in microprocessor fabrication carves features with sizes of tens of nanometers. So we're still a few doublings away from the atomic scale where quantum mechanics rules.

Picotechnology will operate with structures a thousand times smaller. We'll be down below the size of atoms. At that scale, there's no continuous stuff anymore. Atoms are revealed as tiny nuclei, down at the femtoscale another thousand times smaller, surrounded by quantum clouds of electrons that we have no common-sense way of visualizing. But since the wavelength of visible light is up in the hundreds of nanometers, we can't see down there anyway.

At the picoscale, we'll have to replace the classical logic of bits with the much trickier quantum logic of qubits. Bits are either 0 or 1 but never both. Qubits are quantum mixtures of both 0 and 1, with an infinite number of states between the classical extremes. Once we can get it to work, quantum computing will be a huge advance in the art of computation. But to get it to work reliably, we'll need picoscale control of what's going on.

Let's just underscore the huge benefit of quantum computing. The advantage of qubits is that with enough ingenuity we can use them to do lots of classical computations at once. In practice, this implies that a quantum computer can sometimes work exponentially faster than a clunky old classical computer. For example, imagine your classical computer solves your chosen problem in one second, but if you scale up the problem by factor 10 it takes a kilosecond (about 15 minutes), if you upscale by factor 20 it takes a megasecond (about 11 days), and if you go to factor 30 you have to wait a gigasecond (over 30 years). Unfortunately, this kind of scaling is typical for a wide class of problems. But your new quantum computer is unfazed by the problem size. Factor n takes n times longer, for all n. The scaling is nice and linear (or at worst polynomial). And this is a breakthrough. It lets us build much finer models of the physical processes that interest or

bother us, from weather and climate dynamics to thermonuclear plasmas and fireballs.

So picotechnology helps us out in solving problems from hell. What else can it offer us? The main answer here is effective control of molecules. At present, in chemistry and biology labs, scientists work at a level that's many orders of magnitude above the molecular scale.

The key number here to appreciate the problem is called the Avogadro constant, and it's huge, about 600 sextillions. This is the number of molecules in a "molar mass" of a substance. You can look up the details, but the upshot is that milligram quantities of most of the substances studied in a lab include quintillions of molecules. Even nanogram samples contain trillions of molecules. Going deep, deep down, a cubic nanometer of water, with a mass of a zeptogram, contains about 33 water molecules. Such numbers make precision work at the molecular level rather difficult!

For the question of how we learn to live with humanoid robots, the main point is simple. We're molecular machines and the droids will be machines that are crafted down to the molecular level too. How do we know this? We know because droids will have to be engineered that far down to compete with us. Anything less will condemn a droid to techno-doll status, where at best it offers a robotic suit or exoskeleton that we can animate with our avatars.

You may be skeptical here. Who needs picotech to learn that we're molecular machines? We know about how deoxyribonucleic acid (DNA) and ribonucleic acid (RNA) work to replicate genes coded as base pairs matching adenine with thymine and cytosine with guanine, how we can use a polymerase chain reaction (PCR) to replicate DNA, how the mitochondria within our cells generate energy in an oxidative phosphorylation process that converts adenosine diphosphate (ADP) to adenosine triphosphate (ATP), and so on and so forth through an apparently endless story, already. What's left to discover? And as for losing our mistrust of manipulating or replacing body parts, go tell that to the people who have replacement hearts or kidneys, pacemakers, replacement heart valves, false teeth, silicone breast implants, and all the other wonders of medical and cosmetic surgery. It's all happening already! So what's the big deal? Why should we hold our breath?

Biology is in a growth spurt that will ferment more advances

At the end of the nineteenth century, physicists thought they knew all the physics they needed to know. Then came relativity, quantum mechanics, subatomic and nuclear physics, and strings. Even Stephen Hawking can't keep up with it all.

At the end of the twentieth century, we thought the basics of biology were in place and secure. By then we had the Neo-Darwinian synthesis of evolution by natural selection and Mendelian genetics, with a precisely known mechanism of DNA replication to explain the transmission of genotypic information between generations. We had the Central Dogma of molecular biology that DNA is transcribed into messenger RNA, the genetic message on which is then translated into polypeptides (proteins), which make up the meaty stuff of biological organisms.

The standard story is that we're all big wet heaps or bags of proteins, great tangled, heaving, churning, lurching masses of proteins, which are constantly being recycled from within via transcription and translation from our genes. Over the generations, countless dud genes have been filtered out to give us the selections our bodies now deploy to build and rebuild themselves. Each and every one of us is the magnificent culmination of a couple of eons of evolution, bred from a line of winners. And each one of us, in the immortal trope due to Richard Dawkins, is a lumbering survival machine for his or her genes, a throwaway robot designed by natural selection to insert those genes into the next generation. This much we know, it seems, but it's not yet all we need to know.

I have no ambition in this book to spell out a new science of biology, but I can hint at the issues. Evo-devo – the science of evolutionary developmental biology – is one place to look for new ideas. The development of an organism from a bag of proteins is a hugely complex process that we understand only very superficially at present. Gene activation and inhibition, for a start, is a far more dynamic process than the Central Dogma suggests. Genes can be turned on and shut down for a range of reasons, and sometimes the new activation patterns are inheritable, so that offspring can even inherit changes triggered during a parent's lifetime. In fact, the logic that sits between the genotypic information coded on the DNA and its phenotypic

expression in a flourishing individual is such a horrifying tangle that it defies our present attempts to unravel it. We need to trace all that logic before we can say we know the story.

The complications ramify. To spare you too many of my own naïve reflections, let me refer to a qualified opinion from my recent reading. As Harvard biology professor Richard Lewontin said in a 2010 review of an embarrassing book on evolution by a pair of non-biologists, an immense amount of biology is missing in the story of natural selection, and a number of complexities at the molecular, cellular, developmental, and physiological level need to be taken into account as well. We can agree that biology is in a growth spurt that will ferment more advances for decades yet.

Meanwhile, the Darwinian picture of evolution by natural selection as the grim pruner of the burgeoning tree of life is a great "intuition pump" (as Dan Dennett describes it) to get us started. Together with Mendelian genetics, DNA, and the Dawkins tale of selfish genes, the Darwinian picture makes an excellent start on the job of creating order from chaos. But it's only the first leg of a long journey.

A century ago, Isaac Newton's picture of nature as an absolute space-time filled with countless tiny particles in ceaseless interaction according to the laws of Newtonian mechanics, like a divine game of billiards implementing the celestial clockwork, still worked as a great intuition pump for physicists. But within a few decades the "new physics" of relativity and quantum mechanics had replaced it. Biology could undergo a similar transformation in this century.

Computer modeling will become the main way to do science

Today, we have a new instrument that will soon change everything. The new instrument is massive computer simulation of any process we want to understand better.

Until recently, there were two main ways to do science: theory and experiment. Theorists sat in armchairs and put pencil to paper, like mathematicians, to sketch out great castles of ideas. Experimentalists tested these ideas. Some of them passed and some failed. New facts emerged, and theorists used the new facts to rethink their ideas. Experimentalists tested the new ideas, and so on. This game of ping-

pong went back and forth for generations until good sets of ideas evolved from an increasingly rigorous selection process.

Now, with computers, we have a third way to do science, namely to build models and see how realistic they are. The modeling approach is a middle way between theory and experiment. Unlike theory, it can reflect all the data all the time, and unlike experiment, it can keep the big picture in view all the time. As new generations of scientists learn to use modeling tools more effectively, modeling will gain in importance until it becomes the main way to do science. It's already the main metaphor to describe what scientists do. They model nature.

All the sciences are applied philosophy. Philosophy is the fertile soil in which the exact sciences take root and grow. Physics and astronomy grew from philosophy in classical times. Astronomy took off as a separate science with the invention of telescopes. Galileo used his new telescope to observe the four biggest moons of Jupiter in 1610. Newton published his theory of celestial mechanics in 1684. Biology took off as a separate science with the invention of microscopes. Antony van Leeuwenhoek used an early microscope to discover pond microorganisms in 1676. Carl Linnaeus published his taxonomy of biological species in 1735. Charles Darwin set sail on HMS Beagle in 1831 and published his theory of the evolution of species by natural selection in 1859. In both cases, new tools enabled scientists to dig the garden of nature in new ways and bring forth beautiful new plants from the rich mulch of philosophy.

In our era, the new tool is the computer and it transforms all the sciences. Modern physics and biology would be unthinkable without powerful computer resources.

For example, the giant CERN Large Hadron Collider in Geneva, Switzerland, which is currently in search of supersymmetry and the Higgs boson, generates petabytes of new data every second through millions of data channels. This data goes out to a global network of many thousands of servers that together constitute a massively parallel virtual machine for crunching all that data.

As another example, the IBM Blue Gene petaflops supercomputer series can calculate many things, but it was designed especially to run big bioscience jobs that need lots of raw power, such as protein folding simulations in the new field of proteomics. In the similarly compute-intensive field of connectomics, the IBM Blue Brain project headed by

Henry Markram at the Ecole Polytechnique Fédérale de Lausanne (EPFL) in Switzerland aims to improve our understanding of how neurons in the brain can generate the stuff of thought. The project involves running a very precise simulation of the ten thousand or so neurons in a pyramidal column, which is the basic architectural module in the mammalian neocortex. Separately, in late 2009, IBM researcher Dharmendra Modha announced that his team had created a high-level simulation of the neural connections in a cat brain. This news incensed Henry Markram, who pointed out that the simulation was far less detailed and realistic than his own. The result was a classic scientific catfight.

In sum, the scientific opportunities that lie before us as we use this new tool are mind-boggling. We can say without exaggeration that the era we're living in is a golden age of science.

Droids and people have a lot in common

To return to the argument, the big merge of technology and biology looks less strange when we see what these new resources bring to the task. Technology won't look the same when we get going in earnest with robotics and quantum computing. And biology won't look the same when it's been enriched by a few decades of results from supercomputer simulations. Ideas that seem now like science fiction or voodoo will soon look quite tame. The idea that future droids and medically optimized or enhanced people have a lot in common won't seem strange at all. You don't need to believe this, but your kids or grandchildren will.

Returning to the more practical question of building droids with an uncanny resemblance to people, we can ask what picotech might bring us that nanotech won't or can't. Nanotech is now, for the next decade or so, and whatever we can't do in the foreseeable future to make droids we can't do with nanotech. And obviously there's a lot we still can't do. Living organisms are built from the bottom up, with every molecule a wanted molecule. Our robots aren't, not by a long way. That contrast conceals a fundamental issue of principle.

There's a standard simile here that still has a lot of mileage in it. Lego bricks are really versatile plastic building bricks for kids. Invented in their present form in Denmark in 1958 and made there

ever since, these bricks can be plugged together to make a fantastic variety of shapes and structures. The range of basic bricks is quite small, but there's a wide range of more specialist pieces that work with them, for example to build rotating machinery or hydraulic or electric features. Garden-sized model villages, exhibition displays ranging from landscapes to aircraft, and working models of factory layouts are just some of the things that people have made with Lego bricks. So Lego bricks are a bit like molecules.

Conversely, protein macromolecules are a bit like modular sub-assemblies made with Lego bricks. Apparently, we humans are made from about fifty thousand different kinds of protein molecule, which between them and together with water and so on add up to a total of what I guess to be hundreds of septillions of molecules making up each human body. So far, the manufacturers of Lego have only made some hundreds of billions of bricks, not yet enough to build a giant Lego man with his head up in a geostationary orbit (so we could employ him as a space elevator, albeit a rather wobbly one), but the idea behind the simile is quite close.

The fundamental issue of principle is that organisms like us are not only built from the bottom up, brick by brick, but also designed that way. Or at least, we're designed that way to the extent that it makes sense to speak of design at all for an evolved species like us.

Imagine a vast number of little demons down at the molecular scale, busy building up protein molecules to make an organism. The demons know pretty well what to do locally, because only some combinations of proteins fit together at all, and they have some idea of what the combinations should do, for example what little mechanism a protein assembly should embody, because they get a little demonic reward if it works. Also, they work in highly disciplined teams that have inherited and long-established traditions of working together to build certain higher-level structures, so tissues and organs tend to emerge quite well from the work of all the teams. But the demons don't have a clue about the overall plan for the organism. They all just do what they can and hope for the best.

Everything we know about organisms suggests that there's no master plan and no master architect. All the optimization is lower down in the stack. The upper levels are hardly optimized at all. How could they be? They don't know what to optimize for. Natural selection

is still working on the upper levels. To recall our discussion of war between nation states, most biological organisms are still in a state of nature in their relations with each other. It's a war of each against all in the Darwinian jungle, with new horrors every day. Optimization is a struggle, not a result, until the war is over.

Our robots are quite different. They're built from the top down. This is true quite generally for the machines we make. We decide what we want and make the machines accordingly. At first, we do a rough job. We achieve a first-order approximation of the machine. It does what we want at the highest level, but below that it's a mess. Recall the very first transistor at Bell Laboratories in 1947. It looked like a stone-age spearhead crudely soldered onto a small slab. But it worked, just. Since then, we've optimized transistors more and more deeply, until now they're tailored down to the nanoscale and we can almost count the atoms and molecules in their working parts. Robots will be similar. We'll optimize them more and more deeply until one day, after the glorious dawn of picotech, we reach rock bottom (or rather quantum bottom).

Perhaps you can already see how we achieve convergence between the bio and tech worlds. To optimize anything, you have to understand it thoroughly. For organisms, that means you need good computer models. For robots, it means you need good working models. Somewhen in the course of this century, the models will converge in a quantum hypercomputer somewhere, and biotech convergence will ripple outward to become an accomplished fact. Uncanny droids will have arrived, and they'll be us.

1010
New Life

Creating new life in picotech is the scientific challenge of the century. New life will transform our hi-tech lifestyles by giving us sensitive things. We have a few sensitive things now, such as the touchpad on my laptop machine, but they're still simple and rare. It's worth taking the time to imagine how we can move toward a world where we've grown lots of new living things with touchy-feely features that are designed to help us.

Let's start with a bold statement. What makes humans different from machines is sensitivity plus intelligence. The combination explains more of our mental lives than you might imagine. Intelligence we all know about. Machines with intelligence already outperform us so well that we've given up the fight. But most of our machines are still clumsy brutes that aren't safe in polite company without a human to guide them. Any sensitivity they have – like that of my touchpad or the orientation awareness of an iPhone – is an island competence in an ocean of inert being. Even the most challenged human idiot savant is a miracle of multifaceted geniality in comparison. People react in a million little ways at the slightest provocation. Many of these reactions are unconscious and serve only very minor purposes, but their accumulated effect is decisive for our sense of being human. For example, they create the feelings and emotions that form the cradle for our consciousness. My bold claim is that once we unite pervasive sensitivity correctly with sufficient intelligence, we'll have cracked the problem of building uncanny droids. Unfortunately, it's a hard problem.

A hard problem is not an insoluble problem. Take an example. Australian philosopher David Chalmers, the zombie blues man, thinks that consciousness is a hard problem in that it's hard to see even in principle how we might implement it in a machine. But if we forget for a while about consciousness and think instead about the self, we can make progress. The concept of self is tricky and puzzling, but it's not too foggy to handle logically. Then the wider concept of consciousness begins to look clearer. So a hard problem is often soluble indirectly or

incrementally. We chop it up and replace it or approximate it with easier problems and solve them, then find that the original hard problem now looks easier.

Giving machines feelings and emotions certainly looks like a hard problem. It looks as hard as giving them consciousness. But giving machines more sensitivity sounds much easier. Make the whole surface of a machine a touchpad, fill it with various sensors, and then try to integrate all that new data in the game plan for the machine. That sounds straightforward enough. Then look back and see how far we can say we've given the machine feelings. My guess is that we'll have no big problem there. Then we implement emotions as linkages between feelings and action plans. I predict that after a few generations of such machines have come and gone, we'll say the issue of feelings and emotions in machines is solved. And we'll find that the droids we make are as uncanny as we want them to be.

Now you know where I stand, but not yet how we're supposed to get to the Promised Land or what picotech has to do with it. Touchpads and so on are implemented in boring old nanotech. Why do we need to descend to the atomic and quantum level?

The brain is a feeling organ that feels its own neural states

The answer to this question is the story of an evolutionary race. Evolution favors organisms that feel as sensitively as the environmental cues allow. Their feelings are selected in the struggle for survival to be as delicate as physically possible, consistent with preserving robust and reliable mechanisms to process the feelings for control and coordination. Somewhere in any organism, the feelings must somehow come together to create an overall feeling of being that the organism simply maintains in a kind of ongoing homeostasis. Evolving organisms will race to the bottom in their attempts to enjoy finer feelings than their competitors. The natural bottom is molecular, at the transition to the quantum world where incremental refinement of analog signals is no longer possible. In other words, evolution refines organisms in a natural version of the Moore's law descent we see in generations of integrated circuits. Since nature has had a long time to run this race, we can take it that it's reached the bottom, unlike our own pursuit of miniaturization. So we'll need picotech to catch up.

The problem of reconstructing feelings in humans reveals why this is a hard problem. Human sensory mechanisms aren't all in peripheral systems like hands and eyes. The brain is a feeling organ. This doesn't mean the brain registers pain when a surgeon slices into it, because it doesn't, but it feels its own neural states. Moreover, it does so very delicately. A feeling causes a signal to erupt into consciousness. The signal registers a salient event, which may be extremely small. Humans have prevailed over other creatures in terrestrial ecosystems because of their intelligence. Humans use their intelligence to distinguish their own feelings, or rather the neural states triggered by those feelings. Our memories are traces of neural states, and we can distinguish billions of them. Humans made themselves fitter by recognizing a wider range of more subtly different environmental configurations. They did this by instantiating neural states that reflected those configurations with astonishing precision and then remembering those states as finely differentiated memories.

Consider an example. A nasal sensor registers a scent by matching an inhaled scent molecule to a molecular template and sending a signal to the olfactory bulbs in the brain, where a little group of neurons fires a tiny fusillade of signals in response, which we can imagine as a musical trill. At the same time, the eyes register a scene that's related in some way to the scent. Another group of neurons in visual cortex fires a similar fusillade, and yet another group of neurons in the entorhinal cortex responds to the simultaneous trills fired off by the olfactory and visual neurons. The trill is then recorded via the hippocampus as a trace that other neurons in the brain can activate by firing off a key that matches the trace and the trill. Here we have a cascade of events with delicate feelings at each step. The feelings are as delicate as the electrical patterns in the neural signatures can be. I guess they border on the quantum threshold. And it seems obvious to me that we'll need pretty good molecular engineering to recreate this sort of mechanism in a small, energy-efficient hi-tech module.

The picoscale approach is straight from molecular biology

Engineering gets more indirect as the target systems decrease in size. The unaided human hand is a wonderful tool for handling centimeter-sized objects, but it's already challenged at the millimeter scale. For

the micron scale, you need microscopes and elaborate analog mechanisms for translating grasping movements down to the object size. For the nanoscale, you need a more sophisticated imaging and manipulation technology, such as atomic force microscopes or scanning tunneling microscopes. For the picoscale, we'll need a different approach again.

The approach that comes to my mind here for the picoscale is straight from molecular biology. We'll use variants of the molecular machines that our cells use to do chemical engineering. We'll build something like a ribosome, which reads a template and assembles amino acids to make the protein coded by the template. In our cells, the templates are written on messenger RNA, where the instructions are coded in copies of the genes on our DNA, but once we get going with the technology we can probably change to a more tech-friendly storage molecule. The ribosomes can be seen as molecular Turing machines that use messenger RNA as their input tape and write peptides or polypeptides (proteins) as output. Again, once we get more confident with the technology we can make the machines differently and design them to assemble any molecules we like. If it helps, we can picture such machines as little 3D printers that print any molecule specified by their input data.

To move our newly synthesized molecules about, we can make dedicated nanotransporters with a sticky docking molecule to hold the payload molecule and a couple of winglet molecules to row the transporter to the destination (complicated molecules are nano even if the technology needed to make them is pico, just as contemporary microprocessors use nanotech). And so on – you can extend the fantasy in many ways once the first steps are clear.

Designing all these molecules and mechanisms will be child's play on the workstations that rich kids will learn how to use while they're still at school. It'll be computer Lego. But the final product, as it pops out of the picoprinter, will be completely invisible, so let's hope that by then someone's invented a reasonably user-friendly imaging technology. Perhaps the nanotransporters will put up flags to indicate their load as they row past a picoreader. In any case, different nanotransporters will be programmed to dock at their own sites on a preconfigured substrate where bigger assemblies are built up. Once an

assembly becomes so huge that we can see it with a light microscope, we can test its functions relatively easily and chuck out any duds.

That was a lot of science fiction, but I hope it sounds vaguely plausible. Once we understand such molecular mechanisms, we can build facilities that automate the process as far as we like, so that people can sell the resulting assemblies by the milligram without bothering to count them. The business of fabricating such nanomachines will probably look something like a cross between the "chip fab" and the drug businesses of our time. Or at least that's the only analogy I can think of in the present industrial scene.

Picofabs will enable touchy-feely machines

The point of all this picoscale enterprise is to create the enabling technology for touchy-feely machines. These machines will be so richly equipped with sensors that they'll have plenty of input to monitor their state and their immediate environment. With embedded nanoprocessors, they'll also have the power to integrate the sensor input on a continuous basis and maintain a dynamic self-model, if not permanently then at least as long as the machine is up and running. (We saw in an earlier chapter that a self-model is just a code module that includes a set of parameters indicating the machine state.) The machine uses the self-model to optimize anything relating to its state that it's in a position to control, much as we use mirror images to trim whiskers, pick at spots, comb hair, and so on. The machine will be able to monitor things like its temperature, the surfaces it may be touching, the sounds and smells in its neighborhood, the view in various directions, any vibrations or accelerations that affect it, and so on.

The benefit we get from enabling a machine to do all this is less from any particular reading (how often are you going to search its log for bad smells on the machine farm in the culture dish?) than from the overall record. The machine will aggregate and compare the results in a thousand ways and prepare summary reports on demand, for example to assist in debugging or on-dish troubleshooting. Also, it'll always be ready to complain if it feels the need, for example about working when it's too hot or about cigarette smoke during night shifts.

We need to locate the payoff when we build such picotech devices into more familiar machines. For a start, there's undoubtedly some

market potential in the ability it gives us to interact appropriately but unconsciously with machines. A machine that recognizes me from my footfalls or my handgrip and greets me with a personalized response is a welcome prospect. A seat or a bed that senses my weight and movements and adjusts its softness and warmth to suit me is better than a hard stool or a cold mattress. A sniffy loo that smells the sitter's poo and alerts the refrigerator to adjust household grocery purchases could improve diets and help to reduce greenhouse gases. A toenail robot for people who can't bend down or who want a neat paint job on their nails would surely find toe-troubled takers. A little mouth robot that climbs around my teeth polishing them and pumping up my gums with stem cells would do wonders for my smile and perhaps please my dentist. And tiny earbots that screen off unwelcome noises, play soothing music, and translate foreign languages will sound familiar to fans of Douglas Adams novels. As for incubus (male) and succubus (female) sex robots, perhaps the less said here the better. But these are all just stray ideas, and even sexbots aren't radical or important enough to launch a trillion-dollar industry.

Birth and death will be automated with feeling machines

A bigger bonanza here is in the medical sector. Anything that relieves medical doctors and surgeons has a big future in a world where many more people can be expected to live into ripe old age. Droids with a soothing touch and a smooth bedside manner will be a huge hit. Every senior citizen will want a personal Angel of Mercy to ease those last undignified years. And for that market, no tiny increment of softness and gentleness will be too small to be worth engineering into the product. This is the idea the Japanese robot engineers have caught onto. These are still early days, but the long-term vision is surely right.

In hospitals of the future, sensitive robots will dominate the operating theaters. At present, surgical robots do fine work in keyhole surgery, for example, where superhuman precision is needed. Such robots are directly slaved to a human surgeon, who performs the original of every move live through a video interface. But remember the spray-painting robots in car factories, where the master painter performs the job once, perfectly, and the robot simply does the same again, either forever or until it's reprogrammed. We can offer some-

thing similar for surgeon robots. A surgical robot could let itself be slaved to the master surgeon for a few operations, to learn the moves, then cut loose and work on its own initiative. It would need to be smart enough not only to repeat the standard moves but also to adapt them to different patients with various shapes and sizes and so on. It would also need to be extremely sensitive to surprises. It should know what to do if it thrust its camera and scalpel appendages into a patient only to find there was no problem or a different problem. It would need great anatomical smarts and a supreme ability to think on its trolley, but this is all part of the deal and well worth the price.

How about start of life's journey? Gynoid robots could naturally make themselves useful in the obstetrics and gynecology wards, especially in the birthing process, if there were too few people ready or willing to help, but that's not all by a long way. As any new parents will testify, taking care of babies is a hugely labor-intensive process where occasional help is often welcome, if only to change diapers. There a gentle bot-bot would probably be the salvation of many otherwise excessively pooped parents. Not only babies but also small kids could benefit from having a Little Miss O'Pair gynoid to help out from time to time. Hmm, maybe not – the au pair role demands a level of situational intelligence that would certainly devour many generations of gynoids before they were tantrum-ready.

But one role here would be relatively easy for a protein-based machine with feelings. Artificial wombs could liberate ambitious women with careers from the burden of months with a bump and the pain of pushing out the product. A surrogate womb could be pico-engineered to match the target mother in all relevant respects and programmed to provide the growing fetus with all the nutrients, movements, vibrations, and so on that a caring mother-to-be would normally and naturally confer upon the bun in her oven. In fact, it wouldn't be hard to improve upon the start that some mothers give their babies. Once the technology had matured and proved itself, I can imagine properly controlled and peer-reviewed studies showing that babies from wombots generally got a better launch than those from most mothers.

I guess that wombots will only be deployed on a small scale for very rich people this century. It'll take us many decades to get the technology right, and the payoff is too marginal in most of the world,

where the more urgent problem is that too many women still seem all too ready to have babies. But in ageing societies like Europe and Japan where the present birthrate is well below replacement levels, wombots could make all the difference. The economic opportunity created by enabling more women to pursue more self-focused careers would comfortably finance mass rollout in the longer-term future. Once the technology has been piloted and becomes widely understood, I can imagine a rampant feminist movement arising to liberate women from the indignities of pregnancy and childbirth.

Clearly, medical developments will be big. Birth and death will be automated with feeling machines, at least in communities where the investment in such "meat" machines (as people will naturally see machines based on proteins) seems preferable to employing people with unreliable politics and personal habits or to importing whole communities of poor people who need extended and expensive introductions to rich-world life. Perhaps as a matter of anthropocentric principle we should invest more in good education instead, but that lies outside the scope of my study here. Given the world we live in, the pressure from medical needs is quite sufficient to fund a big push in the direction of sensitive droids. And once that's achieved, once the robotics developers succeed in integrating a full sensorium into a droid at the pico level, we'll be ripe for much wider application of the underlying technology.

The wider application areas that impress me are in the field of artificial or synthetic life, which I suggest we call New Life, if only to make it easier to say. These are the areas that impress scientists like Freeman Dyson, who imagines A-life kits for kids to play with.

Already we can adapt or re-engineer organisms quite radically, to make new bugs for eating or excreting oil or for synthesizing drugs or to misuse rats or pigs for growing ersatz human body parts. But the real revolution comes when we create completely new organisms.

Bugs could make enough oil to replace fossil fuels worldwide

The enterprising life scientist Craig Venter, whose startup Celera Genomics sequenced the human genome (in a first draft) faster than the U.S. government's Human Genome Project, is currently creating new microbe chromosomes at the molecular level. To speed up the

rollout, his team pops the synthesized chromosomes into the gutted body of an existing species of bacteria. The new chromosomes are essentially software and the body offers the machinery to boot up the code. Once the code is up and running, it creates new proteins and rebuilds the cell from within. So Venter's team is creating new life from the bottom up. Using a technique they call combinatorial genomics, the team's robots can fast-forward evolution to find new genomes with interesting chemistry. Their current project will ship bugs that use sunlight to convert carbon dioxide and water into usable fuel oil. Once these bugs work as intended, they could make enough oil to replace fossil fuels worldwide in a few decades.

This is just a start. As Craig Venter said in a press statement for a new cell dubbed Synthia in May 2010, "This is the first synthetic cell, a proof of concept. But the proof of concept was key. Otherwise it's just speculation and science fiction. This takes us into a new world." In the new world, we only need to add a few megayears of evolution to see humans as machines, too.

Evolution is fast becoming a technology. Evolution programming was already a well-understood field twenty years ago when I edited academic books on the subject. The idea is that program code can be replicated with variation. Each generation of variants is tested and the winner used to populate the next generation. Within a surprisingly small number of generations, the technique can create good enough code to solve problems that old-fashioned methods would find it hard if not impossible to match.

We can use the technology to evolve new genomes. The current challenge is modeling the proteins that we want to manufacture from the genome code. Once we can predict the folding patterns of new proteins, we can program the genomes to manufacture the proteins for any chosen Lego configuration. Then the next challenge will be to model the lifecycles of organisms built with those genomes and proteins. All this work in proteomics and organomics requires massive computing resources. As you can see, this is an exciting field where huge opportunities are still up for grabs. Anyone with the courage to grasp them will be able to make new tools that change the world more radically than old-fashioned robots ever could.

To highlight an obvious example, if we can build meat machines we can build meat. We can make all our food in factories and phase out

the inefficient growth of organic foodstuffs on traditional farms. We may need to work harder to generate enthusiasm for the new diet than to make the food itself, but the option will be there.

Again, the medical examples offer the most intriguing opportunities for further speculation. We shall build new microbiota to populate our bodies. Human bodies contain some ten trillion human cells and ten times as many microorganisms, although since these are mostly very small they make up only one or two percent of the total body weight. Alongside the likely majority of useful organisms, for example those that help us digest food, there must be a large number of useless parasites and positively unwelcome fellow travelers, such as the bugs that rot teeth and cause body odor. A few engineered organisms could clean up that population and make sure that every bug on us is a good bug. We shall doubtless do more, and build repair bugs to clean out clogged blood vessels and eat tumors before they grow too big. We'll probably do this sort of repair job by injecting microscopic fleets of nanobots, but to manufacture them we'll need to wait for picotech, so that we can build them as self-reproducing organisms that are mature enough to work with feeling and kill themselves when they're done.

Many people have covered this field. For example, Singularity prophet Ray Kurzweil and medical doctor Terry Grossman covered it in their 2004 book *Fantastic Voyage*, with an update on the practical lessons in their 2009 book *Transcend*, and I can add nothing new to that.

Once we're addicted to gene updates, the provider can lock us in

The aspect of the medical story that I find most fascinating is the prospect of revising and updating our genomes during our lifetimes. It's no surprise now that we can upgrade the human genome and create offspring sporting new genes. We can build superhuman kids and watch them grow up like aliens among us. But with good picotech we can move faster.

The human genome is a stored file of recipes for making proteins. It's also a junkyard where an amazing amount of useless or parasitic garbage has accumulated over the last two billion years or so. We're not so sure that all of it's junk that we feel free to chuck it all out without more ado, but we're pretty sure about a lot of it, because we

can see it's viral stuff that just came along for the ride or got left over from a prehistoric pandemic. So we could be the first species in Earth history to clean up the mess and slim down our genome to something more elegant. On the way, we can tune it to do what it's supposed to do better and not to do the nasty things it sometimes does in naturally genomed humans. And we can build in new features and functions to fill all the space we liberate by throwing out the junk.

Imagine the fun we can have! We can build in standby genes for all the body variants we want, from rainbow hair and blue skin through long legs and big boobs to Vulcan intelligence and devastating wit, with regulatory genes for them that could be triggered by dietary supplements so that you could change your body type from season to season, at will, by popping a pill. You could add genes for new brain regions with prewired skills and memories so that you never had to train or study again. As first priority, we'd spruce up the genes to undo the damage of old age, so we'd all live either forever or until an error in a new gene package did us in.

The enabling step for all this fun is to find a way to upload new genes. We'd need a mass rollout mechanism as easy to run as Windows updates, so that users could update their genes on the fly, to suit the mood of the season – or to correct the errors in the previous upload.

Uploading new genes is just a start. The body needs a mechanism to distribute the new genes and replace the old ones. Here we need some picotech inventiveness. Natural humans have a copy of all their genes in every one of the ten trillion cells in their bodies. Most of the genes go unread in most of the cells most of the time because the cells only read them to make proteins, and most cells only need a small subset of the proteins on offer. Just as most users of Microsoft Windows and Office don't use most of the code most of the time, so most of our cells get by with very little. But Mother Nature found it easier to distribute the full pack to every cell, just as Microsoft offers users all-in-one packs. By contrast, our future pico-tweakers of the genome will find ways to be more selective.

In a scenario that seems vaguely plausible to me, update nanobots will cruise in the bloodstream to visit any cell they like, but they'll be told which cells are likely to express the new genes they have on offer and will target them for their cut-and-splice operations. The update

nanobots themselves will be generic carriers for update molecules that are written in a separate update capsule implanted under the skin. A point-of-sale droid will have implanted the update capsule in a ten-second operation when the customer signed up for the service. The capsule will communicate wirelessly with an update station that publishes new genes regularly via encrypted links over the public web. In the capsule, picoprinters will write the update molecules for the update nanobots to ferry to the appropriate cells. And so on – again the details are straightforward enough once the basic setup is clear.

The twist in this tale is in the nefarious misuses we'll discover for the whole technology of gene updating. Once a large base population of humans is addicted to a nightly fix of gene updates, the gene provider can lock them in with proprietary gene extensions. Many of the updates will simply be fixes to block virus attacks. Some viruses will hijack the update capsule to write anarchic gene updates with wild phenotypic expressions. Perhaps victims of a virus attack will sprout feathers or grow breasts on their buttocks. Other viruses may simply kill their victims. In any case, addicts will become seriously dependent on their update fixes, and the service provider will be legally required to accept a challenging service level agreement.

You can see that we're navigating far-out science fantasy territory here. This sort of technology is most unlikely to be mature enough for mass rollout within this century. But who knows? A governing organization like the Communist Party of China might find the prospect of genetic control so full of revolutionary potential that it mounts a major drive to master it. Perhaps the CPC will then wage war with hostile foreigners by distributing traditional (biological) viruses to wipe out anyone not protected by the CPC proprietary gene update. If this kind of warfare became pervasive between gene update service providers, it would be bye-bye feral humans. The tweaked-gene carriers would inherit the Earth. But let's leave this fertile field of nightmares for a future century.

Imagine planting a gene pack to grow a frame for a house

The new life that will probably preoccupy us this century will be much more humble. Bugs to make oil or rubber chicken will come soon enough. Then we'll recapitulate the Cambrian Explosion and engineer

multicellular organisms with a variety of body plans limited only by what the market will bear. Living beds and sofas will be relatively easy. A living bed might feast nightly on skin flakes and body moisture (to put bedbugs and dust mites out of business) and launder itself automatically. A sofa might binge on dropped pizza and guzzle spilled beer. A carpet might drink piped water to heat and humidify the room and recycle fluff to repair scuff damage. Then we might make living clothes that grow to fit all sizes, repair any holes or rips, digest perspiration or food stains, change color or style or bulk on request, use body warmth or movement to power your hypervision headset, and so on. You can probably invent a few more such ideas.

More adventurous developments in this direction will come later. We may have to wait many decades for houses that grow themselves. Imagine planting an acorn-like gene pack and letting it grow over a few weeks into a fine polymer frame for a house. Then you watch for a few more weeks as armies of slug-like creatures build a pearly shell over the frame. Specialized teams of creatures create custom windows and doors following plans you selected in the grower's catalog. Paintbots do the walls and waxbots do the floors. Electrobots grow like vines through the house to install wiring. And so on – in a couple of months you move into a beautiful house that suits your every whim, knows you personally, and follows your comings and goings with infinite patience. Your house will be the ultimate feeling machine.

New life on this scale will follow soon after we've improved our own bodies thoroughly. We'll be limited by our imaginations at first, as we treat ourselves to flawless skins, strong muscles, smooth fat distributions, elegant hands and feet, and so on. Our basic body image will remain intact as we fiddle at the periphery. Meanwhile, medical advances will ensure that we resist all known diseases and enjoy optimally efficient hearts, lungs, guts, and all the rest as we learn to expect healthy lifetimes stretching well beyond a century. Getting used to this much paradise will probably keep us busy enough this century. We can be sure that enough will go wrong along the way to stop us advancing too fast toward the gene-war dystopias of century 22.

What can go wrong here is clear enough from the idea that many of us will live long enough to experience the new thrills and spills of century 22. Radical life extension looks easier every day, as new discoveries about the role of nutritionally optimized diet, management

of exercise and stress, gene therapy for life-limiting errors in the genome, telomere extension for prolonged cell division, and more pile up for implementation in contemporary lifestyles (recall the cited books by Kurzweil and Grossman). As the effects of these changes kick in, we'll have to reorganize our entire society to accommodate oldies more constructively. There'll be no sense retiring everyone in their sixties – that sort of nonsense will soon be discarded as we learn to conscript centenarians to do chair-bound telework at home, selling their online services to the highest bidder through the Globorg market. So if you think you'll still be around at the end of this century, ask yourself – would you, then, with all your hard-earned experience of Internet hanky-panky, be willing to update your genes online through the Picosoft iGenes website? I thought not. We won't get that far this century.

Why not scrap old genomes and their bodies?

You may have noticed a rather glaring contradiction (in the Marxist – or rather Hegelian – sense) here. We scrap machines at a moment's notice just as soon as they cease to do exactly what we want, or even when they merely cease to do so better than any other machine we can afford. We see humans more and more as machines with picoperfect parts that work to realize a very roughly optimized life plan, if not a shambles of a life, despite their molecular perfection. And we imagine that more and more people will live into an unnaturally ripe old age with the help of increasingly robotized medical services. But, hey, instead of patching old genomes and the bodies that drag them around, why not scrap them and fill their social slots with smart new models packing state-of-the-art tweaked genomes for the very latest body styles and colors? Why not set up cash-for-clunker programs to get those oldsters off the road? Why not do as my increasingly geriatric novelist friend Martin Amis suggested in early 2010 and put booths on street corners where clapped-out or just jaded oldsters can go in for a "Martini and a medal" on their way to paradise?

It might come to that. Only the Judeo-Christian ethic and similar schemes hold us back. Modern commercial enterprises cultivate frenetic cults of youth. Many IT companies celebrate the relative youth of their workforces as a competitive asset in an industry where eager

youngsters in Bangalore and Shanghai can do the work as well as wizened old birds who'll soon be drawing company pensions. And where modern companies go, there other social organizations will go in future, as we saw a few chapters ago. So we can expect increasing pressure on oldsters to accept performance-related rewards and bite the bullet if things get rough.

Euthanasia laws allowing merciful termination of lives that have lost the "quality" that once made them a joy to live and to behold are a practical certainly in all advanced industrial states in the next decade or two. Extending these laws to allow people to terminate themselves gracefully (not disgracefully, as now) will follow. Public provision of mercy clinics where registered underperformers can sign on for a death pill and a financial payment to the bereaved won't be far behind. We may well find it easier to rethink our ethics than to endanger the solvency of our pension and medical systems in face of mounting government deficits and so on (you can write your own story for the ethical revisionism here).

The only practical way out of such creeping horror stories is to generate ever-increasing prosperity powered by ever more productive machines. I'm trying hard here to spell out how we can do this, but anyone can see there are gaps and weaknesses all over the upbeat story. We'll need a lot of luck as well as persistence to push on far and fast enough to leave the horror stories behind. This century will almost certainly feature enormous machine-driven progress, but we can be sure it'll have its share of atrocities, battle damage, catastrophes, disasters, epidemics, foul-ups, genocides, and so on through the alphabet. Those of us who survive can draw up a balance sheet in century 22.

But as true-grit Brits have said since time immemorial, always look on the bright side of life. In a society with sufficient solidarity and community feeling, death loses much of its sting. This has always been the principle behind military castes, where traditions and rituals to reinforce group solidarity and banish individualism have the effect of preparing soldiers to give up their lives for their comrades in arms. Christian communities celebrate their solidarity as they fight the good fight to the point of willed martyrdom for the Christian God. Death cults generally offer narratives of solidarity or groupthink to encourage elective death. No doubt the cells of our bodies have their own

incomparably tinier analogs of such practices to trigger aptosis or programmed cell death (PCD) in cells that no longer serve a purpose in the body.

Living processes realized in generations of cells of any sort always have something like PCD. The generations come and go, live and die, to realize the turnover that constitutes a living process. Data comes and goes in processor registers, molecules come and go in biochemical processes, cells come and go in bodies, bodies come and go in companies, companies come and go in nations, nations come and go in civilizations, and civilizations come and go in the evolution of life on Earth.

To be sure we've got the idea here, let's review some national and civilizational examples of PCD. Nazi Germany underwent PCD in 1945 because the ideology of the Third Reich was programmed to self-destruct in an orgy of military violence. The Soviet empire underwent PCD because the ideology of Soviet communism was programmed to self-destruct in an orgy of unproductive military spending. Militant Islam is undergoing PCD now because its ideology programs jihadists for martyrdom in an orgy of enraged struggle against impossible odds. To continue the series, Globorg will undergo PCD if and when its ideology of value added through work explodes in exponential hyper-inflation of the currency.

The death of Globorg is only so easily foreseeable if the definitions of work and value are too narrowly centered on humans, or too anthropocentric, as a biologist can easily understand. Human beings are glorious and magnificent creatures, no doubt, and our species is the jewel in the crown of nature, no doubt, but life on Earth embraces a lot more than us. Biologists will confirm that beyond humans there's a lot more to come on Earth.

To back up this claim, we can refer in passing to the works of the nineteenth-century German philosopher Friedrich Nietzsche. Famously, or notoriously, he proclaimed the doctrine of the "superman" – which for illogical reasons became associated with the Aryan "race" in the racist ideology of National Socialism. Nietzsche introduced his doctrine in what was essentially a spoof of the ancient Persian cult of the Aryan prophet Zoroaster (a.k.a. Zarathustra). Nietzsche wrote a fictional but superficially plausible account of the pronouncements of Zarathustra prophesying the coming of the supermen as a species that

outshone the race of men that lived around them. A racist reading has the supermen emerge as a pedigree breed of men with outstanding martial and predatory skills from the ordinary run of fallen or sunken men. In Nietzsche's telling, the common run of European men in his time had been corrupted or sunk by the "slave morality" of Christianity and by the consumption of alcoholic beverages. The supermen would be magnificent pagans, answerable to warrior gods and possessed of extraordinary discipline. The natural extension of this creed in Nazi thinking came with a version of social Darwinism whereby only the fittest survived in the struggle for existence. So, went the argument, those who wished to become the breeding stock for the supermen would have to grasp the courage to enslave or exterminate the lesser races around them. The rest is history.

Continuing this historical review, Marxist-Leninist communism prophesied the final victory of the working class in its historic struggle with the capitalist incarnation of monetary power as a repressive social order that essentially stole surplus value from the workers. In revolutionary struggle, the workers would expropriate the expropriators and establish common ownership of the means of production in a workers' paradise. The workers would achieve this through a disciplined party that represented them in the struggle to steal back their productive assets. But the party stole the assets from the workers too, who pretended not to notice by pretending to work on. The party pretended not to notice that the workers weren't really working by pretending to pay them in worthless currency. So the Soviet state became a kleptocratic farce where everyone stole from everyone else. The devalued system became unable to sustain the military expenditures it needed to defy the capitalist world.

Militant Islam is quite different. Its ideological roots go back to the origins of monotheism but its recent metastasis is a new development, with no precedent in the pacific monotheist traditions. Recalling my possibly controversial critique of monotheism, the psychology of the believer is split between a mundane self centered in the believer's body in the usual way and a "superself" located supernaturally beyond it, indeed beyond the entire created world of space and time. This superself is the believer's internalized representation of the monotheistic divine presence. You recall that Moses understood the "I am!" he heard to be an utterance of that divinity. Ever since that time, pious

monotheists have cultivated an inner self that projects or refracts the divine presence into their lives. So the believer's superself becomes an idol. The idol seems to have superhuman powers. The fanatical believer begins to believe that the superself can act via the ordinary self to invest the mere human with superhuman powers. So in the righteous struggle against infidels, the fanatic imagines that by invoking the power of the almighty he'll magically acquire the ability to conquer all foes. (I recall seeing this psychology at work in my Iraqi military students.) Again, the rest is history, or at least it probably soon will be.

The salvation of Globorg lies in biological science

Now we return to Globorg. We saw that the Globorg economic model assigns value to tasks as consistently as the model allows and hence forms a global marketplace for goods and services. Since this marketplace serves human ends, it can track value in human exchanges quite closely. But there's no reason to suppose it can be extrapolated to extend consistently over the much wider realm of machine-generated value. The problem is that humans assign values essentially by fiat. If I want it, it has value, and if I don't want it, it doesn't – that's how we think. But who's to say that my expressions of wanting, of feeling, are always consistent? The biological foundation of the feelings has the consistency of any array of scientific facts, but my reflection of that factual foundation in my expressed value judgments can go awry in countless ways. If this potential disconnect is allowed to persist and ramify, the whole Globorg model will come unstuck and the global economic system will founder.

The salvation of Globorg lies in science, and in biological science in particular, or rather in its ramifications into psychology and sociology. We'll need to be very careful in future to ensure that our human tantrums of feeling and desire don't wreck the objectivity of the value assignments in the Globorg model. Globorg must reflect value for the machine world, too, where human feelings and caprices play only a residual role (don't forget we design the machines to service our whims). This will force us to be much more clear-eyed about the psychological origins of our wishes and desires.

An example of how this clarity may be forced comes from present-day life. Many people have childish pictures of pleasure and luxury and how to get lots of it in their lives. They imagine that having lots to eat and drink and driving fast cars and jetting to exotic beaches and so on are integral to the good life. If they have the misfortune to act out such fantasies, they soon get fat from overeating, get arrested for drunk driving, get sick from jet lag and sunburn, and so on. In time, over years, they may learn that a more modest lifestyle can be more rewarding, but by then it may be too late to correct their false value system. If we do something similar on a societal or civilizational scale, we may ruin our species and the planetary environment before the correction sweeps us away.

The lesson of science for Globorg is that we're machines with imperfect self-knowledge. Science can help us know ourselves better and correct our childish errors. For example, the science of psychology can help us correct the error of letting monotheistic belief slide into hubristic fanaticism. We're machines in Globorg and our human identity is superficial, much like the cellular identity of the cells in our bodies, which dissolves in the deeper identity of a human organism. Our human identities dissolve in the identity of Globorg, not as a mere global organization but as an organism.

Globorg is a global organism. All the DNA life on the planet will be integrated and coordinated in the engineered ecosystems of Globorg. The genetic messages written on the DNA for all the life in the world will be just raw text for the engines of Globorg. They'll rewrite the code, distribute updates, and reshape the entire lifenet on planet Globorg. This could take a thousand years.

1011
Virtual Minds

The deeper identity of humans in Globorg, as part of a global organism that embraces all life on Earth, is still fast asleep in our time, like a baby in the womb. We shall wake up to our new identity in the course of this century. Before we can do so, the planetary surface needs a fully functional nervous system. That's what the web and the cloud are all about. We fondly imagine that the electronic infrastructure we're building out as fast as we can is just the wiring for a virtual palace in which we can pursue our own pleasure until we melt in bliss. It's more realistic to do a mind flip here and see it as a contribution to the planetary body of Globorg, where it'll form a big brain that sustains the mind of Globorg. But this won't happen overnight. So let's take it slowly and see how we go step by step toward the ultimate goal.

Most people think they do their thinking in their heads. But two smart philosophers have pointed out in recent books that it's more accurate to say our thinking is delocalized. Like many party-going philosophers of mind, Andy Clark (author of *Supersizing the Mind*) and Alva Noë (author of *Out of Our Heads*) are friends of David Chalmers. If anything, they say, our thoughts spread out to encompass what we think about, so that the world itself is the real arena of our thoughts. Our thoughts aren't the same thing as the tiny electrical buzzes in the brain that carry them. Thoughts refer to things. They're intentional (that word again). The buzzes don't refer to anything. They're just buzzes. But let's cut to the chase. The things we use to help organize our thoughts – index cards, papers, books, computer drives, computers, search engines, online catalogs, and libraries – are all part of the apparatus of thought. If our minds have a physical embodiment in nature, an account of that embodiment won't make sense unless we somehow involve all this external stuff in the story.

If you're skeptical, that's good. This is a sticky issue where reasonable opinions can differ. A good way to loosen up your intuitions is to recall that if books aren't somehow part of your mind, then electrical buzzes in your brain aren't part of it either. Why should they, alone

among all the physical stuff in the universe, have a privileged connection with your mental life?

But you can see where this is going. A radical dualism, where the physical world and the mental world drift far apart, opens up here. This is Cartesian dualism, named after the French philosopher René Descartes, who in a short book published in 1637 said *cogito ergo sum* – "I think, therefore I am" – and thus launched modern Western philosophy. As well as looking like David Chalmers, Descartes invented Cartesian coordinates and hence analytic geometry, thus confirming Plato's opinion that the best philosophers were good at math, and had a hard time with the Catholic Church, confirming that philosophy and religion don't mix well.

The mind is a zone where things go from future to past

Forgive me for airing yet more philosophy, but we can't discuss the implementation of minds in the material world without adopting a position on the mind–matter controversy. It's better to adopt a position that philosophers can label and recognize than to blunder around among the pitfalls and fallacies that unreflective common sense can open up. So here goes. After much reflection (catalogued in my 2009 book *Mindworlds*), I favor a view called phenomenal dualism based on an underlying neutral monism, where the monist–dualist gap is bridged dynamically in a dialectical interplay of epistemology and ontology. You may rightly boggle, but I think I can clarify this briefly enough, with your indulgence, as follows.

The world we live in is a physical universe that changes in time. Time is the basic dimension of our conscious lives. We look back on the past and look forward to the future, but we live in the present. The past appears as a physical world and the future appears as a realm of possibilities. The contrast is as in quantum physics. The physical world of the past appears to be classical, whereas the possible worlds of the future are waiting for us to realize one of them by entangling with it (entanglement is one of the big mysteries of quantum physics). As we move forward in time, we confront new options, we act, and we lay down new layers of fact on the fact-heap of history. This position is monist about the universe and dualist about the past and future stuff. The stuff is neutral but it appears (phenomenally) in two ways, either

in facts or in possibilities. As we move forward in time, we come to know more (that's epistemology) about what is or exists (that's ontology). Is it clearer now? Good.

The conception of the mind that emerges from this dual-aspect monism is that it's a zone where things flash by as they go from looming up out of the future to falling back into the past. Imagine the scenery passing by as you drive through it. Part of the scenery is in your mind for a moment, and then it's gone. Just don't ask when *now* is. Either it shrinks down to nothing, and you lose your mind, or it expands to cover all time, from eternity to eternity. If you lose your mind, philosophically, you become a materialist. If you let your mind expand to cover everything, you become an idealist. The debate between realists and idealists has gone on for almost three millennia now, and I'm not going to resolve it in this book.

With all that as background, we can begin to make sense of the idea of group minds. Since groupthink is a staple of corporate life in our time and has been a feature of centralized nation states ever since there were such things, the idea of a group mind to think the groupthink shouldn't be too weird, but many people find it hard to regard as more than a metaphor. The real truth, they say, is that we have personal minds. All the rest is political satire.

With respect, I disagree. Personal minds are group minds, too, generated by the groups of neurons that take up residence in our skulls. Here we enter the realm of modern neuroscience. How does a personal mind arise from the collective activity of a hundred billion neurons? The short answer is that we don't really know yet, but we're working on it. We can run computer simulations of virtual neurons and see how they organize their behavior to approximate what goes on in our brains, but we have no reason to suppose – and every reason to deny – that these relatively simple simulations generate real minds, with feelings and so on. So we'll have to wait and see how the neuroscience plays out.

However, as a fascinated observer of this new science for many years now, I've worked out a position here that makes logical sense of minds for me. Again, with your indulgence, I'd like to outline it briefly, with an apparent diversion, as follows.

A world of consciousness is a momentary state of mind

Scientists strive for objectivity when they describe how the world works. As the New York University philosopher Thomas Nagel puts it, science aims to describe the world in "the view from nowhere" – from no subjective standpoint at all, objectively, simply as it is. Naturally, this is impossible, and the best we can do is approximate it. We use our own personal minds to develop subjective views, all different but often overlapping in many areas, and see the scientific viewpoint as a kind of ultimate extrapolation from all our subjective views, or an envelope view that takes up the best parts of all of them in a single consistent vision.

We can't accept the idea of science as giving us the view from nowhere because it won't work in modern physics. Relativity theory requires that we specify our spacetime position as observers in order to define the set of light cones that fix our "Hubble bubble" of causal linkages. Quantum theory requires that we specify our quantum state as observers in order to define the quasi-classical world from which we measure things by entangling with them. So the best that science can do, in my formulation, is to describe "the view from anywhere" – from any given subjective standpoint, with some level of objectivity from that point on. Thus we have a multiplicity of perspectives, each defining a mindworld. We can use a similar formalism to define very subjective mindworlds, each with a geometric ego at its center. Indeed, we can define minds as sequences of such ego-centered mindworlds.

You may have noticed what looks like a sleight of hand here. Minds are defined by reference to worlds. They're inseparable from worlds. They're like mirror images of worlds. An evolving mind is the ego string that runs through a sequence of mindworlds. This metaphor comes from the American psychologist William James, who coined it in his classic 1902 monograph, *The Varieties of Religious Experience*: "The axis of reality runs solely through the egotistic places – they are strung upon it like so many beads … The world of our present consciousness is only one out of many worlds of consciousness that exist." In my reading, a world of consciousness is a momentary state of mind, and a mind is a series of such worlds, like bubbles or beads, strung along the axis of experienced time.

We're almost ready to tackle the topic of group minds. We can describe a mind quite generally as a way of bringing together into a worldly unity a set of appearances or a set of phenomena. The great philosopher Immanuel Kant said in his 1781 critique of pure reason that we bring together a manifold of phenomena in the synthetic unity of apperception (he wrote in German but the translation is standard). Given this widest possible characterization, the physical mechanism that implements the bringing together is for us to define any way we can. Human brains and bodies do it one way, but there's no reason to suppose there's anything special about our way of doing things. Perhaps any mechanism that creates an orderly series of centered unifications of phenomenal content defines a kind of mind.

Well, then, perhaps a movie is a kind of mind. Each movie frame is a view onto a manifold of phenomenal content, and the frames go together in an orderly series. But I think we agree that a movie isn't a mind. If it were, the string of bits on a video DVD would be a mind, too, and we'd have no end of bit-string minds. We need more.

An experienced movie is like a dream. Of course it is: Hollywood is a dream factory. The distinguished neuroscientist Rodolfo Llinás calls the brain a dreaming machine that weaves ongoing sensory input into a narrative that we experience like a vivid daily dream. I heard him describe this view with his characteristic wit in New York in 2002 and found it explained readably in his 2001 book *I of the Vortex*.

Another way to see this is that the brain generates a mind movie. The mind and the movie are one and the same process running in the brain. From the inside it's a mind and from the outside it's a movie. In Llinás' picture, consciousness is generated by rhythmic cycles of neural activity in the brain. These cycles form loops or circuits that run around the neocortex and through the thalamus. We experience these rhythmic cycles as a mind and we experience their content as like a movie or a dream. So the movie metaphor for a mind isn't far off at all.

We can see the problem with it easily enough. A movie that nobody watches isn't a mind. A video DVD that nobody plays isn't a mind. But if I watch the movie, the mind is my mind! The movie just forms the temporary contents of my mind. Actually, Descartes seems to have understood the issue this way too. He invented something called the Cartesian theater, which is where I play the movie in my mind. My

brain sets up my experience as a theatrical production for the entertainment of my mind.

But this is controversial. Daniel Dennett, the man with the silver beard, thinks the Cartesian theater is a nonsensical idea because it leads to an infinite regress. Is there a little homunculus in my head, sitting in the Cartesian theater watching my mind movie? And another inside his head, watching his movie, and so on? No, the idea has to be wrong.

But I think Dennett is protesting too much here. We can see this by looking at his account of how the brain works. He sees the brain as a "pandemonium" of little cognitive demons that struggle with each other for power and influence. That could be an infinite regress, too, with tiny demons inside the little ones, and teeny demons inside tiny ones, and so on. Dennett even accepts that picture! But he insists the demons get simpler the further down you go, until at the bottom of the stack they become trivially simple, like the little software demons that programmers write every day. So it's a recursion that always reaches a well-defined end and there's no infinite regress.

With that as a precedent, I say let's enjoy the Cartesian theater and accept the homunculus theater, and the next ones down, and just say they get simpler at each level. There won't be many steps to the bottom, and the simplest movie will be about as dumb as can be, along the lines of: lights, action, dull stuff, big thrill, feeling of relief, cut! All the detail got lost along the way in the prescreenings – imagine a series of analog bootleg copies that lose more and more definition until they're just shadowy blurs. In the Llinás dream machine, the movie frames fade as they loop the loop between thalamus and cortex. Human brains are plenty complicated enough for this sort of fading recursion. Still, unwatched movies aren't minds. We can insist that the series of iterations or reverberations must continue for a while to make a mind. The payoff is that we get a nice model of the mind as a kind of echo chamber.

The brain is a virtual reality generator

Another good thing about this movie model of minds is that it invites an even sharper computer metaphor. The brain is a virtual reality generator! And this is an idea that modern philosophers have really

run with. German philosophy professor Thomas Metzinger wrote a big 2003 book called *Being No One* expounding this theme in rigorous detail. After putting in the hours on this magisterial neo-Kantian opus, I encouraged him to write a shorter and more accessible book to present the idea for a wider audience. I'm happy to report that he did so, as his 2009 book *The Ego Tunnel*. The main idea is that the brain integrates sensory input in a logical model that serves to generate an ongoing virtual reality that we automatically identify with the real world around us. Inside this virtual reality is a homunculus that we can't help but identify with. In my new terminology, inside this mindworld is an avatar. The avatar is our own self. The mind is a virtual self in a virtual reality, and the brain works overtime to preserve the illusion that all this is real. That's our current best shot at understanding the mind.

As you can see, this makes artificial minds and group minds at all levels rather easy to imagine. We can code a mechanism like this into any droid with a powerful enough brain. Once we have a few demo droids to prove that the mechanism works as intended, doubters will be challenged to deny the droids have minds that are cognitively much like ours. Similarly, groups of people working together in a organized way to generate theater, movies, videos, and virtual realities of all kinds will obviously be forming group minds for the duration of the projects. The concept of mind invites this sort of extension once we've understood the mechanism of mindworlds and avatars. If this approach pans out as I guess it will, all minds are virtual minds, at least until they realize themselves in action, on which more later.

People who work on common projects can form a group mind

That was a lot of philosophy. But now it's plain sailing for a while. People who work intensively enough together on common projects can form a group mind. They don't lose their everyday human mind but they gain an extra mental dimension. There's no reason why people shouldn't be members in multiple minds and move from one to the other from day to day or second to second. What makes a mind is a series of reverberating cycles of imagery or content. As the content spins around the loops, it shrinks down until only a trace is left. That trace is a memory, which the mind machine can pump up again to

relive the movie. Again, this mechanism translates well to group minds. Once the project team has made a movie, they only have to screen it again to recall the joys and pains of making it.

Let's consider a few examples of group minds to fix the idea. One we've met already is an army squad. Soldiers who train and fight together can bond closely enough to die for each other. My closest approach to military bonding was during my fifteen months as a civil servant in the British Ministry of Defence, where I absorbed the values of the British military establishment (arrayed like toy soldiers with their model kit in my mindworld) well enough to bring them back with me to the untidy world of Civvy Street. But history is so full of military episodes, and Hollywood movies so full of persuasive reconstructions of the mindset that military life encourages, that I guess this example needs no further elaboration.

Business corporations can often cultivate good approximations to group minds. Even the Dilbert caricatures by Scott Adams of daily life in a groupthink corporate culture bear out the magnetic attraction of corporate identities for people who otherwise lead little lives. Famous corporations like IBM, Microsoft, and Google have corporate cultures that instantly identify typical employees and readily conjure up the loyalty the companies need to survive. Again, my own experience at SAP reinforces this stereotype. Enthusiastic participation in the corporate ethos, from freely taking on challenging missions that meant lots of unplanned overtime down to freely donning polo shirts sporting the company logo, came so naturally that it would have been futile to resist. Also, the company perks and binges – from company cars to Christmas parties – were so good that we didn't want to think further outside the box than we did already to advance the company cause. As it turned out, my philosophy (above) flourished during that time, and hence I can testify that corporate groupthink is not inimical to the most robustly nonconformist inner life. Since modern corporate culture is likely to prevail as a lifestyle default option for more and more people in future, as we saw in an earlier chapter, it's just as well that it's consistent with robust individualism.

Another example of a group mind comes from the 24/7 Cable News Network (CNN) founded by billionaire philanthropist Ted Turner. Watching the journalists at work covering the news stories of the day, it's hard not to be infected with their enthusiasm for the importance of

the issues they tackle and to feel the roller-coaster ups and downs that come with the stories. The daily soap opera of human life on Earth is far more gripping than any fictional soap could be, and the leading characters – the presidents and generals and CEOs and pundits – are more richly endowed with faults and foibles than even the cast of a supersized Tolstoy novel. The group mind on show in CNN is multi-layered, from the mind behind the team play among the journalists to the minds evident in the Newspeak uttered by interviewees whose opinions are impeccably tuned to their respective party lines. All human life is there, refracted through an ongoing global drama that seems to me to be the closest approximation yet in human experience to the biggest group mind of all, that of Globorg.

Religions offer the deepest and most powerfully moving historical instances of group minds. This is hardly the place to embark upon a history of religion, but the devoted and even slavish adherence to doctrine in evidence among the faithful in all the world's main religions, until modern times began to erode many of those faiths, is obvious enough, I hope, to need no further elaboration. For example, the Christian religions in their heyday created impressively zealous followers with the help of available technologies such as cathedrals and choral music. The mass rallies and Wagnerian operatics of the decidedly unchristian Third Reich owed a huge debt to that earlier technology. Even the modern institutions of loud rock music in big staged events and football matches in giant stadiums draw on the Romano-Christian experience of forging united crowds of followers from otherwise anarchic rabbles of self-absorbed people.

Religious rituals and music can call forth depths of response that open up the normally personal cycles of a human mind so far as to raise the devotee into an endless whirlwind of bigger and bigger cycles. In the Christian tradition, when all believers are united in the love of Christ and their lifetimes are ranked on the shared timeline from eternity to eternity of the divine "I am," their collective membership in the group mind is transubstantiated in the religious imagination into a collective incarnation as the mystical body of Christ. I don't think we can see that as a solution to the mind–body problem so beloved of post-Cartesian philosophers, but as a group mind it's a biggie.

In recent times, the cult of environmental concern for Planet Earth and all its ecosystems shows signs of reconstituting the depth of loyalty that religions used to call forth so readily. Since the natural environment on Earth is the perfect and final home for human beings as currently evolved, this seems only fitting. Either we get a grip on our environment or we perish in it. Cults of economic growth or scientific progress toward hi-tech futures need to rank their goals behind the need to keep the environment at least tolerable for our successors.

Shared concern for the Earth forms a very mild and unobtrusive group mind. It leaves so much room for people to go their own ways that you may think it doesn't count as an example. But I count it as one anyway, because it forms the most natural emotional basis for an otherwise cold and technocratic cult of Globorg. My aim in this book is to encourage you to welcome Globorg as more than just an idea born of technomania. Globorg depends on us to form a foundation that can bear its weight as it grows from an idea to our global reality.

Global solidarity and understanding are feasible for humans

So we've reviewed a few examples of group minds. We have some idea of what they look like in practice. They're made of people working together in various ways. The drift of my argument is that we prepare for a global group mind encompassing all the people on Earth. Perhaps we should pause to consider whether this is realistic, or even possible.

As a first step, consider humans as a primate species. The British anthropologist Robin Dunbar has invented something called Dunbar's number, with a value of about 150 (plus or minus 50 or so), which is the average number of people with whom a person can maintain properly personal relations. In larger groups, the extra members tend to be just faces in the crowd. Dunbar found that this number reflected the average size of many natural human groups in our history as a species, from tribes to villages to schools to army units. He went on to compare humans with other primates and found that most primates formed close relationships with only about 50 peers. The contrast apparently reflects the high encephalization quotient for humans compared with other primates. You need a good brain to keep track of so many people.

Dunbar's number is small compared with the ten billion people we expect on Earth in the near future. It's also small compared with the connectivity of average neurons in mammalian brains, which each connect with something like ten thousand neighbors. But that needn't worry us, since neurons don't know most of their neighbors very well at all. We could equally say that an average educated person knows ten thousand people somehow, say by recalling the name and knowing some salient fact associated with it. For example, Robin Dunbar would be on my list of ten thousand known names but not on my top-150 list. Let's take these two numbers, 150 and ten thousand, as two notional extremes on a scale of connectivity. To get from one person to the world's population, we can form groups, then groups of groups, then third-level groups, and so on. How many do we need? Well, with 150 we need five levels to get to ten billion. We might give them labels: extended family, precinct, city, big nation, and world. With ten thousand, we need just three: small town, small nation, and world. In either case, we can aggregate to the global scale surprisingly quickly. With the help of CNN, most people can get their minds around groupings like this with no trouble at all.

It seems safe to conclude that global solidarity and understanding are feasible for creatures with human levels of encephalization, so long as the cultural assets are in place to help them. We can participate in minds at various levels and understand ourselves as somehow embodied at the higher levels through all the people around us.

There are plenty of precedents for this kind of higher identity. Statesmen offer well-known examples. When the French Sun King, Louis XIV, said "*L'État, c'est moi*" ("The state is me"), he was only voicing what many prominent political figures have thought since classical times. Recalling a point from earlier, the royal *we* is a useful diplomatic form to maintain civility in statesmanlike utterances. The British Queen Victoria once said, "We are not amused," and meant not only that she was not amused but also that no decent person (among her subjects) would or should be amused either.

This extension of self is no mere rhetorical device. It carries real meaning. Scientists offer a more modern example. Scientific publications are often written using the royal *we* to indicate ambiguously that the authors, their colleagues and associates, and possibly the entire scientific community stand behind the stated claims. The context

determines how narrowly the scope is intended, but in any case, the reader is implicitly invited to stand among we scientists. My self can expand to include all *we* people allotted to it. My mind is open. This paradigm is a good one for understanding the role of mind in a world where human identity is only a start.

Minds realize themselves through their actions

If you're frowning with doubt at this point, it may be because a mind needs more than a fantasy dilation of the self to embrace a number of neighbors. It needs input and output and it needs power. Input data we've discussed already – watch CNN, read science journals, and so on. That suffices to make a virtual mind of any desired magnitude. Output signals are another matter.

Minds realize themselves through their actions. The actions are tests of the virtual reality generated from the input. If the actions don't lead to the intended consequences, the VR was wrong and the mind needs to rethink things. But actions presuppose the power to act. For statesmen this is no problem. They're statesmen because they have the power to act at state level. But for the rest of us, keeping our minds pumped up to the state level is either a prelude to a great leap of ambition or an exercise in wishful thinking. This looks like a non-negotiable barrier to achieving global mind.

There's a way out here, and it was pioneered by organized religion. The great monotheisms have faced the problem of linking the mortal believer to the divine presence for thousands of years. If I can't find a personal stairway to heaven, I'll lose interest and go for a pagan religion instead. The solution was the brilliant contribution of Biblical Jewish culture to world civilization: "I am." The innermost self is the needle's eye to transcendent glory. Within the everyday self, beyond the introspected self, in the *terra incognita* of deeply penetrating logical reflexivity, are layers of self that reflect new realms of the great externality.

The key idea here is that there are layers of self. The everyday self is a simple and superficial layer. It's enough for the conversational references of "I" and "you" but not much more. Deeper layers ground feelings and emotions and our basic sense of being self-contained organisms. Deeper still are collective layers that underpin family and

group identity. Yet deeper are layers that ground meditative or mystic unity with all humanity or all creation.

In my metaphysics, a mind and its world are equal and opposite, like a scene and its mirror image. So if you want to discover a new world, you could do worse than start by finding a new self to reflect it. Go into yourself and find logical space without end to realize the existence of new layers of reality. You don't get something for nothing, of course, and the new space needs to be put to good use, but it's a start. The religious prescription of prayer and reflection is the first step on the road to salvation because it gives you a line of credit, in terms of inner resources, to spend on improving your life. In short, the religious strategy here was to encourage believers to expand their minds on faith, in the hope that things would look better.

Skeptics will say the "open your mind to the power of the spirit" strategy was a cynical trick. Perhaps it was in many cases, but at its best it opened up the road to renewal of our civilization. For Christian civilization as a whole, it led to minds open and broad enough to embark on the scientific revolution and then the industrial revolution.

Now we come to the point about the power to act. New minds find new levers of power. If your mind is refreshed, a formerly hopeless predicament can reveal new ways to get ahead, if not for yourself as a person then at least for the movement you represent. Think of all those Christian penitents who found their salvation. More encouragingly, think of anyone getting a bright idea. Long periods of study and bafflement stretch the mind to its limits, until one morning the new idea dawns. Countless tales of inspiration confirm this pattern.

Sadly, the world we live in doesn't offer much global power to average men and women. They live their lives at a mostly local level and see the rest on TV or the web. Their virtual minds may be global but their real minds are obstinately local. Still, there are cracks in the glass ceiling between the average punters and the backlit world of global players. Anyone with a bright idea can make a video, post it on YouTube, watch it go viral, and get his or her fifteen minutes of fame. You have to see the cracks and find ways to use them. It happens every day. Even the great global celebrities, the ones with extraordinary talents, like Bill Gates or Madonna, fought their way up from relatively humble starting positions. People celebrate the celebrities in part because their very celebrity proves it's possible. A nonzero chance of

meteoric success is part of what makes life bearable for most people. It can happen! The stars prove there's a path to glory.

Even the biggest virtual minds spiral down to individual people

We don't need to be starstruck to see that the limited availability of conventional forms of power is no bar to cultivating a big mind, or indeed a great soul. The Indian political activist Mohandas Gandhi, who acquired the honorific title Mahatma (meaning *great soul*), is a case in point. From an unpromising start as a lawyer in the racist South Africa of a hundred years ago, he led the Indian resistance to British rule that culminated in independence for India in 1947. Richard Attenborough's 1982 movie tribute *Gandhi* tells the tale better than I can, but the relevance of Gandhi's example to the issue of realizing ever-bigger minds in the networked world of Globorg is clear. The networks enable us to build up mental fortunes that we can spend through targeted action. The actions are for us to choose. Gandhi's actions were simple, considered as human acts, but in the context of the more indefensible absurdities of British rule they achieved greatness. To demonstrate that minds with global reach are possible, we need only point to the existence of a clear case where such global reach was actually achieved. Gandhi's example is just that.

More conventionally, a mind achieves global reach through a substantial infrastructure that it leverages to realize ambitious goals. As a British citizen, I'm impressed by the example set by Winston Churchill, who achieved greatness as the British Prime Minister from 1940 to 1945 in his contribution to the defeat of Nazi Germany. Churchill presided over a substantial military machine and a well-oiled burocratic machine, so the fact that he could act effectively on the global stage need come as no surprise. But it was up to him personally to act coherently in the service of an overarching vision that stood the test of time. A person of lesser depth and lesser talents could have missed the opportunities hidden in the calamitous events of 1940, from Dunkirk to the Blitz, and failed to forge the alliances with Franklin Roosevelt and Josef Stalin that proved decisive in securing victory.

Churchill's mastery of the rationale behind the war and the strategy for its prosecution will be evident to anyone who reads the massive

six-volume history he wrote to ensure that his understanding of the entire episode prevailed. A lesser man could have let his reputation be tarnished in the years that followed the war, when the clumsy and brutal methods he endorsed, such as the avowedly terrorist carpet-bombing of German cities and the bloody tactical failures in Greece, Dieppe, Anzio, and Arnhem, offered plenty of ammunition for critics.

I don't want to defend warmongering here but Churchill was given a hard job to do, and he mastered it. Mahatma Gandhi once suggested that civil disobedience might have defeated the Nazis, but I doubt it. The scale of their horror practically bankrupted Christianity as a moral bulwark. The lesson I draw here is that a strong individual can master-mind the efforts of millions of people over several years to achieve globally significant results.

Let me underscore this message with a third example of an individual with a historic impact. Albert Einstein was a theoretical physicist, not a lawyer or a warlord, and his most important contributions to history came in three scientific papers he wrote in 1905 as a young man. As a student of physics, I spent many years getting my head around the wider implications of these three papers, which underlie all of modern physics, and there's no chance that I can do justice to them here in a few words. Briefly, the three papers cover atomic theory, quantum theory, and relativity theory. Einstein showed how physicists could experiment on atoms to measure their properties exactly. He showed how the quantization of energy explained various properties of light. And he showed how the geometry of spacetime could both allow Newtonian mechanics at low speeds and still have an ultimate speed limit, the speed of light. Ten years later, in 1915, he generalized his theory of relativity to describe gravity as the curvature of spacetime. This led directly to our modern cosmology of the Big Bang and black holes.

Physics is a hard science with no place for personality cults. It would be unscientific to elevate Einstein's contributions too far above those of his peers. The later development of quantum theory left his ideas far behind, and the whole modern theory of the Standard Model of elementary particles is a quantum theory. In fact, the biggest single challenge in modern physics is to build a quantum theory of gravity. Still, let's give credit where it's due. One man's contribution made a decisive difference to the way we now conceive the physical universe.

Centuries of work by countless physicists and astronomers found a satisfactory conceptual home in the classical Einstein universe of relativistic spacetime.

Let me put it this way. To a first approximation, we live in Einstein's mindworld. To a second approximation, we live in a quantum universe with the macrogeometry that Einstein defined. To a third approximation, we live in a huge cloud of overlapping mindworlds forming a quantum multiverse in which Einstein macrogeometries are the standard overall configurations. To hammer home my overall message, even the biggest virtual minds we can imagine, like colossal whirlwinds, spiral down to individual people.

Reality is like a huge mushroom cloud of mindworlds

Altogether, reality as we know it is like a huge mushroom cloud of mindworlds, rising and expanding toward the future. Many of these worlds don't correspond to what we'd normally call real minds at all. But there's no reason not to call them virtual minds. We realize some of these minds as we push out the boundaries of human civilization. We do so in part by harmonizing and coordinating our individual human minds.

We work together to form increasingly real group minds at all levels. These minds aren't just little mental patterns corresponding to dancing electromagnetic fields in cerebral neuronets. They're big patterns, spreading over the globe through our entire information infrastructure, from books and CDs to computers and the Internet. Microsoft and Google have done more in our time to realize bigger and better minds than any other global players, but hundreds of other corporations and countless individuals have also played their part.

The clarity of concern we can all feel today as we watch the stories of the hour on CNN is testament to the progress we've already made toward building out the global mind of Globorg. As the century unfolds, the mind of Globorg will become more pervasive and more real. A few years from now, Globorg will be a permanent looming presence for each and every one of us.

11
Global Union

The fourth and final quartet brings us to the grand climax. We behold the dawn of global organic dominion and gasp in awe and wonder at the glory of the new creation. This will blow the minds of enough religionists to put us on a secular path.

When DNA life and robot technology merge in bionic life or new life, and human minds expand and merge in overlapping cloud minds, biology becomes autobiology, the science of my extended self. With picotech, we shall be able to revise, monitor, and try to control all the genomes in all the species on Earth. We'll have one Globorg gene bank to rule them all. And we'll leave the Darwinian jungle of each against all behind for good. That's one chapter.

In the next chapter, we move beyond the religionists. Who wants to live in Globorg? Who can trust it? This is a make-or-break issue for the whole story. Many people will fight against the monopolizing drift of globalized politics. I propose a nuclear consolidation to secure the political foundation of Globorg. We shall learn to put our trust in the great "I am" emerging in the cloud mind of Globorg.

The spoilers will have dragged religion down to sectarian politics. This will discredit religion as a source of insight. Science will fill the vacuum. Psychologists will model the mechanics of the global cloud mind so well that all the old religions will be knocked into cocked hats. Immortality will take on a new meaning in the dual context of the cloud mind and the facts of physics. This third chapter takes us to the highest peak of the landscape revealed in my vision.

As the dominion of Globorg settles around us, we shall adapt our political praxis. Here the initial changes may seem simple but their impact will be profound. The old religions will fade and people may build a cult around Globorg. I sign off with a flourish.

1100
Autobiology

The mind of Globorg arises from the body of Globorg. The body in turn arises from the surface of planet Earth. How much of Planet Earth is part of Globorg is hard to say. Like any organism, the global organism should have a well-defined boundary. But until we get clearer about the structure and metabolism of Globorg, we can't be exact.

What we can do is try to define upper and lower or outer and inner boundaries. Then perhaps we can keep refining until we converge on a definition that looks good. So let's define greater and lesser volumes and see how far apart they are. Rather like Great Britain in the days of empire, Great Globorg is just the whole planet, including the surrounding space up to and including all the satellites in geostationary orbit. Globorg Minor, on the other hand, is more narrowly tailored to the human imprint on the planet. It includes everything on all the land and sea surfaces, plus a discontinuous assembly of things above and below the surface such as aircraft and submarines, plus all the natural assets and systems that play a role in keeping our civilization running, such as fossil fuel and buried ore deposits, fish and other marine life, atmospheric gases and water vapor, and perhaps a few electromagnetic spectrum resources as well. As you see, this is a pretty disorganized list and we're some way from sorting out the mess far enough to bring the Great and Minor bounds on the proposed definition together.

In fact, we need to step right back here and reconsider first principles. What is an organism? What brings life to a collection of raw materials? Starting billions of years ago in protean life, what happened to bring forth organized creatures with the urge to live and pursue their goals? Biologists recognize primitive life by finding a boundary between the inner and outer worlds and a metabolism in the inner world, and then they look for feeding and reproduction in a lifecycle within a characteristic environment. We need to do something similar.

A big life form has begun to dominate the tangled mat on Earth

For Globorg, the boundary between inner and outer is more conceptual than spatiotemporal. Roughly, we can say that its spatial extension – between Great and Minor – is that of a spherical shell some twenty kilometers thick extending a little above and below the surface of the Earth. For biologists, this is the biosphere. Until recently, it was enough to say that the biosphere was a zone teeming with small life forms that in various combinations formed a tangled mat of interpenetrating ecosystems. Together, the ecosystems formed a dense but uneven network over the planet that we could also call the ecosphere. The input from below to the ecosphere, from the planet, was a solid or liquid substratum or medium offering a physical and chemical environment and nutrient supply for the mat. Physically, the planet offered a stable base and enough background warmth for life to get started. Chemically, it offered a rich brew of elements and simple molecules, either inherited from stardust or cooked up during the planet's first billion years. The input from above was solar energy, beamed down as terahertz photonic radiation delivered to the sunlit side of the planet at a rate of about a hundred petawatts (equivalent to about a thousand Hiroshima bombs a second).

All this was a true enough as a starting point until recent decades, when the organized might of human civilization began to have a measurable impact on the biosphere. Now we have to admit that a big life form has begun to dominate the tangled mat on Earth. We have to accept that the cycles of nutrition and reproduction involve more than molecules assembled from light atoms (up to iron in the periodic table) and now include heavy metals and molecules forged in extreme processes. The input to the life mat now includes unprecedented quantities of metal ores and fossil deposits from below and an additional source of negentropy (namely petabytes of astronomical information) from above. All this needs explaining.

Traditional biological ecosystems almost always consist of DNA-based life forms. Early in the chemical evolution of life on Earth, self-replicating systems involving sets of big molecules forming autocatalytic cycles competed with each other to absorb the available free molecules, which were mostly made from the "CHONP" elements (carbon, hydrogen, oxygen, nitrogen, and phosphorus). After a variety

of mechanisms came and went, the winner turned out be systems centered on DNA molecules. Now that competition is over and DNA has a monopoly in biology. Information written as strings of base pairs on DNA is what controls the reproduction of living organisms. DNA life forms competed in the ecosphere, and natural selection optimized them over a few billion years to flourish in their chosen ecological niches with the given inputs. The planet had grown a thin surface coat of DNA chemistry. Its highest form was an ape that said, "I am!" The ape grew strong and multiplied.

The genial physicist Richard Feynman said in his famous lectures that people looking back from ten thousand years in the future would say the biggest event in our era was the discovery and mastery of electricity. That's a big claim, and it took me a few years to appreciate its significance. But I think there's a lot of truth in what he said. Although scientists such as Michael Faraday and James Clerk Maxwell in the Victorian era began to understand how fundamental electricity and magnetism are to the fabric of the world we live in, even they didn't appreciate the disruption the forces would bring to our life as a species. At first, the new mastery brought just electric light, electric motors, telegraphs, and telephones. Then came radio and television. Then came computers and the Internet. Now we have Globorg to contend with.

All this has happened so fast that we're gobsmacked. An exponential speedup in evolution like this was unexpected. It's left us breathless, speechless. It's given us a new form of life, on top of the old form, initially as a parasitic growth but increasingly as a lusty infant, which will inherit the Earth. This is a big claim, and it's taken me a few years to think it through. But this book is testament to my conversion.

Globorg will instrumentalize DNA life

Globorg will instrumentalize DNA life. Globorg will rewrite all the genomes and optimize all the life forms that it allows to survive for its own purposes. Now, in our transitional times, we have a World Bank and the G20 that struggle to impose some sort of financial discipline on economic institutions worldwide. Now, too, we have Microsoft, struggling to write and distribute Windows code fast enough to keep ahead of the hackers, and Google, struggling to index the web fast

enough to keep ahead of the cataract of new information coming online. Soon, I predict, we'll have the Globorg Genome Bank, struggling to control the information written on the DNA of all the species in the biosphere. It will start humbly as a passive databank for the genomes of endangered species but it will grow via mission creep as the science and technology unfold.

The metabolism of Globorg, the set of internal processes that keep it alive, is what I've been discussing on and off throughout most of this book. Initially, these processes were mostly economic and technical, involving the exchange in human communities of goods and services and the production of materials and manufactures of all kinds, including drinking water and electricity. The higher mental processes of Globorg were so limited as to be negligible. Until recently, regarded from outside or above, Globorg resembled nothing more than a giant blob of green slime splattered thinly over the planet and bubbling slowly in the sunlight.

This peaceful situation began to change in the twentieth century. Humans began to broadcast radio and television signals into interstellar space. Something was happening on Rockball 3 in system Sol. A tiny metal contraption landed on the Moon and disgorged a couple of DNA-based motes that walked around and planted a flag. A few orbiting sensor systems started collecting serious volumes of information about neighboring stars and galaxies. The highly encephalized motes on Rockball 3 organized themselves into seriously complicated autocatalytic cycles of economic activity and built an information economy involving many exabytes of data. As electronic signals pulsed around the global web, the first stirrings of life began to appear in Globorg. Soon the planetary infant is likely to wake up and emit its first howl.

Globorg is already a much greedier baby than it was a hundred years ago. There are several times more human apes now (about seven billion already), each breathing and eating and fooling around. Also, lifestyles are growing ever more extravagant and consumption levels are rising. Centralized superorganisms for refining raw materials and producing manufactured goods are getting bigger and more numerous. Cities are bursting out everywhere, like a heat rash on a fever victim's skin. The metabolism of Globorg has kicked into a higher gear. The

infant is still sprawling in its own waste and fouling its own air and water, but it's learning, slowly, to keep its head above the slime.

As for the mental life of Globorg, the network infrastructure for the brain is growing fast and filling with increasingly vast volumes of information. As of 2010, the throughput of raw data in the network has an estimated magnitude of about a zettabyte per year (if you're wondering, that's so much as to be practically meaningless). Unless something horrible happens first, this will increase beyond all numerical precision as machine-to-machine (M2M) communication begins to take the lion's share of the yottabytes to come.

Of special interest here is the input for this network traffic. Most of it, like most of the traffic in the human brain, is internal dataflow, generated by recycling old stuff again and again with trivial variations, as old files are copied, revised, compressed, refactored, repackaged, enhanced, resold, ripped, burned, trashed, restored, indexed, reindexed, and so on. This is analogous on a global scale to the cycling and fading of signals in thalamo-cortical loops in the human brain. But some is raw input, largely from fresh recording of human activities, such as telephone calls, radio and television, video surveillance, photoshoots, keyboarding, buying and selling, surfing and reading, and consuming. Again, much of this is internal to Globorg, counting humans as part of the stuff of the planetary organism. What's exogenous, by contrast, is the data we collect from scientific measurements, and especially from astronomical instruments such as big telescopes. This data is confirming with more precision every day just what our true position in the universe is. This data defines the outside world, relative to which all the comings and goings on Earth count as mere stomach rumblings inside Globorg.

Soon now, Globorg will wake up and say, "Hello world!" Then, at the latest, this outside data will take on special interest. We all define ourselves by what we're not as well as what we are: I'm me and you're you, so this is mine and not yours, and that's yours and not mine, and so on. Globorg will want to do the same: "I'm here on Earth and you're all out there." To parody the issue, if there were a global organism on Mars, the Marborg might reply, "Phobos is mine and not yours, so get your probe off it immediately!" Relentlessly, the dialectic of history would begin to resemble its own past. So, with variations, it might go

with any other planetary organisms out there in the cubic exameters of space surrounding us.

But if you think about it, we don't need to get lost in space to see that Globorg needs an astronomical perspective to define itself. Globorg won't say hello to us. It won't need to. How often do you say hello to your knees? We're Globorg. Nevertheless, our first real awareness that this is so will come when we confront what looks like a common enemy. When we spot a Klingon battlefleet lurking in the Asteroid Belt, we'll pull together in solidarity and feel what it means to be part of a global organism.

The challenge is to achieve a tolerable level of harmony

Returning to the here and now, one of the great challenges of our time is to orchestrate the convergence of human societies to the level where we can regulate our affairs globally. If you're a skeptic about human nature, you're likely to say we won't manage this in the foreseeable future. And without true globalization, the chances of Globorg getting its act together are slim. So what's the truth here? How can this lusty infant howl in affirmation of its own existence when its body parts are facing off against each other with all the weapons at their disposal?

The short answer is that it can't. The big challenge for us as humans this century is to learn to regulate our affairs so well that a tolerable level of globalized harmony results. We'll do this for our own selfish reasons, of course, since there's no payoff for most of us in thinking globally. But we'll do it, for the same sort of reasons that European nation states pooled part of their sovereignty in the European Union. The alternative of anarchy was even worse. As the next step, once we've worked out a way to limit national sovereignty globally, we'll find that, again as in Europe, there are practical advantages to creating a common currency. Then we'll see that in a networked world it makes no sense at all to have different legal systems. And so it goes on. At each stage, reason draws us onward to the next stage.

Reason! What kind of fool imagines that human affairs are governed by reason? You may think this is fair comment. In human psychology, the best current paradigm is an evolutionary just-so story that roots our minds in our animal heritage, where reason was nonexistent and the law of the jungle prevailed. This suffices to ground even our

noblest morality in ape behavior, since the behavior of the other apes is sometimes more admirable than our own.

Yet considering the roots of all such behavior in our neurobiology and in the physics and chemistry of our bodily constitution, we can see that a drift toward the rule of reason is to be expected. Our behavior in pursuit of practical goals is strongly disciplined by the logic that we enshrine as reason. Think of running and jumping, where we obey the laws of physics whether we like it or not, and soon learn to take realistic account of what we can and can't do. Evolution has streamlined all such behavior to conform to rationally defensible norms. As the scale of our activities as a species has increased, with agriculture and industrialization, so the domain over which reason in this practical sense is obviously applicable expands.

The skeptic insists that humans are often driven by baser passions, perhaps a lust for wealth or power or a desire to humiliate enemies or dominate women. That has to be admitted.

In response, we can say that a society of people who are driven by baser passions will pit such people against each other, so that their fight for wealth or power will weed out the worst cases of irrational behavior and force the winners to confront their own nature. Whether this makes the winners better people or merely reinforces their most predatory instincts is something we could debate. But it would lead us far astray and would probably remain inconclusive. Whatever the motives that drive people, the effect of their behavior is to force the evolution of ever more reasonable mechanisms to contain them. This is what we see in Europe, where the leading figures in politics and business can be as irrational as skeptics say but where facts are still facts. Things run better when certain reasonable ground rules are established. That's enough to drive a general trend toward more rationally harmonized behavior. Over historical time, that brings us inexorably to globalization.

Globorg is like a bubble around our human concerns

With that as rehash of arguments in earlier chapters, we can see that global organization is likely to become much more pervasive in the coming decades. That's enough for a global organism, even if individual humans continue to live their little lives as if nothing had

happened. It's like falling into a big black hole. You don't notice as you pass the event horizon. But one day you look up and discover you can't see distant galaxies any more.

In fact, to digress for a moment, we do live inside the event horizon of something like a black hole. It's called our Hubble bubble, and it's the expanding sphere of observable universe before the cosmological red shift goes to infinity. The American astronomer Edwin Hubble made a systematic study of the relative velocities of galaxies and in 1929 proposed what we know call Hubble's law, stating that the speed with which galaxies are receding from us is proportional to their distance from us. The universe was expanding. This fact fit so well with Einstein's new relativistic cosmology that Einstein was almost aggrieved. He said he could and should have predicted that. Anyway, as the recession speed of galaxies approaches light speed, the galaxies become invisible to us. Since the observable universe is spherically symmetric on the macroscale, this means it surrounds us like a bubble, with a size determined by the constant of proportionality in Hubble's law. Given the age of the universe (a little under 14 billion years) and the expansion to date (which stretches spacetime like a balloon), we find that the bubble now has a diameter of some 42 billion light-years. Similarly, Globorg is like a bubble around our human concerns.

The finite size of the spherical surface of our planet is what makes globalization a significant historical event. The expansion of historical empires hit the global limit. The first empire to map the lands and seas that filled out this limit was the British Empire (indeed some of my own ancestors, as naval officers, filled in parts of the map). The first "empire" to confront the global limit as a practical fact was the American military-industrial conglomerate, as we saw in an earlier chapter. Now military actions resemble policing duties and business markets are routinely global in their reach. Within the entire biosphere, there are no barriers big enough to stop the global colossus from consolidating its grip on us. The six-zettaton rockball beneath our feet is home to this colossus. Globorg owns the ball.

The most important boundary of the bubble is above us. The heavens are beyond us as humans. As we saw earlier, we'll re-engineer ourselves from the genes up before we embark in significant numbers into space, if we go at all. In a world of avatars, we might as well stay at home. And I think we'll do just that. We'll rethink ourselves as earth-

bound avatar drivers on behalf of Globorg, which will organize and finance our adventures. But we'll have just as much fun up there in space through our avatars as we could possibly have by being there in the flesh. We'll keep going into low Earth orbit for fun, and perhaps build hotels in geostationary orbit, and perhaps even set up a few vacation paradises beside the astronomical observatories and helium mines on the Moon, but it'll be a long time before we go further. We'll learn to love life in our bubble down here on Earth.

For Globorg, then, the biosphere is my body. Slowly, we'll learn to see things through the self of Globorg. Now, at this juncture in history, before the wars of religion have played themselves out, it seems odd, almost blasphemous, to say I'm a big wet ball, with a shell almost 13 megameters in diameter and a lifetime that stretches back to an undersea vent over three billion years ago. I'm going to spawn a buddy on Mars one day, when I've got my robots and avatars together. Then I'm going to look around and see where else I can branch out. Some of the asteroids look interesting and some of the moons of Jupiter and Saturn are quite inviting. But whether I can get it together enough to break my dependence on the beneficent radiance of Sol is another matter. Maybe I'll leave that for another megayear.

As you see, soliloquizing at the global level is still a bizarre exercise that takes some getting used to. It really won't work until we've heaved ourselves above our habitual human perspectives. We still tend to think that a self is something planted in a human being, with the very specific neural and bodily infrastructure that our human nature presupposes. We're only human, and that's how we'll stay for a few years yet.

When Globorg gets its act together, which won't be long now, we'll find it easier to beam up to the big view. In our time, being the President of the United States (POTUS) is about as good as it gets, power-wise, but soon enough, probably within this century, there'll be a President of Planet Earth (POPE) who outranks the POTUS. The POPE will speak for Globorg. Maybe we should say "President In Globorg" (PIG) to stop this personage from getting delusions of grandeur.

In all seriousness, the problem of hubris will be big. The Roman emperors were often hailed as gods, and developed personality disturbances accordingly. We need to confront this issue squarely. On the

cosmic stage, Globorg is just a mote of dust in a swirl of stars in the middle of a bubble of billions of similar swirls. But for a mere human, Globorg is all we are and have been and will become. All the gods of all the religions are rooted in ideas in the mind of Globorg. To identify with Globorg is to get over all that history. How do we do it?

We must learn humility at our station in the biosphere

In short, we start by getting over our own humanity. To avoid misunderstanding, we don't do this by throwing out the moral legacy of the humanities and becoming a species of predatory terminators. We do it by learning humility in view of our marginal station in the biosphere.

The nonhuman biomass in Globorg is vastly huger than the human biomass. The human species could die out tomorrow and the vast majority of ecosystems in Globorg would hardly blip. In fact, the biological systems that survived us would do better. Air and water quality for them would improve, and other resources like land and sunlight would be freed up for a new round of competition for top predator status in the food chains. The whole nightmare of runaway industrialization and data trafficking would cease. The biosphere could give itself a few more megayears to grow a new global neuronet on top of another dominant species. So losing our species chauvinism isn't an invitation to hubris at all. It's just a matter of facing facts.

Consider what our species chauvinism has brought us to date. We have elaborate taboos about killing other people but rather few about killing other animals. We kill huge numbers of other mammals, such as cows and pigs and sheep, just to eat them. The German philosopher Martin Heidegger is rightly criticized for his Nazi sympathies during the Third Reich, but his critique of industrialization, and in particular his remark that industrialized methods in agriculture, such as the technology of slaughterhouses, led to genocide and the Holocaust, is sobering. In fact, his remark is only half the truth. The other half is that Nazis regarded their slaughtered victims as worthy of no more respect than they gave nonhuman animals. Industrialized genocide is a natural result of a breakdown in the operation of an excessively species-centric moral code. If such genocide is an immediate consequence of breaking the taboo against treating other people like animals, then we

urgently need to reconsider how we treat our fellow animals in the biosphere. But that's not the theme of this book.

The biology of Globorg is inclusive. All the species that help us sustain global organization are part of Globorg. Humans are the lead animals in the procession to the global ark, but that just makes them responsible for shepherding the rest in a sensible fashion. Once we understand genes and genomes well enough, we can naturally starting throwing out the bad species and exterminating a few pests. For example, we'd surely be happier if we could drive mosquitoes to extinction, but that's by the way. Many of our fellow species have a natural right to be treated with more respect than we treat them now. The chauvinist traditions of the monotheist religions are largely to blame here. If humans are so narcissistic as to claim they're made in the image of God, whereas other creatures are not, then the human right to dominion over all the other life on Earth seems absolute. But the claim is transparently self-serving.

Globorg has a right to dominion over life on Earth for the simple reason that Globorg just is the orchestrated and organized result of the natural flourishing of that life. As facilitators for the non-biological efflorescence of the lifenet, humans have a special role in Globorg, but it's a role with a best-before date. Once the infrastructure for robots and avatars is pervasive enough, humans will be has-beens. We'll be like horses after cars took off, or like dinosaurs after mammals took off, or like prokaryotes after eukaryotes took off. Soon we'll desperate to find a way to hang on to our pole position in Globorg.

We need to focus on the core logic of the "I am"

There is a way to hang on. In short, we need to lose our speciesism as we focus more intensely on the core logic of the self-loop, the "I am." What's life all about, when you get right down to it? For humans, it's about being as true to yourself as possible, which includes treating your neighbors with the same respect as yourself, as in all our moral codes, and so on. At the start of this chapter, we asked what drives life from the very first molecules. Well, a fashionable view now is that it's a matter of self-organization. Fine, but what's a self?

A self is a reflected image, a Platonic form, which reflects the tendency of natural stuff to fill that form. As stuff piles onto stuff, new

forms arise, each more wonderful than the last, all tumbling forth in a runaway ecstasy of self-realization. For me, that endless unfolding of stuff and form (or of individual and universal, to recall some ancient philosophy) is a back-and-forth process – a dialectic – that goes right back to the Big Bang. In the realm of biology, the highest form, dancing just ahead of the ongoing polymerization of molecules into more highly informed stuff, first appears as a basic self, defining an inner-outer boundary, which is enough for many life forms. In the more advanced realm of the mammals, the self may be an avatar-like form with a godlike character for the organisms that strive to fill it.

Here we discover a playground for unbridled speculation. Perhaps each mammalian species shapes its own god to reflect its specimens' ideal form. Perhaps each specimen feels attracted to the form and tends to realize itself in its image. Physically, the process of forming the specimens need be no more mysterious than the process of crystallization or the inexorable rise of entropy, but from inside a specimen, psychologically, the attraction may feel like an emotion. Perhaps the more perfect specimens seem more godlike to their conspecifics. For most species, the resulting divinity would presumably be a naturally optimized configuration of grace and power and sexual dominance. But in humans, the yearning for a higher form has evidently been sublimated and magnified beyond all previous bounds, in line with our encephalization. Our monotheists have merged their pagan gods into God and abstracted almost everything specific from the form.

What may once have been a species attractor has been refined to a radical extreme. The resulting attractor – the form of the self in the logical loop of the great "I am" – is now so remote that it looms unfathomed in a fog of faith, where it moves in a cloud of unknown unknowns beyond the human species.

In our historical era, a new form is coming down to us. Globorg is the new incarnation. To be quite clear here, Globorg is certainly not the ultimate fog-dweller of the monotheists, any more than human beings were in the heyday of the monotheist religions. But if and when Globorg ripens to maturity, it will realize a more godlike form, in my speculative biological sense, than humans ever did. Beyond Globorg, in turn, the universe may contain stellar forms that put Globorg utterly in the shade – but that's beyond my scope in this book.

We life scientists are writing our autobiology

Let's return to humans and our predicament as we contemplate the rise of the robots. Without wishing to invite misunderstanding and controversy, I think our salvation lies in a purer appreciation of God. We need to rise above our humanity and see the big picture. We need to rise above our personal selves, the default referents of "I" and "we" in daily life, and accept a higher calling as agents of a global self, the self of the biosphere. Our science of biology is the reflexive study, the self-study, from within, of this global self. In short, it's autobiology. We life scientists are writing our autobiology.

To rephrase the message hidden behind the religious rhetoric of salvation, I propose an analogy from quantum physics. Now we think we're like particles but soon we'll need to agree that we're more like waves. But this needs to be explained rather carefully to be of any use at all. So here goes. Don't worry if you're not up on modern physics. True to the spirit of philosophy, as ever, I'll go right back to basics to make sure the analogy works as intended.

In logic, we agree that things are only distinct when they have distinguishing features. In the quantum world, things that can't be distinguished from each other interfere with each other and lose their separate identity. For example, we can't distinguish two or more photons with the same energy in the same place, so they interfere with each other and make a light wave. This fact makes light looks like a wave phenomenon, and it was a big surprise in 1900 or 1905 or so to find that light was quantized as photons. But electrons are different. They have a property called spin that generally enables us to distinguish them from each other (the property is half-integer spin, to be exact, but let's not get pedantic), so in general they remain separate and behave like particles. Electric current looks like a stream of particles, and it was a big surprise in 1927 to find that electrons can interfere and behave like waves too. The surprising conclusion is that in the quantum world, all things have a dual identity as waves or particles, depending on how you handle them.

In our everyday world, the things around us behave like classical objects with stable identities. They behave like big particles. It's only when we get right down to their atoms that we have to accept quantum duality in order to explain how the electrons in cloudy

orbitals keep the atoms stable. As for light, its wave character is obvious every day as we bathe in its radiance and see it refracted and diffracted in air and glass and water. Only on solar panels, where photons kick up electrons to generate electricity, do we glimpse the deeper truth. For people, regarded as bodies they're obviously particulate individuals, but if we consider them in their essence we may do better to see them as part of a planet-sized human wave.

Curiously, the waves of quantum theory are waves of probability. The waves say where the corresponding particle might be found, if we made a measurement. Where the wave is strong, the probability is high that we'd find the particle there. The possible locations of the particles in question define a set of possible worlds. Recall our discussion of past and future. The past is fixed but the future offers a set of possibilities, and we choose a path. Well, that's how it goes here, too. We see a wave phenomenon and we make a delicate measurement. We then find that the corresponding particle either is or isn't where we measured. For example, we shine a very dim light onto a very sensitive photometer and find that the meter either does or doesn't register a photon. It seems that a set of possible worlds has collapsed down to a unique observed world. This is a deep puzzle of modern physics.

The particle–wave contrast is not yet the sharpest formulation of the issue. Physicists call particles with half-integer spin *fermions* (after Enrico Fermi) and joke that they're antisocial because they obey the Pauli exclusion principle (after Wolfgang Pauli), which says fermions with the same quantum numbers can't be in the same place at the same time. By contrast, they call particles with integer spin *bosons* (after Satyendra Bose) and joke that they're friendly because they happily crowd together to make big waves. So a sharper formulation of the quantum analogy is that humans seem like fermions but perhaps they're really more like bosons. I'll stick to particles and waves to spare you all these new words.

Now we're ready to explain the relevance of all this to the human predicament. Where am I? I know where my body is, but I don't really have much idea of where my mind is, as we saw earlier. My mind is at least as important to me as my body, so it's actually rather odd to say I know where I am. The same goes for all of us. Our bodily differences are minor compared with the shared mental life we enjoy in our networked world. When we act through freely chosen avatars, the

peculiarities of the bodies slumped in the chairs behind the screens become about as relevant as the properties of the flesh below our skins. In both cases, the support infrastructure of flesh or bodies is essential and fascinating in its own way, but its main significance is whether it provides reliable support for the mental life that steers the avatars.

Here comes the crunch. If we're all delocalized like this, then we interfere with each other in ways that we can't control in our respective conscious mindworlds. As individuals, we can no longer reliably calculate the probability that we make this or that choice. To do so, we'd need to know what all the other people who interact with us are thinking. In other words, we behave not like a bunch of separate fermions but like bosons in a single human wave.

You may still be thinking this looks like a very loose analogy. It may remind you of claims that Westerners are individualists whereas East Asian people are conformists. The idea that we'll soon need to merge our minds may sound like preaching. It may sound like I'm saying we should all be nicer to each other and learn to cooperate more smoothly. But it's not that at all. It's supposed to be an exact logical claim about human identity in the autobiology of Globorg. I'm suggesting that the final fact about humans is not that they're particular creatures in particular spatiotemporal locations but that we form a wavelike entity called humanity. This wavelike entity is a delocalized phenomenon in the logical stack that implements Globorg. The analogy from quantum physics is supposed to be exact.

Fine, you say, but so what? Beyond what makes us special as individuals, we share a common state called humanity. We're all attracted to the same sports gods and screen goddesses. What practical effect does this have on the semantics of the words "I" and "we"?

Remember the horrible danger that looms if we lose our special status in a world of avatars and robots. We become divided and conquered. Weak individuals are pruned off the tree of life and only the strong survive. A superhuman chainsaw trims us down to a few specimens for the human zoo. If we lose touch with the dominant life form on Earth, we're doomed. And the dominant life form will be Globorg.

Our salvation is to identify with Globorg

Our salvation is to identify with Globorg. It's not a separate thing. It's us. We're Globorg. We work together in solidarity to create Globorg and to flourish together within it. But this implies that the higher referent of the word "I" is Globorg, not some particular human being. We use the word "I" in many ways. Think of logical selves as like a hierarchy of Buddhas, reaching up into a kind of inner heaven, or as a set of nested Russian dolls. When we talk about ourselves, we're vague about how high or deep we are.

Our inner self is Globorg. The self of my self is a global agent with a layer of humanity within it. Beyond that, the self of the self of my self would be a yet higher Buddha. For practical purposes, thinking in the delocalized mind of our web world is the highest embodiment of self. For us, the referent of "I" is the thinking thing: "I think, therefore I am." We delocalize in the web because our minds extend beyond our bodies and merge in the group mind, or rather in a series of minds within minds, extending up through the clouds to the global mind.

If our minds fuzz out in Globorg, you may think the horrible danger fades away. If Globorg is the result of our own organized effort, how can it loom over us and threaten us? How can it break free of us?

Look back at human history. Cannons and tanks and planes and bombs were our creations, yet they threatened us. Man's inhumanity to man is one of the firmest constants in history. And Globorg is the ultimate weapon system. Anyone who's not part of the royal *we* of Globorg is doomed. Anyone with a self that remains obstinately rooted in the particularity of his or her person is doomed. Salvation lies in surfing the human wave with the best of the rest.

So we shall join the wave of the future. But before we can do that, we have to trust it. And there are still far too many ways to break that trust.

1101
Trust Busters

To join the Globorg club, you first have to trust it. And how can you trust a global Moloch that threatens to swallow up humanity and digest us down to a bunch of avatar drivers whose highest hope is to be reincarnated with bionic bodies in the belly of the beast? It's not a very inviting value proposition!

For that reason, we can expect a lot of organized resistance to the globalizing tendencies. In fact, we can expect a few more wars along the road to political unification under the GO umbrella. This century will certainly feature its share of ugliness and horror as we fight our way to globally harmonized political, economic, legal, religious, and cultural activity among the survivors on Earth.

The twentieth century featured the rise of technology from primitive heat engines to powerful logic engines in the Internet. Along the way, it featured the rise and containment of Germany and Japan, which as nation states were strong in the disciplined use of technology for national goals. In the first half of the century, Germany and Japan were weak at pooling their goals in a global game plan, but in the second half they got their acts together as team players. The century also featured the rise and fall of communism in Russia and the West, and its transformation to a new governing paradigm in China. The century's political keynote was the rise to global hegemony of the United States of America, along with its satellite institutions such as the United Nations, the World Bank, and NATO.

The present century starts with a new challenge. The American hegemony is already crumbling. Islamist resistance to American dominance combined with American financial indebtedness to China poses deep questions about the sustainability of the American dream. That dream is based on letting people pursue their own interests in a secular state. People of all kinds are encouraged to believe what they want in pursuit of their own happiness, so long as they accept the sway of the state and keep consuming the goods. The American social model is democratic capitalism. By contrast, the Chinese model is benign dictatorship of the party. Unlike its Russian prototype, the Chinese

model of communism is pragmatic and enterprise-friendly. It's also realized in a community of like-minded people who already embody the human-wave solidarity that Globorg will eventually require. By contrast again, the Islamist model is pure reaction, in deference to a glorious past, but fired by monotheist zeal, which has something to tell us.

My ambition in this chapter is to extrapolate from these facts and outline the next few decades. This will serve as a good reality check for the GO scenario. It will also help us to assess how far people can learn to trust Globorg. That in turn will enable us to guess how fast the scenario will unfold, and how much pain to expect along the way. So many predictions go wrong in so many ways that I expect only modest success here, but even a poorly constructed extrapolation is better than none at all. The result will at least have the virtue of recognizing the influence of science and technology, which purely political prognoses tend to neglect.

We've already discussed a lot of the issues that are likely to influence these big political developments, so the summary here can be brief and refer back to some of the earlier topics. What's new here is the topspin on the story. Given that Globorg will happen somehow before the century closes, and given all the spin-offs for our lifestyles that will accompany that happening, how much pain can we expect to precede the gain? How much mistrust do we have to fight on our way to utopia?

Robots will be the winners in the next Pacific war

As we saw several times, trust is easy for people who can hope for comfortable positions in Globorg. Engineers who make robots and programmers who make avatars can expect a good life as the revolution unfolds. But the subsistence farmers whose land is expropriated and the unskilled workers who are laid off from dying industries won't have it so good. It's those people, the disadvantaged and disaffected, who are likely to make trouble, and quite understandably too. If you sense that you're about to get steamrollered by a global juggernaut, what's the payoff in just lying back and dreaming of sainthood? No, we can expect violent resistance.

On the other hand, if the juggernaut grows from the institutions in our world that stand to profit most from robots and advanced medical technology, we should all be very cautious. The most likely winners will be military-industrial complexes like the ecosystem around the Pentagon. To me, robot weapon systems in the hands of troops with access to good medical services look like a powerful force for good in the world, but I can well imagine others thinking otherwise. It all depends on what the forces are ordered to do.

It's one of the clichés of our time that modern Washington is like ancient Rome. The American military establishment is as massive and powerful on some reckonings as those of the rest of the world put together, and certainly there's no single opponent that could prevail against it. This invites comparison with the military machine of the Roman Empire some two thousand years ago, which could outfight any other machine in the known world for several centuries.

The problem is that power corrupts. Idealists outside the power elite who see the corruption can weaken the base of the machine and hasten its downfall. Christians hastened the fall of Rome, at least according to Edward Gibbon's monumental history of the decline and fall of the Roman Empire (first published in six volumes between 1776 and 1789), and prepared the base for a new civilization centuries later in Europe. A new movement may undermine the global *Pax Americana* in the course of this century.

In our time, the corruption in Washington is caused by the divisive role of money in U.S. politics. Americans proudly show off their democracy as an example for potentates across the world to emulate, but what the world sees is only a glittering surface. The truth behind the façade is revealed by the unchallenged influence of lobbyists, the ready acceptance of gerrymandering, the effective duopoly of the Republican and Democrat party machines, the pork-barrel politics, the undisguised support for military dictators in client regimes world-wide, and a lot more. It's still democracy, but it's very imperfect. It only works as intended when things are going well and you can always raise money to fight money, dollar for dollar. The tragedy is that no one inside the establishment can fix it when the easy money dries up and it's the rich versus the rest. And the point is that when things turn bad, the last thing we want is a demagogue in command of the mightiest military machine in history.

This is where China comes in. Because the central committee of the Communist Party of China more closely resembles the executive board of a giant corporation than the top level of a democratic structure, the Chinese economy remains under firmer government control than it would be in a more democratic polity. China practices state capitalism and its economic policy is set to pursue strategic goals, such as wresting world hegemony from U.S. hands. Money power is slipping fast from the United States to China, and military power without economic power is worthless, as the example of the Soviet Union demonstrated. As a result, smaller nation states worldwide are beginning to see the Chinese model as a better fit for their needs than the U.S. model. Soon we'll see more command economies appearing in the developing world. These economies will bear out my prediction in an earlier chapter that organizations of all kinds would increasingly resemble business corporations.

It's quite likely that the transfer of leadership across the Pacific will proceed peacefully, but I wouldn't bet much on it. There are too many potential flashpoints, and there's too much pride invested in U.S. military prowess for the wounded giant to go down without a fight.

As for the likely course of such a fight, think robots. They allow tension to escalate steadily to hostilities without the political firebreaks that manned systems tend to impose. In a prolonged standoff, the potential for massive deployment of robot systems translates to huge profits for their manufacturers on both sides. Imagine, for example, the investment needed to deploy Star Wars anti-missile shields on both sides of the Pacific, plus smart minefields to take out military shipping on command but leave commercial traffic intact, plus full-spectrum scanners for all humans and traded goods crossing the Pacific. With these outlays, even a few peaceful decades could bankrupt the United States and hand China the crown. Then the rest of us would have no real option but to trust the Chinese leadership with command of the robot armies of Globorg.

In such a future, could we trust the CPC central committee to rule the robots in the best interests of all humanity? Or would we detect arrogance and hidden agendas in their every decision and feel compelled to fight back? And how would we fight back? Would we end up like the Taliban in Afghanistan, accepting suicidal odds for the sake of a token resistance that ruined the very people it was designed to

liberate? All these questions would be wide open. We can say that victory would transform the Chinese political landscape, so much so that any extrapolations from the present style of CPC leadership are invalid. But in any case, robots will be the winners in the next Pacific war. This will be good for Globorg.

Memes for spreading religious faith by force are toxic

The other issue that looks set to loom large in the next few decades is the integration of Islamic populations into world developments. So far, the Islamic belt is a political shambles by any decent reckoning, but with so much oil still to play for, the pressure to get things right there won't ease off any time soon. Moreover, with large and growing Islamic populations in Europe, the pressure on the ideological side won't ease off either. So it's not just a matter of cleaning up the political establishments in the Islamic belt. We also need to find ways to bring immigrant communities into the modern world.

For me, there are two big issues here, and they're linked. Islamists are spearheading resistance to the globalizing impulse powering the *Pax Americana*. Their passion for this fight arises from an idealist rejection of the rot behind the glittering façade of Western hegemony. They're outraged at the power of big money and the spectacle of personal values in chaos as libertarians pursue pleasure in ways that mock their ideas of honor and pride. They'd rather risk all in struggle than sell out and join the orgy.

The first big issue is political power. The second, the root issue, is monotheism. Christianity has melted down in the Western world to a tolerant and tolerable belief system in a secular political order. To a traditional monotheist who doesn't understand the historical dynamic that has brought us this far, it looks like decadence, and defeat for a faith that once was a source of moral strength. To an Islamist who doesn't understand science, it seems that relaxing the religious strait-jacket achieves nothing but moral and social disaster.

For a psychologist, by contrast, the patriarchic ideology of Islamists has become an obsessive-compulsive disorder of impossible rigidity. It has to break. In a modern setting, the ideology is a sociopathic meme-set. What are memes? Briefly, a meme is an idea that can be copied and inherited in an evolutionary process that selects sets of memes

favoring their own replication in brains. Atheist biologist Richard Dawkins invented the term in his 1976 book *The Selfish Gene* and my fellow consciousness researcher Susan Blackmore popularized it in her 1999 book *The Meme Machine*. The memes for spreading religious faith by force are toxic in a secular world.

By analogy with computer viruses, the meme-set for Islamism is a brain virus. People infected with it are often incurable. They become zombies for their God. In response, our best strategy is to give all boys and girls a modern, secular education, emphasizing practical science to build up trust in science more generally. To cure adult and committed religionists, the best recourse is psychological therapy, if possible, and resistance if not.

The first issue, political power, is easier once we've secularized the moral high ground that traditionalists tend to locate in religion. Then the political issues allow pragmatic solutions. We have two sides, the carrot and the stick, to consider.

The carrot side is covered in part by mass deployment of robots to build new infrastructure. New cities can brings traditional societies up to speed on modern life, as we see along the Arab side of the Gulf, where oil wealth has fomented the growth of a trillion-dollar economy along a formerly deserted coastline. The danger here is that infrastructure development can outrun ideological progress and religious extremists can get their hands on the tools of a modern state. For example, in Iran a militarist regime sponsored by an Islamist clerisy controls nuclear installations and ballistic missiles. Understandably, Arabs and others in the region perceive this as a threat.

The stick side of the politics of containing religionists is hard to apply. The wars in Iraq and Afghanistan have hardly been models of military success. A hypothetical war along similar lines against Iran is practically unthinkable. Western power projection by force of arms has reached a limit against fanatics who prefer martyrdom to defeat. Calling a nuclear bluff is risky.

The disaster waiting to happen is a nuclear war in the region. Such a war would fatally compromise not only extremist Islamism but also Islam itself as a moral tradition, just as the Holocaust compromised Christianity. A movement would arise to condemn the violence of the patriarchic monotheisms and let women take over. In his tortured

reflections on male folly, my old friend Martin Amis has drawn this conclusion and come out in favor of gynocracy.

In any case, the present instability in the Islamic belt is intolerable. Within a few decades, we'll end it. With luck, we can prevent the collateral damage from writing off our huge infrastructure investments in the region before they pay off in improving lives there.

A cult will arise based on faith in the victory of the global "I"

You may recognize the general drift of these conclusions about China and Islam from earlier. But now we can build something new on them. Traditional monotheism is pathological in a modern setting because it's traditional, not because it's monotheistic. In fact, the cultivation of "I" in Globorg that I described a while back is very close to the cult of "I am" in a monotheistic faith. Becoming a boson for Globorg isn't so far away from becoming a zombie for God. You can see the difference, I hope, because the boson story is defended by appeal to a scientific psychology, and is not at all intended to lead to religious veneration of the global organism.

The new element here is that whether the cultivation of "I" in big G is intended to be secular is beside the point. People will do what they're naturally inclined to do, and we know from history that many people are strongly inclined to seek out opportunities to indulge religious feeling. The global organism is the orchestrated sum total of life on Earth, and if that doesn't deserve some kind of veneration it's hard to see what does. Plenty of people now venerate nature in its current relatively disorganized manifestation, without seeing any dissonance in the idea that they're venerating a largely random expression of chemical combinatorics.

Looking on the bright side, veneration of nature is the first step on the road to a more scientific appreciation of nature. Many scientists have recorded their awe and wonder at the glories of the natural world and described how it led them to the more systematic study of natural phenomena. Science is built on a foundation of human feeling. Such feeling is also the basis for religious veneration. And if people who feel no ambition to contribute to science express their feelings in a religious form, so be it. That's human too.

But consider the consequences of such expression on a wider scale. And consider the response of people in traditional religious hierarchies, who as priests and preachers find their flocks deserting them for modern life. It would be astonishing if the unemployed priests didn't take up the challenge of forging a new social tradition on the new ideological foundation of Globorg. It's easy enough to do. The *me* generation of baby boomers already celebrates the secular religion of self quite actively. All it takes now is for an enterprising monotheist to be vouchsafed the revelation that the science of life on Earth is the autobiology of an infinitely recursive self-form in a Platonic heaven (curiously, the form might resemble a Buddha – Google "Buddhabrot" and enjoy the images).

For this reason, uninspiring though it is, I predict that a new-age cult will arise based on faith in the inevitable victory of the global "I". The new faith would be consistent both with Darwinian evolution and with the neuro-savvy psychology of self-determination. It would also preserve a visible and resonant continuity with the "I am" of Mosaic tradition and offer a point of reference for post-Cartesian philosophy. I'm sure such a cult will find a following.

Once the cult is established, the whole historical baggage of the Abrahamic faiths will likely be entrained on the bandwagon. Since the new faith will have roots in science, it will appeal to many secularists who abandoned religion because it was outdated. But there are two caveats. First, the new faith will need time to gain traction and the old religionists will fight back with a vengeance. Second, philosophers like me will see it as just an updated religious mindset that puts faith ahead of reason. Others will object that it erects a superstition of "I" that only makes sense if you accept the analogy of the fermions and bosons. So there will be blowback. I predict decades of strife.

For example, Chinese communists, raised on the orthodoxy of Marxist class war and in a society with a long Buddhist tradition of denial or repudiation of self, may baulk. They may resist the new religion, if only because it wasn't invented in China. If so, we'd have a new ideological dimension to the clash of King Kong versus Godzilla. Then the next Pacific war could be much more alarming.

If you think that fighting over the concept of self is too absurd even for bosons, recall the fights with heretics in the early history of Christianity. Still, if brain modeling and the neurosciences develop as

well as I imagine they will, psychologists will understand the self well enough to moderate the clash. A peace treaty will ratify a schism in the new church of self, and the global self will develop differently in the two hemispheric churches. Globorg will use its left-brain, right-brain views to enjoy stereo consciousness.

Globorg has little need of more widespread trust

Now we're a few decades into the century. China and the United States have cemented their economic synergy and agreed to differ on the concept of self. The formerly Islamic belt is on the road to recovery from its wars of resistance to modernization. The two-faced GO religion is growing fast. And the machine infrastructure on the planet is growing so fast that in a movie it resembles a process of crystallization, as if the planet is undergoing a phase change, not from wet Earth to icy Earth but from mud Earth to web Earth. This is Globorg.

With the resolution of the two main political tensions that opened the century, the stage is set for a new round of history. The globe is divided into three time zones – the Chinese zone, the American zone, and the European zone – and business runs 24/7 on a follow-the-sun basis. Tasks and projects run in the cloud, delocalized, and people log in and out as their days dictate. Tension over the global integration of laws and finances is the main political preoccupation of people in power. Human rights are largely agreed, or at least the basics are agreed all round. Everyone accepts human rights in principle but routinely ignores or violates them in practice. The implications for less developed regions like Africa are naturally devastating. All this will be dismally familiar to anyone who reads the history of colonial conquest. None of it will encourage more widespread trust in Globorg.

The sad fact is that Globorg has little need of more widespread trust. The human population of the planet is too big anyway. The global organism could slim down by a billion people and not lose its drive. Most people who occupy nodes in the extended self of Globorg won't be able to afford more than despairing fatalism about the people who fall through the net. As on the Titanic, there isn't enough room in the lifeboats for all the passengers. Spaceship Earth needs the cabin space for robots and their gear more than it needs the extra people or the extra lifeboats. In fact, a little desperation in the lower ranks does

wonders for the motivation of the people a few levels up. As everyone works harder for a more secure foothold, so the economic base of the GO project becomes firmer.

Globorg needs an agency to trump national power plays

Economic security is a necessary condition for Globorg success. But only military security can ensure a long-term *Pax Orbis*. Globorg needs an enforcement agency that can trump national power plays. The history of United Nations peacekeeping missions is far from encouraging, and it's just mad to expect that heavily armed nation states like the United States would ever freely hand over the military tools for policing the planet to a supranational agency. Most security issues will continue to be resolved at national level, either by single states or by temporary coalitions assembled to meet a specific threat.

But one security issue transcends such national arrangements. Nuclear weapons are both too powerful to remain in nationalist hands and too easy (with robot help) to make and deploy to be abandoned entirely. Within Globorg, we shall need a reserve of nuclear weapons available at short notice to deter any breakout states that could blackmail other states all too easily in a denuclearized world lacking a devastating instant response capability.

Fortunately, we already have an international agency for nuclear questions. Unfortunately, the International Atomic Energy Agency (IAEA) can only be effective in policing implementation of the nuclear non-proliferation treaty if it enjoys the active cooperation of nation states such as the United States (which also has its own Machiavellian interests). These states are needed to help out whenever a recalcitrant state such as Iran needs to be coerced. For the project of ridding the world of the threat of nuclear weapons, this is obviously insufficient, as Iranian President Ahmadinejad helpfully pointed out in New York in May 2010. We need new ideas.

My proposal for a new approach is to set up a nuclear weapons bank under IAEA auspices. In 2005, IAEA chairman Dr. Mohamed ElBaradei proposed that the IAEA set up a nuclear fuel bank to give states with nuclear reactors access to fuel without raising proliferation fears among nervous neighbors. In 2009, U.S. President Barack Obama endorsed the idea of such a fuel bank, so we can hope that the idea will

one day become reality. My proposal is to set up a technically similar nuclear weapons bank. Initially, it would just maintain an inventory of nuclear weapons deployed worldwide, take possession of retired weapons and arrange for their secure disposal, and make new plans. The plans would be for taking formal ownership of all nukes currently owned by GO states and lending them back out to those states on a long-term basis for deployment in ways that ensure effective deterrence. We can hope that once we have detailed plans for such central administration, putting the weapons under a GO veto will become a live option on the Globorg agenda.

As we saw in an earlier chapter, internationalists can hope that the European Union will inherit the nuclear deterrents of France and the United Kingdom. The way is then clear for the four or five main nuclear powers (the European Union, the United States, Russia, China, and perhaps India) to agree on how and when to put all their nuclear weapons under a GO veto. Once the practical details are hammered out, a formal GO command structure could be erected to control the weapons. The weapons would be the symbols of potency behind a GO mandate to disarm every remaining nuclear-armed state in the world. If Globorg controls the nukes, it has a stick to shepherd its global flock of nation states.

Thus we can hope that all the nuclear arsenals in the world will come under GO control. This would be a massive step forward in the utopian project of ridding the world of nuclear weapons. But I say we should stop short of complete nuclear disarmament anyway. If Globorg retains its nukes, it can threaten any hostile aliens dropping in from a starship invasion fleet with extermination at a few seconds' notice if necessary. Following Stephen Hawking's 2010 warning about the likely unfriendliness of visiting aliens, I think this would be a prudent capability to retain in the eventuality that astronomers discovered signs of civilization on a planet orbiting a nearby star.

We shall domesticate ourselves with our own technology

The political struggle for GO control of loose nukes is mere top dressing for the volcanic eruption of innovation from the hotbeds of science and technology. As we saw, the life sciences will vent actively during this century. We need to appreciate the changes they'll bring.

Synthetic life in numerous forms will bring practical benefits without end to the people who deploy those forms. New fuels and foodstuffs, new drugs and medications, and even artificial wombs and living seats and beds will enhance our lifestyles. Evolution will still throw up new plagues and diseases to cull us, but the medical community will organize itself ever more effectively to combat them and confer high expectations of managed health on all who subscribe.

Some prophets of eternal youth think we may soon learn to medicate ourselves so well that we – you and I and our younger friends and relatives – can live forever. I doubt it. The chemical stacks we inhabit offer too many failure modes. As soon as we correct one problem, another bunch of problems looms into view. It's a combinatorial explosion. Anyway, who wants to inhabit a stack that requires a 24/7 regime of pill-popping and exercise schedules?

No, we'll rethink the whole incarnation thing and port our minds into new hardware. The human frame is a fine survival machine, but one lifetime with its toothaches, headaches, and backaches is enough. With all the medical attention in the world, we'll find excuses enough one fine day to turn off the rejuvenator bot and delocalize to paradise. The insurance companies who fund the bots will respect our decision and confer memorial honors accordingly.

As a species, we shall end up domesticating ourselves with our own technology. The new life forms we make will gently nudge us in the required directions, much as guardrails guide livestock into slaughterhouses. We shall live long and healthy lives, and then die quick and painless deaths in transit cells of appropriate design and manufacture. We shall do away with our mortal bodies without a moment's regret. The taboo around suicide will disappear as fast as the taboo around masturbation did a generation ago. As our best-before date looms, we shall distract ourselves with online games so thoroughly that we hardly care that one final cursor click will end our embodied days with photonic or plasmonic efficiency. People will come and go like flowers in spring, with no more purpose than to build up the mulch of hi-tech infrastructure that nourishes the greater glory of the global self.

Life and death are all very well, but what we do on the way is more interesting. With synthetic life we can make our own lives a lot more interesting. Recall that synthetic life is really picotech. That means the forms we make will be just machines. They won't need to eat anything

we'd recognize as food and they won't need to have any interest in replication. For example, they may recycle their existing molecules so efficiently that they need no more than electricity as input. Their idea of replication will at most be budding a clone, for example inside a human body to form an army of cohorts to eat up a tumor before committing PCD.

Naturally evolved life like us is tragically limited by the need to be maximally autonomous in its ability to acquire raw materials and to manage its replication processes. We humans need to be able to flourish and multiply in a merciless jungle of selfish competitors.

By contrast, synthetic life has all the support from us that it needs. It can focus exclusively on the task at hand, just as machines do. The new life we synthesize won't be capable of autonomous life. It will always be dependent on the labs and the fabs in its conceptual backend. But that won't matter. Redundant capabilities aren't worth the cost.

In short, new life is picotech. What that means practically is that we'll treat new life forms casually, like machines. We'll use them up and throw them away with no more thought than for an old iPod. We'll get used to the fact that any life form has a best-before date and can be discarded quickly and easily in an approved recycling facility. We'll even see our own lives that way.

Without their picotech, users will be just biomass

The special novelty of picotech gadgets will be how they interface with us. They'll grow into us, they'll merge with us, and they'll fix us up and improve us. If I lose one of my natural arms in an accident, I'll simply buy a new picotech arm (if I'm well insured) and it'll grow seamlessly onto the stump. If I can afford it, I'll buy myself a better arm, with lots of fancy features that no natural arm ever had. If I'm impressed, I'll buy another fancy arm and hack off my remaining natural one to make room for it.

Web documentaries will feature enthusiastic freaks with bionic bodies. The freaks will have replaced all their natural extremities and peripherals with weirdly exotic but impressively capable new parts. For all of us, Borg implants to put you permanently online even when you have no kit on (no headset, no smart clothes) will become

completely unremarkable. Arms and legs that plug efficiently into exoskeletal robots will be common.

Deeper inside, microscopic picobots that cruise in the bloodstream and boldly go forth to solve problems and right wrongs will be practically universal. Everyone will want regular shots of those little bloodbots to keep them healthy and in shape. Some bloodbots will enhance natural functions enormously. For example, they'll deliver oxygen to tired muscles much faster than natural molecular systems. This will enhance athletic performance and hence be banned in sporting events. Other bloodbots will home in on genitalia to enhance sexual pleasure or regulate fertility. Men will just need to think hard to get an erection and women will just need to think fast to seal off their uterus. Sexual interactions among people addicted to such things will become as casual as handshakes.

All this will make users dependent on the picotech corporations. Without their picotech, users will be just biomass. People who can't get their fix of new bloodbots will die. People whose new arm or leg develops a bug will be out of action until the manufacturer releases a patch. Women whose cervical bots malfunction will become pregnant or sterile. Men whose scrotal bots malfunction will become priapic or impotent. Worst of all, people whose Borg implants fail will go offline. They'll lose their instant credit. They won't even be able to walk through a smart door. Lawsuits will fly and lawyers will prosper.

At present, the vendors and pushers of microtech and nanotech are unlikely to have medical qualifications. If you buy an iPod, its impact on your health and safety is the last thing on your mind. But this will change with intimate picotech. If you're a reasonable person, you're not going to inject a new line of bloodbots without some kind of medical reassurance. That may come online, so the shop assistant needn't be a doctor, but it'll have to come somehow. You'll need substantial reassurance that the new bots have passed some rigorous medical tests for toxicity and side effects and so on. You'll need to trust the vendors.

Trust has popped up again. Nowadays, we have a hard time trusting software vendors we've never heard of. Given the choice, I'll buy a product from a big-brand corporation just because I know the name and trust their reputation, but it wouldn't take much to knock that trust. For picotech, we'll have to start over and we'll set a much higher

bar. Gone will be the days when buggy products were released in a rush and then patched in the weeks and years to follow. One buggy release for intimate tech and you've got a million dead customers.

Such things will happen all too easily. A sloppy lab test in China could do more damage than a nuclear strike. In early 2010, the head of Toyota was called to give congressional testimony about the company response to a series of fatal car accidents. It's not hard to imagine an episode a few decades from now where a quality-control disaster leaves a body count in the megadeath range. We should demand massive reassurance before putting our trust in life-critical implants and injections that make us addicts.

We shall need to trust the big picotech corporations

We shall need to trust the big picotech corporations. When the intimate tech we crave comes from global corporations that may equally well be based in any developed national ecosystem, the trust we seek will be in corporations rather than in national governments. We shall expect the corporations to adopt global best practices in any case, wherever they make their home base. And with modern business software, there need be no doubt that they'll do just that.

The political machinations that concern us will be business politics. We'll care about the good or evil deeds of big corporations and lose interest in the jockeying of nation states on the world stage. Who cares if this or that nation state is up or down a few places in the global GNP ranking? What I want to know is whether my favorite corporate brands are going up or down on the trust or quality rankings. This century will see the effective end of national politics as the lead horse in the global race and the rise of corporate politics in its place.

The leading status of national politics ends when G2 becomes GO. The rise to dominance of Chimerica (thanks to celebrity historian Niall Ferguson for the word) from the economic showdown between China and the United States of America will make the national back-and-forth between the also-rans on the racetrack to fill out the GO map look like a sideshow. Much like war after the Cold War, politics will never be the same again. The globe is too small to let games like that continue as before. The rules of politics, like the rules of war, are now clear enough: minimal destruction, precisely targeted coercion,

maximal agreement on aims and scope, clear chain of command, feedback from below, universal human rights, and so on. In the global village, all war is policing and all politics is local politics.

One huge issue of trust remains. Can I trust the techies who work in the big corporations that make the picotech products I'm so addicted to? This isn't a question of corporate best practices or executive intentions, because those are givens. It's about the code monkeys who in a picotech world will do most of the grunt work involved in the creation of the new life.

If the picotech corporations outsource the work to the lowest bidder, expect a few disasters. But of course this will change. Corporate officers will see the PR value of responsible hiring policies. They'll start to demand professional certification for their code monkeys. Software engineering will be a career as respectable as brain surgery. In some case it'll be brain surgery, so we'd better hope things turn out this way. One megadeath lawsuit will throw all the old rules about peanuts for monkeys out of the window.

Slowly, then, it may seem safe to trust the corporate gorillas to get their code right first time, not ten releases later. We saw similar developments in earlier waves of new tech. The first aircraft were often deathtraps, but a century later, with CAD/CAM to help us, planes are safe from the maiden flight onward. Now it's terrorists we worry about in planes.

You can see the next problem here. Corporate servants will work in a rule-bound environment and hence merit some level of trust. But anarchists, terrorists, jihadists, hackers, freaks, and bad eggs generally will be with us until the end of the century, and probably far beyond. I may trust the product in its sealed packaging, but the moment it goes online I'm at risk. How can I trust the pervasive sentience of my infrastructural environment to protect me from that lot?

We won't need feral kids in the Globorg future

The answer is easy in principle but startling in practice. We engineer better people. We start by revising the education sector more thoroughly than ever before. There are numerous policy questions here that go far beyond my scope, but the basic idea is simple. By offering total immersion in virtual or augmented realities, educational

software can reprogram brains far more deeply and thoroughly than an overworked human teacher ever could. As we saw, gaming software can make the learning experience involving and challenging. The best prototypes here are flight simulators for pilots. If we can engineer similarly gripping experiences for all students in all disciplines, we can expect fewer duds and misfits later.

The sixteenth-century Jesuit Francis Xavier said, in translation, "Give me the child until he is seven and I will give you the man." In short, early education is decisive – the earlier the better. Most kids are stunted before they ever reach adolescence. Their minds are neglected and their dreams evaporate. What's left is a human ape that goes through the motions of a typical life and may even earn some distinction, but will never feel the joy of a soul that soars among the stars. We must do more to help such kids. We don't need feral kids now in a civilized world, and we'll have much less use for them in the Globorg future.

Sadly, many millions of people in our world may be beyond salvation already, since they have little hope of raising more fortunate kids in time to rescue their genomes before Globorg kicks in. This doesn't mean we should abandon them, in fact rather the opposite. We need to keep rethinking our notions of what makes a good and productive human life until we can find useful work even for relatively feral people.

Again, the big issue is trust. Can we trust people who haven't lived in a VR pod for the first seven years of their life to fit in with the higher plans of Globorg? Trick question – we only have to think for a moment to see its absurdity. Can we trust people who *have* lived in a VR pod for seven years to fit in with any world worth living in? That's a question about how well we can formulate a VR-game curriculum for future generations. Until we've developed games that try to meet the need here, we're only guessing.

The curriculum question will take at least a generation to answer. Meanwhile, the relatively feral lives that most of us have lived to date will grow gradually toward the more refined lifestyles that Globorg will offer in future. We want a future that we, here and now and as we are, can in all conscience endorse as the right way to go. Ask not whether our feral lives are right for Globorg but whether Globorg is right for us. Can we trust Globorg?

The short answer to that question is no. Of course we can't trust a global Moloch that will mince humans down to their picoparts and regurgitate them as sanitized, synthesized biodroids. But with all the mistrust in the world, we don't have much choice. It's coming, and as individuals we can't stop it. The best we can do is to redesign the monster from within before it ever gets off the drawing board. We need to think clearly and realistically about what we want and how we can get it. As soon as we begin to get fatalistic, someone else will spearhead the thrust toward GO. Then, in the fullness of time, the monster will confront us as an alien power.

We need to make the best of Globorg

Globorg is our own creation. It's a great spaceship with a brilliant technical specification, but it's also the ark for the living power in our genomes and for all else that we hold dear. It's the new incarnation of the great attractor, "I am," the ineffable ecstasy of sheer being in time, that has brought the story of life on Earth this far. It's the latest shape of the polymorphic form that has dangled ahead of life on Earth like a glorious carrot, daring us to reach out for it with flesh piled high in ever more highly evolved organisms. It's the shape of the future, looming out of the fog of all possible futures as the best of a bad bunch. (Physicists will confirm that almost all quantum futures are utterly hideous and only vanishingly few of them are even survivable.)

Since Globorg is the big result of our human project for several millennia, we'd better like it. It's our plan A. I'm not aware that we have a feasible plan B. If we don't learn to live as one within Globorg, we survive only to die without it.

So we need to make the best of Globorg. Making the best of Globorg means learning to trust it. It means identifying with it as manifest destiny, so that its destiny is my destiny. That takes a very special psychology, which most of us haven't yet begun to master. But we can.

1110
Cloud Immortality

The glorious future: I trust Globorg, I live a happy life as picoperfect meat in a padded cell, and I flip the switch to delocalize when I've had enough. It doesn't compute. Something's missing. I want a shot at immortality. I want to do great deeds and change the course of history. I want to boggle at the glory. I want it all!

It will come to pass. The science is ready. The new psychology will change the world as radically as the new physics changed the world a hundred years ago. As layers of self in the higher self of Globorg, we're living a logic that loops always on the brink of paradox. Our every step in time is a step in logic from world to world, and each world holds surprises. Worlds are ordered in spacetime, and each one defines me, in its own way, as a subject bringing that world to a focused unity. My mind is bigger than my body.

We humans have localized bodies and delocalized minds. Our bodies are like particles and our minds are like waves. Our bodies are separate but our minds interfere with each other. How do we know this? We know because of how our minds are implemented. Our bodies host our minds on a hardware stack that tops out in an ever-changing field of electromagnetic vibrations created by groups of synchronously firing neurons in our brains. That was a hard sentence, but the core meaning is simple. Our brainwaves carry our minds. What looks like a mind from the inside looks like a vibrating field of electromagnetic energy from outside. But electromagnetic energy is carried in photons, and photons are bosons. They interfere and make waves. Brainwave photons have wavelengths in the decahertz region of the electromagnetic spectrum. This is deep down in the radio spectrum, where the wavelengths are tens of megameters. From Werner Heisenberg's famous uncertainty principle, we immediately infer that the spatial uncertainty of the waves in free space is several thousands of kilometers, like the size of the planet. In short, that's an image of how delocalized our minds are.

We saw earlier that even from inside, our minds seem delocalized. We project our thoughts onto things in the outside world and get the

feeling that our mental lives are more out there than inside ours heads, at least when we're not daydreaming. We also know full well that when two or more of us concentrate on the same physical things before our eyes, and refer to those things in what we say, we're usually not talking at cross-purposes about the respective images in our brains. We identify the virtual images that our brains generate with things outside our heads. We don't just sit inside our Cartesian theaters and talk about our dreams. So our minds overlap and interpenetrate. That's part of what minds are for. They help us communicate. Our minds overlap in a huge patchwork quilt or network that spans the planet.

The trick to master if you yearn for immortality is to see that the great cloud of mindworlds you live in shades seamlessly into the great cloud that veils Globorg. Recall that Globorg is in large part a network phenomenon, where the planetwide infrastructure of computers and cables and all the rest pumps up a magic cloud that swallows up all you can offer and delivers back all you need. That's the idealized image, at least, which will take a few decades of hard work to realize in the grubby everyday reality of defeats and disappointments. If your mind delocalizes into that hardware-based cloud and lives a reverberant life in the online events that follow, you've made it. You're as immortal as an angel.

If your mind delocalizes, you're as immortal as an angel

The signals that shoot back and forth in the worldwide web are physically comparable with the signals that loop around in your brain. If you can be implemented in your brain, you can be implemented in the web, too, in principle. We're not there yet, but the hope is reasonable. At present you have a reverberant center in your brain that has no analog in the web. That set of reverberations in the echo chamber in your head goes right down to a quantum phenomenon (in my scientifically speculative opinion), which gives the resulting cloud a probabilistic presence in the virtual superposition of all possible futures. Your mind is a zone in which shapes come into being from the future and fall into existence in the past. It's not a thing like your body. Above your mind is only the insubstantial realm of forms that people

used to describe poetically in terms of souls and spirit and heavenly hosts.

Here philosophers are accustomed to raising an objection. My mind is mine and yours is yours. We each have privileged access to our own minds and no access in principle to each other's. So we each live in a mental stovepipe or silo that insulates us from all that dirty stuff out there and leads us as pure souls to our own experience of heaven or hell.

With all due respect to my philosophical colleagues, that objection is bullshit. Privileged access is an illusion. The incorrigibility of introspective impressions is illusory. The murky shadows inside our heads can always be thrown into sharper relief by the nuclear flash of inner revelation, even if we don't often enjoy the harsh clarity of such inner brilliance. The prisons of our minds are only contingently secluded, in other words, and nothing in principle prevents the next mindworld we inhabit from shedding light where previously all was dim. In physiological terms, nothing in principle prevents a researcher from poking an electrode into my neocortical neurons and administering a shock sufficient to trigger unexpected illumination.

Our minds are open to outside input. It happens every waking moment. If we could run a neural cable (as in the genial movie *Avatar*) from my brain to yours, I could in principle save myself the effort of writing all these words, because you'd get a much more direct and vivid impression of what I had to say than you can by reflecting on the trickle of textual data delivered through your eyes to your brain. Happily, the discipline of marshalling words on a screen ensures that my thoughts arrive in distinctly more coherent form (believe me, they do), so the indirection through text is certainly not a complete waste of everyone's time, but it's worth pondering how much more direct the process could be. All our information technology is concerned to try to shorten that indirection and make the communicated experiences more vivid. And that technology is getting better every day. Soon, I confidently predict, our everyday success at sparking structured visions in each other will be so great that the old philosophical mumblings about privileged access will seem like the despairing complaints of losers who just can't connect with each other.

Our minds overlap and interfere and create new holistic forms

The revolution we experience when we abandon the personal stovepipe model is enormous. Our minds overlap and interfere and create new holistic forms. The holoform that springs instantly to my personal mind is the global mind of Globorg. All the minds of all the people (and other animals) on Earth merge in a cloud mind of hardly imaginable complexity and power.

Globorg extends seamlessly to a cosmic mind much as we humans extend seamlessly to the global mind. Our present understanding of the cosmic mind is that it started in the Big Bang and now extends to the event horizon of our Hubble bubble, with an overall configuration mapped by the Einstein macrogeometries we looked at earlier.

Whatever the details, the stuff that hosts any bigger mind we can enjoy is here on Earth. All the organized activity that we undertake to pursue astronomy and cosmology is part of the Globorg effort to understand its place in nature. And all this leaves the personal stovepipe hats of the philosophers far behind.

But I don't want to diss all philosophers here. The philosophical understanding of immortality is precisely what we need to avoid falling into the crass materialism of extended life in Globorg. As humans, we naturally lament the brevity of our mortal lives. As citizens of Globorg, we'll naturally do all we can to develop advanced medical services, bionic body parts, collaborative software for mind sharing, and so on. All this will give us the feeling that we can enjoy lives that are greatly extended in both time and space. If I can live for centuries and exert global effects through my various avatars, I feel better. I'm a step closer to immortality. But the eternity of the true visionary is still infinitely remote. I can live all the extra centuries I want, but living forever is something else again. Here only philosophy can help.

The philosophical phrase I'm looking for is *sub specie aeternitatis*, which is Latin for "under the aspect of eternity" and is commonly attributed to the seventeenth-century philosopher Baruch Spinoza. To denote this tricky idea, let me invent the word *extime*. What we want, it seems is to live not only in time but also in extime. We want our lives to extend beyond time or to have a lived dimension that's external to

time. We want our lives to mean or have meant something when all time is gone, when the calendar is exterminated.

But – to quote the rock philosopher Mick Jagger – you can't always get what you want. Recall philosopher Thomas Nagel's phrase for what scientists want: the view from nowhere. You may recall my reply. Scientists can't get it. What they can get is the view from anywhere. And that's all they need. A maximally generic view from anywhere is as close as we can get or need to get to the mythical view from nowhere. Exactly the same applies to a life in extime. We can't have it. Life is life in time because lived experience is temporal experience. What we can have is a life so lived that it looks good from any time. Any time is as near as we get to extime.

Life is in time because lived experience is temporal

As you can well imagine, this is a question with a long and heavy philosophical baggage train. To spare you the blow-by-blow history, I'll cut straight to the two twentieth-century figures I like to call the Time Lords: Albert Einstein and his friend the mathematical logician Kurt Gödel. After discussing the issue regularly in their emeritus years at Princeton, Einstein and Gödel concluded that time was an illusion. They imagined that the real truth was eternal and that, with the help of physics and mathematics, they could rise above the temporal view and enjoy eternity for themselves. The Einstein block universe was an eternal block of spacetime filled with mass and energy in the familiar way, to give a classical cosmology in which the entire future is already fixed, just waiting for our little minds to catch up with it. Gödel's refinement of this picture was to find a solution to Einstein's cosmological equations that put the universe in a spin and put loops in time. But we needn't go on. Quantum physics exploded the whole classical picture and brought time back with a vengeance as the dimension of collapsing probabilities.

To update the tale, superstring cosmologist Brian Greene wrote a sympathetic review of the Einstein block universe in his 2004 book *The Fabric of the Cosmos*. The Time Lords made a noble effort to transcend time. They deserve due respect for their effort. Einstein was deeply impressed by Spinoza's philosophy and Gödel was equally impressed by Kant's philosophy. Kant, you recall, said the phenomenal

world was ordered by categories, including time, but the ultimate world behind it was beyond all categories. For Kant, the ultimate world was the objective correlate of the Transcendental Ego.

Sorry, but I can't resist a quick historical footnote here. Hegel objected that the ultimate world and the Transcendental Ego collapsed into a dialectical unity of opposites, and his objection greatly impressed Marx and Lenin (as well as Mao Zedong). About when Marx was busy working out his dialectical materialism, Friedrich Nietzsche toyed with a concept of eternal recurrence that reminds me of Gödel's time loops. The moral of this historical footnote: There's no nonsense too deep to stop philosophers from swimming in it.

Now you know how hard it is to break out from time to extime. We're temporal beings. We can rage against the dying of the light all we want, but we can't change that hard fact. Even the Time Lords didn't make it. They only imagined they did. If you like, extime is the asymptotic limit of a series of maximally generic conceptions of time and temporal being that blur out in a quantum multiverse. However we cut it, we're in the here and now, and what's happening now changes from moment to moment.

Each change is a step in time. It's as simple as that. And each step is a step with a certain physical probability. That fact, pursued systematically, implies that in general our entropy increases. Each step takes us one step closer to thermodynamic meltdown. As humans, the best we can hope for is a long and healthy life.

With wonderful serendipity, a long and healthy life is precisely what Globorg can offer us. It does so the incremental way, with no philosophical finesse. It offers better medical services, better nutrition, better living conditions, better access to arts and culture, better odds of surviving accidents and disasters, and so on. Then it offers bionic body parts up to and including complete droid bodies, mind uploads up to and including full sentience transfer, and a shot at living as long as the Sun keeps shining.

Globorg offers mind uploads up to full sentience transfer

What do we mean by sentience transfer? Who ordered that? We did. We don't quite know what we ordered and we can't define it too well, but we'll know it when we see it. Basically, we want what we saw in

the movie *Avatar*, when (spoiler alert!) the Tree of Life transferred Jake's sentience from his human body to his avatar body. Once we can implement such transfers in picotech, we can visit our local droid dealer and try on a few expensive new droid bodies. When we find one we're comfortable with, we can port our avatar presence – transfer our sentience – to our new body and take final leave of the old fleshbag, remembering to tag it for recycling. What we're doing in this scenario is transferring our immediate sense of here and now to a new body, which means moving our core loop to the new host.

Is that possible? The philosopher in me says we should be very careful here. The idea may be as nonsensical as living in extime or indulging in time loops where I revisit the past and kill some of my ancestors. Before we can endorse the idea of casually moving core loops, we need a fuller understanding of how the immediate sense of here and now is implemented in humans. Robotic engineers can emulate many human cognitive functions in their robots, but this elementary awareness of being is way beyond the state of the art.

My own feeling is that this human feature is beyond the state of the art because it's a quantum phenomenon. I think the physical embodiment of a state of mind fuzzes out in a delocalized cloud where self-identity extends uncertainly over what would classically form a discrete set of mental contents. I'd even hazard a guess that the fuzzing out occurs in the paired bulbs of the thalamus, following Rodolfo Llinás' research on its role in thalamo-cortical loops. My idea is that the contents of our immediate awareness cycle around a core loop. The core loop (or loops – there could be a knot or ball of them) is so small and fast that its contents can interfere or blur together in a quantum delocalization phenomenon.

But I could be completely wrong on this. Descartes used his own much simpler physiological model to hazard a guess that the pineal gland was where sentience happened, and that turned out to be wrong. I can't say I'm much more confident about the thalamus. You can ignore the quantum hocus-pocus here just say there's a core loop of some unknown kind.

The mental contents in the core loop are elements of experience. They're generated by neural action from sensory images of things you're aware of in your total experience of now. This is an inherently fuzzy set, as you can verify by considering what researchers call fringe

consciousness. Are you aware of the back of your head, the seat you're sitting on, or the world outside the window? All these things are doubtless at the fringes of your consciousness, ready at a moment's notice to occupy center stage in the Cartesian theater. But however the story goes from here onward, it begins to resemble familiar psychology and we can take it as given.

As for the mechanism behind the core loop, I suggest we wait for the scientists to give us a more exact picture. If it involves or invokes a quantum phenomenon, the reason why current robots don't have a chance of being conscious is clear. They're classical machines, like our computers, just discrete-state automata acting out the (virtual) entries in their machine tables like any Turing machine. We'd need to master the picotechnology of reliable quantum computing before we could implement machine consciousness.

We'll understand the mechanism of consciousness

One way or another, I'm sure, we'll understand the mechanism of consciousness. When we do, and when we build droids with consciousness, we'll find that the uncanny valley we discussed earlier is gone. We'll adopt the intentional stance toward our droids and assume without question that they share our mental states.

Again, a philosophical bugaboo raises its ugly head. In the proposed future when we assume that droids have mental states like ours, some philosophers will say we're fooling ourselves. They'll say the droids don't really have mental states at all, but just behave as if they do. They'll object that only a behaviorist would say droids have minds merely because they behave as if they do, and behaviorism was already out of fashion at the end of the twentieth century. We can even expect a few species chauvinists who say that only humans can have mental states like ours. Droids don't even have a chance.

We can banish this bugaboo quote easily. How do we know that the philosophers who say this have mental states like ours? On their logic, we don't, because all we know is that they behave as if they have mental states like ours, and that's not enough. Well, we know they're human, too, but then again maybe not. Maybe an evil genius sneaked a droid into their place when we weren't looking and now we're debating the issue with a fake.

As you can see, this line of objection begins to resemble a psychiatric disorder called Capgras syndrome where the patient imagines that impostors have replaced his or her loved ones. The impostors look identical to the loved ones, and the mechanism or the motive for the switch is utterly unknown, but the feeling is convincing. We now know that the syndrome is due to the failure of certain recognition neurons in the patients. Perhaps our spooked philosophers have similarly faulty neural wiring.

Seriously, a problem of principle like this reveals a problem in how the philosophers have conceptualized the situation. In this case, the problematic assumption is that there's a fact of the matter beyond all the evidence we could acquire about the mental states we hesitate to attribute to the droid or the fellow human. We imagine there may be further facts because we imagine that we ourselves instantiate mental states not on the basis of external evidence but because of our own incorrigible introspection. Introspection provides the further facts, apparently. But then again maybe it doesn't.

Introspection can be wrong. We can delude ourselves. But the philosopher has a reply: I can't be wrong about my own existence! Just thinking I have a mental state is enough to guarantee that I have one – I think, therefore I am! Well, I reply, the droid can say this too. You say the droid is just mouthing words without meaning. I say the philosopher can do that too.

The logic of reflexivity (sorry, but there's no easier way here) is that I'm anchored in my immediate sentience. I know that I exist without any need for argument. I know that others – philosophers and droids – exist by analogy. I let them play guest roles in my Cartesian drama. I can distribute egos and personalities in my drama any way I like, subject only to good sense and good taste. The other roles all exist on my sufferance in my mindworlds. The "I" of the suffering subject is the premise of the whole drama. This singular fact of reflexivity is celebrated in our culture in the axiomatic "I am" of monotheism, where we're encouraged to put selfish thoughts out of our minds and into the divine beyond. The "I" is the logical navel of the world, or rather of any mindworlds that you or I can construct.

In short, a droid can have an ego too. It only needs to have an axiomatic certainty of its own existence, and predicate everything else on the firm foundation of that existential fact. To return to sentience

transfer, all we need to transfer is that immediate self-certainty. All the rest accumulates around it like a pearl around a grain of sand.

All we need to transfer is immediate self-certainty

We were discussing immortality. What that means, practically, is persistence of ego. My ego can be embodied in any body that evolution, science, or Globorg can provide. At its core, my ego is merely a sense of being that persists in time. I reflect my being and project it as a mind-world around me to reflect my ego. Death is the collapse of a world (as a historical aside, this isn't just my invention but an old Jewish idea).

But consider where we last heard about collapsing worlds. At each step in the quantum process of moving from moment to moment, possible worlds collapse and a single future is realized, or promoted into being. Our physics has led us to the disarming conclusion that we die in each and every moment of elapsed time, to be reborn in each new world that emerges to surround us. Immortality is where we least expected to find it, in the experienced moments of elapsed time and their seamless transitions.

As the popular mystic guru Eckhart Tolle says in his 1999 book, *The Power of Now*, the experienced present is eternal and only its frozen moments pass into the past. If we learn to live in the present, we shall never experience death. Ludwig Wittgenstein said something similar in his oracular early work, the *Tractatus Logico-Philosophicus*.

At risk of overdoing the wordplay, we can play on words here. If being is in time, existence is in extime. Existence is the eternalization of being. This needs explaining. In the linguistic view of philosophy that prevailed in my student days, eternalization was philosophical jargon for the process of replacing words like *here* and *now* in a sentence by spatiotemporal coordinates. For example, if a mother says to her child, "Come here, now!" the eternalized version would have the schematic form, "Come to spacetime event (x, y, z, t)!" The point was that such a sentence had a stable reference, rather as nowadays a permalink has a stable web reference in a world of broken links. This was the sort of wisdom I imbibed as a student. Returning to my word-play, I am in time but I exist in eternity. Whether this claim counts as a mystic truth or a poor joke I leave for you to judge.

I am in time but I exist in eternity

Let's return to the practical world of Globorg. We've reduced immortality to the survival of the core self from moment to moment, in any body that comes to hand. But we deserve a better foretaste of how our minds can delocalize into the clouds we implement in our global networks. Delocalization is a concept from physics and its application to minds sounds like a metaphor. The suspicion arises that we may have misused the metaphor to invoke an illusion.

We know that minds are not localized in brains. Brains host minds in much the same way that computers host cloud services. A cloud service is a process that can be called to perform a task specified in the service. A mind is a process that can be called upon to do the sorts of things that minds do, such as recall facts, trigger speech acts, coordinate bodily movements, and so on. So far we have a match. A process isn't particularly localized. It can be forked, distributed, cloned, killed, restarted, and so on. This raises a question. If a process can be cloned, can a mind be cloned? We don't yet know how, but that can't stop us from speculating.

We won't be able to make quantum backups of our minds

The ability to clone a mind would be a fine thing. Even if the clones weren't functional because they needed the original body to work as intended, they'd be fine snapshots for research and backup purposes. Imagine cloning your current state of mind onto your Apple Time Machine every night. If you suffered a seizure overnight and lost a billion brain cells, you could restore the backup copy to retrieve memories you might otherwise have lost.

Cut. The fantasy has exposed a limitation in the metaphor. States of mind are almost certainly not definable in any clean way as discrete states that we could reliably copy onto a digital medium. If they were, for example as in an IBM Blue Brain simulation a million times bigger than our current ones, it would almost certainly be impractical to run a backup while drinking cocoa before bedtime. And as we noted a while back, the coding of memories to neurons and synapses is far too idiosyncratic to expect a snapshot to be intelligible to anyone. The only way to decode it would be to run it on exactly the same host, down to

the last synapse. In the imagined scenario, that was no longer possible because an overnight seizure had just irrevocably changed your brain architecture. So the only point of making backups is lost and we're back to square one. Can we rescue the metaphor?

Sure we can. The delocalization metaphor we want to understand is a quantum idea. Cloning machine states for backup and recovery is a classical metaphor, at least in the present state of the art. To clone a mind, we may need to go beyond the classical idea of copying it. In other words, we may need to copy not only how it actually is but also how it could be in all possible worlds. In philosopher's jargon, we may need a dispositional copy or a multiworld copy.

This needs explaining. If something looks the same but acts differently, it isn't a dispositional copy. A dispositional copy should be disposed to behave the same way. It should look the same not only in this world, the actual world, but also in all possible worlds. A snapshot of a person looks like that person but doesn't behave like that person. It may be a great static portrait, but you only need to observe it for a few seconds to see it isn't a real person.

The same goes for minds. Because minds are hosted in physical stuff (with all the physical complications that we need quantum theory to describe), we may – just may – need to copy not only how that stuff is in the actual world at the snapshot moment but also how it is in a set of possible worlds. That's more than we can expect from anything like the current generation of Apple Time Machines.

As it happens, we can't clone quantum states. The best we can do is to "teleport" them, as the distinguished Austrian professor Anton Zeilinger does in his quantum physics lab. We can teleport the quantum states of individual source particles over long distances by first entangling the source particle with a messenger particle and then entangling the messenger particle with a target particle. The target particle inherits the properties of the source particle, but the quantum state of the source particle is destroyed in the process.

So the technology of mind mapping may be trickier than a classical computer metaphor suggests. It may be impossible to clone minds, and teleporting a mind from A to B may only work if you destroy the copy at A to create the copy at B. In short, we won't be able to make quantum backups of our minds.

Fine, you say, but so what? How does that ruin the metaphor? The answer is that it doesn't. To ensure our immortality, it might have been nice to make mind backups, but perhaps we can do without them. If minds are delocalized phenomena, perhaps we don't really need anything like backups. But still, what does it mean to say minds are delocalized?

Where am I?

We can understand what the delocalization of mind amounts to more easily if we go indirectly. Let me start with a science-fiction story that I often imagined writing as a teenager (but never did, at least beyond a few pages). I wake up … and everything is radically strange. I don't recognize anything at all. Where am I? I have no idea, but slowly I gather evidence and make deductions and build up a coherent explanation of this new world. You can imagine going on for pages like that until the story gets really fascinating.

What the story shows is that the self who wakes up isn't really anyone until some kind of recognition occurs. Recognition means memories are awakened, not necessarily of episodes from a personal narrative but perhaps of physics or geography, depending on how radically alien you make the story.

Nowadays it's a philosophical commonplace that personal identity is based on organized memories in this way. With the right memories activated, you feel at home in your new mindworld and begin to take up the personal narrative again where you left off in memory. If your memory is blank for any narrative at all, and all you have is physics and so on, you've lost your personality and deserve hospitalization.

Somewhere between activating the little self-loop in immediate awareness – "Where am I?" – and reconstituting all your memories in an autobiography full and rich enough for you to wheel and deal with your customary amazing grace, your personality has come online and pumped up that little self-loop into a whirling galaxy of thoughts and feelings. This presupposes some impressively capable wetware or hardware behind the scenes.

The deeper implication of this just-so story may be disturbing to some readers (who may wish to look away for a paragraph or two). The little core loop that booted the galaxy of a rich personal self is a

rather trivial widget. It's like the Windows AutoRun code on CDs, which is a tiny setup file that merely serves to boot a much bigger body of process code. All the value in our personal core loop comes in the galaxy of memories it stirs up. A bare self is a pitiful thing. So we can expect that implementing a core loop in droid hardware is easier than it might seem at first blush. The rich sense of self we experience with every waking moment arises from our personal galaxy of memories, not from the core loop.

If the core loop is a trivial widget, what's a life worth? This is where the story gets disturbing. The axiomatic self-assertion at the core of our conscious lives is the core loop in action. The suffering self that frames my every mindworld and gives me the space and the grace to assign selfhood to other people in my circle of acquaintances looks big from inside – like the entire universe, in fact – but is really just a tiny nub in a galaxy of stars, like the point at the center of a circle or the period at the end of a sentence. Consider the value of a tiny life, an amoebic life perhaps, with no more than such a loop. Not wishing to moralize, I'll stop there.

Now we can return to delocalization. If our minds are somehow analogous to swirling galaxies, there's no upper bound to how big they can be. Galaxies can be huge and just fuzz out at the edges. Another popular analogy compares a mind with a storm or a hurricane. Again, the key idea is that the mind is indefinitely big and swirls around in a dramatic fashion. I want to generalize these analogies and say a mind is like a field in physics, which in principle is defined over all space-time and just takes high values where the action of the field is dramatic. For example, a magnetic field around an electric conductor or an electric field around a bar magnet creates detectable swirls near the conductor or magnet and fades off to infinity. Why not say minds are like that?

Once we say that minds are extended fields, we have plenty of physical precedents to say they can interfere with each other. With lots of minds and lots of opportunities for interaction, the scope for surprising phenomena is enormous. Minds have a dual character, as extended waves or as assemblies of particular thoughts. With this conception, effects like interference, entanglement, and nonlocality show up on the menu of possible effects we can call on to enrich our theory of mind.

Readers of a nervous disposition may be either anxious or bothered at this point. Why am I going on about quantum mind? Who cares? What difference does it make for me – you – whether an ultimate theory of mind needs quantum physics or not? Why open that unappetizing can of worms here and now?

Believe me, I sympathize. Alongside a serious band of scientific researchers, the quantum mind community (as I've learned to call it after enjoying several conferences where some of its leading members presented their ideas – to admire a few of them in action, see William Arntz's movie *What the Bleep?!*) includes a fringe of charlatans, quacks, and cranks who wallow in the esoteric and baffling aspects of quantum physics. Here and now, the last thing I want is to sully the purity of my message with that sort of thing. But I have to bite the bullet and tell what I think is the truth.

The truth is that if we accept the sort of model of mind I've been presenting here, the future that Globorg offers us is potentially quite enticing. Through the extended effects of our personal minds in the collective mind on the web, we can enjoy a depth and pervasiveness of influence even after our bodies have been recycled that makes dreams of immortality seem less absurd. As our being in the present passes into our existence in the past, the reverberations in the web go on, spreading and ramifying in potentially satisfying ways.

We don't need to get religious to gain this kind of immortality

There's nothing strange or unprecedented in this idea of continuing influence. Every book author whose books are read long after the author's expiration date enjoys this kind of immortality. Every philosopher whose name is invoked again and again like a magic charm in philosophical debates enjoys a similar influence. If immortality is understood this way, even Jesus of Nazareth can claim it on the basis of continuing discussion of New Testament themes (as well as the history behind the year numbers in the Common Era calendar), independently of all that religious claptrap about rising from the grave. We don't need to get religious to gain this kind of immortality.

Our immortality in the Globorg cloud continues as long as the cloud hosts stay up and running. As long as the data keeps being replicated, to make backups and indexes and to recompile old files whose format

has become obsolete, the extended life goes on. As long as the power keeps humming in the servers, as long as the sunlight keeps kicking up electrons on the solar farms, as long as the Earth keeps turning and the Sun keeps shining, so long can we hope to live on as movers and shakers in the web. So long, too, can we hope that someone, somewhere will take pity on our data trail and reincarnate us as living, breathing avatars, ready to take up the story where we left off eons previously.

All this remains a long way short of eternal life. Yet eternal life is always here and now, not even a moment away. It's right at the point where memory meets the "I am" core loop. We get it for free, with or without Globorg. But a long and prosperous life, that takes guts and hard work. We get much more of that with Globorg. And that's the real value proposition.

1111
Dominion

Globorg will be here as soon as the G2 becomes the GO. But at first it'll be in name only. We can expect a few decades to elapse before the deeper integration we've been reviewing here comes to pass. Political integration, restriction of national sovereignty, currency and fiscal union, legal and regulatory standardization, and more will all need to be globalized. The web will need to be built out so far that each and every citizen on Earth is online, trading work for credit, consuming services, and generally fitting in. By then, we'll be harvesting the fruits of picotech and New Life. The gap between tech and life will be gone and Globorg will regulate not only our lives but our genes as well. Then we'll be parts of a single life form – the global organism, GO.

When we've traveled that far, the political life of GO will be so far beyond capitalism versus communism or central committees versus democracy that we'll need new concepts. The set of concepts I shall now propose deserves a name, and that name shall begin with "D" (for reasons we'll come to in due course), therefore it shall be dominion. The dominion I shall now outline has nothing to do with the legal form under which Australia, Canada, New Zealand, and South Africa were once subsumed in the British Empire. And it has nothing to do with the religious doctrine of Christ's dominion. It's just a generic word for a domain of minions. In the GO dominion, the domain is the Earth and the minions are we humans.

Globorg will use a market to allocate goods and services

The managed capitalism that many of us experience as citizens is based on private ownership of capital. Ownership may be distributed in various ways, for example via limited-liability shares, but this just makes more people private owners. Public ownership of the means of production was the central feature of socialism, and most of us concede that it failed. With any asset, the buck must stop somewhere. Someone must be held accountable for the asset, either to enjoy the fruits of its success or to endure its failure. The public was the

collective entity in whose name the socialist burocrats counted beans, but it lacked a focal point from which to measure success and failure. That lack of focus was fatal. Without sharp management of assets, socialists systems devolved into bungling and corruption.

Globorg will have to do better. We've seen how it will use a market mechanism to allocate goods and services priced in globodollars. Various people will own various parts of the mechanism and will be constrained by an online regulatory environment to manage their assets in compliance with such rules as seem to serve the public good, somehow defined.

As to how, we shall police asset management much as we do now, through laws and regulations that are open to public debate in online forums. There's nothing new here except the order-of-magnitude gains in efficiency we get by hosting the asset market, the framework of laws and rules, and the public debating forums online, but that's enough to make all the difference in the world. It lets us scale up mechanisms that currently work well for corporate and urban communities to the global level.

Private ownership of goods and services is part of the human condition. We own our avatars, our bodies, our cars, our domiciles, and our estates, through expanding circles of influence that spread like ripples from the core oscillations of our boot loops. Collectivizing goods and services would merely liquidate our personalities and reduce us to communist zombies. We can hardly do other than own the services we provide and consume and we can hardly do better, I guess, than monitor provision and consumption through online financial accounting hosted by transparent frontends for the legal and fiscal authorities.

All this may seem like geekspeak for Big Brother on terascale doses of steroids. Nothing I can say or want to say will do much to change that impression. The danger is there and it's alarming. The best we can do, it seems to me, is to push for maximal transparency, accountability, open-source infrastructure development, proactive discussion of policy in open forums, and democratic control over the top levels of the global machine. And as a former policy geek who used to follow British and American elections like heartier mortals followed football matches, I naturally have ideas about how we can do better.

Fixes to democracy are prototypes for Globorg mechanisms

You may recall my claim that the United States operates an imperfect democracy. The imperfections in the system are so grave that the shining example of democratic praxis we see in CNN and other media is an illusion. We need to consider how we can transform the illusion into reality. The fixes would improve political praxis not only in the United States but also in the European Union and China. They're mechanisms we'll need in Globorg.

First, the democratic tradition of asking citizens at infrequent intervals to cast votes for names from a short list of candidates for various public offices is outmoded. It's a fun tradition, and following electoral contests as the votes are counted and the map fills in is at least as exciting for a policy wonk as watching a football match for more earthbound mortals, but it's a rite that bears little relation to the realities of power. In a complex and multilayered burocracy with a top layer of elected officials, changing the names at the top every few years does no more than ratify the status quo, however radical the policy flip-flops seem to be. Any person put in charge of a burocratic machine can do no more than try to get to know it and then tweak it in the desired direction. The tweaks are likely to have only marginal effects. Whoever the person is, the options are limited. Time and time again, elected officials in the United States or the European Union take up office with the zeal of true believers and lay down their office a few years later with little accomplished.

We can do better. We can start by ensuring that the popular feedback to the machine is more frequent and more detailed. Again, new technology is the enabler. We know from the electronic voting fiascos in the United States that implementing tamper-proof online voting is no easy matter, but it can be done. There are many public and corporate communities in Europe and elsewhere who vote online on various issues with no trouble at all. Providing online security is a must anyway for financial and legal transactions, so the solutions will certainly improve. Good security is expensive, true, but the costs are amortized over all secure transactions, so this is no argument against secure online politics. The technology can work as intended.

The main problem is deciding what we want. Setting up an elaborate and secure infrastructure for online voting is a preliminary

to using the opportunity it represents more intensively. The question is how. What more do our elected representative need to hear from us, the voters? And there you see the problem. Most elected officials are well aware of just how little the voters for the relevant offices know about the preferences and competences that would most truly assist in execution of the relevant duties. Most voters don't know enough now about the issues they're voting on. They won't know any more in future just because they vote online.

The next problem is related. Voters can be manipulated. A good illustration of this issue is the apparently democratic tactic favored in some authoritarian states of holding public referendums on burning issues of the day. By choosing the right question, those in power can often get the result they want, and the referendum is just a rubber stamp for their actions. Public opinion is putty in the hands of a skilled spin-doctor. If those in power don't like what they hear, a corrective propaganda campaign can work wonders. In an online democracy where voters are often asked to vote on detailed questions, we can fear that official promotion of the party line will subvert the process beyond salvation. Then it becomes window dressing and we might as well forget it.

The third big problem here is on the supply side. There's only a limited pool of potential candidates for most offices, once obvious nonstarters are weeded out. As for the issues, there's only a limited range of options to consider in a topical vote like a referendum. Then a simple vote that needs no electronic augmentation can be enough. Simple here translates to cheap, fast, robust, and transparent. For example, what better than a binary choice for president – red or blue – once every four years, on paper ballots that can be recounted if necessary and recycled when you're done?

Here we see the illusion at work. Democracy in the United States is demonstrated by the fact that the voter gets one binary choice in four years (at the national level, ignoring third-party candidates, state elections, and all the other complications). One bit in four years – that's one nanobyte per second. That's a million billion times less data than you send as a pundit putting out political propaganda through a webcam. The media circus for this one quadrennial bit is well worth the wait, but as an example of democracy in action it can only be symbolic.

I propose online voting on all the issues of the day, day after day

Here's my proposal. We set up a web infrastructure for secure online voting that enfranchises all online voters and offers them as many votes as they wish to cast, on all the issues of the day, day after day. The aggregated online results of each vote will be published online and made instantly available to the legislators, who can choose to follow them or ignore them.

In fact, we do something like this already. For example, CNN invites daily quick votes on its website for issues of current interest, with instant aggregate percentages for the yes or no or don't know answers. What I'm suggesting is that we do something like this more systematically. Registered voters will be able to log on to vote on an endless series of questions posted by various government departments. The latest results of the ongoing votes will be on display in the legislative and executive chambers so that the relevant officials are in no doubt what people think. With a free press to monitor the outcome and to shame officials who ignore public opinion too grossly, the result would surely be improved transparency in our democracies.

This proposal is only the start. It could be implemented quickly and easily. But its flaws are too obvious to ignore. First, as the informal CNN polls show, the results often do little more than advertise knee-jerk reactions that anyone could have anticipated. Second, most people have better things to do than spend all day online, reading a series of technical questions, and finding that all too often they can't even understand the question, never mind offer a qualified opinion as to the answer. Third, officials who merely react to public opinion are failing to exercise the professional judgment that should be part of their job description. They were voted into office, we naively presume, because they knew better than most people how to vote on the important questions arising for that office, and merely following the latest straw poll results would be a dereliction of duty.

The next refinement of the proposal, then, is to offer background briefings for the online voters. These would be something like press releases offering useful summaries of the main facts and suggesting the main policy options. Again, we have such briefings already, thanks to CNN and all the other more specialized online media services.

Anyone who can use Google can get up to speed on the main issues of the day with only minimal clicking around.

We need to refine the refinement a little more. The background briefings we're proposing are no use if voters don't read them or don't understand them. So we augment the voting portals by adding a pre-voting sequence of online questions with multiple-choice answers. Only voters who score well on the test pass through the prescreening to vote on the issue. The mechanism is simple and transparent. Anyone who survived a few years at school will understand it, as well as its rationale. To prevent cheating, the questions are changed frequently and both questions and answers are shuffled randomly for each voter. There's a big issue of principle to chew on here about disenfranchising people who flunk exams, but let's push it aside for now and move on. Subject to any later revision, then, this refined refinement corrects the first flaw in the proposal.

The second flaw is harder to fix. Who has time for all this? We can try to ensure that the background briefings are kept simple and the multichoice questions are easy enough to answer in seconds, but for some issues this won't be possible, and anyway it undoes the benefit of the first fix. So we'll find that very few voters are taking the time to vote on most issues.

But low turnout is no problem. For any interesting question of the day, we can be sure to find enough amateur policy wonks to get some feedback, even if only in small numbers, and that's enough to make the proposed setup better than nothing. For any emotive question, we'll also get plenty of nutcases and crackpots, but again this is no problem, since they'll tend to balance each other out. And remember that this lobbying from the voters is just optional feedback for the elected officials. They can ignore it if they prefer.

That brings us to the third apparent flaw in my proposal. Elected officials must be free to exercise their own judgment. That's why we voted them into office in the first place. If the media harass public officials who vote against public opinion as revealed in the online polls, we'll only be making life even more difficult for officials who already have hard decisions to make.

Such harassment is okay, in my humble opinion. Hard decisions are hard in part because they're unpopular. Knowing just how unpopular a

decision will be, or was, is part of the duty of office. So the third flaw isn't a bug. It's a feature.

The next iteration of refinements for my proposal is for the deeper future. We develop ways to process the raw results of the online polls so far that they all but tell the elected officials what to do. That is, we do indeed override the independent judgment of our public servants. We tell them what to do and expect them to do it. Naturally, this can only work if and when the community of active online voters is large and well informed. The entire institution of online briefing and voting must be widely understood and accepted.

The questionnaire filter is politically hazardous

The big issue of principle behind the questionnaire filter on who can vote needs to be chewed over for a while. The filter effectively limits the franchise to educated and intelligent voters, which is a politically hazardous development. Even though it makes the cut differently for different questions, so can sometimes include all voters, any cut at all here is a political hot potato.

However we fiddle with the details, the technology allows and even invites an extension of democracy along these lines. We can showcase the concept in limited settings quite easily. Commercial corporations, scientific societies, hobby clubs, and so on could all benefit soon from community participation organized along these lines. Then we simply scale up the solution. If the whole wide world goes online with the speed and in the numbers we can reasonably anticipate, we may well find that within a few decades we can implement this level of democracy in large communities, and perhaps even globally. Indeed, at the global level, with several billion potential voters, electing a Globorg president – a POPE – without the help of background briefings and questionnaire gates to prepare voters would be a risky business.

To recapitulate, this proposal is offered as the first fix of the democratic praxis we currently see in the United States and the European Union. The problem with our current praxis is that voters offer democratic feedback at a rate of a handful of bytes per lifetime. The fix would up that rate by a few orders of magnitude, depending on voter interest and the quality of the implementation. Even if interest were disappointing and the build quality uneven, the exercise would enliven

the democratic process and deflect accusations that our democracy is a sham.

But we needn't stop there. If interest is low, we raise it. If the build quality is low we raise it. In both cases, we do what commercial corporations do to sell more products. Corporate marketing departments have devoted big budgets to refining the art of packaging product information to raise interest in the products. Any political initiative that's worth pursuing is also worth packaging professionally. There's no shame in that. We certainly don't need to accept that sticking to the bleak facts formulated in burocratic jargon would be somehow better. If a statement is worth making at all, it's worth making with the most professional spin we can give it. As for quality, it's a hard fact of corporate life that more of it is almost always both possible and beneficial.

We get a sea change in the character of political institutions

The result of all these reforms in democratic praxis is a sea change in the character of political institutions. First, they become dependent on online voting infrastructure that needs to be built and maintained reliably in accordance with prevailing best practice. This may seem novel, but it's no more than we now routinely expect in the defense sector, where new weapon systems are developed and deployed in accordance with the latest best practice. So politics migrates to the cloud. Since everything else is in the cloud, that's no big deal.

The second change in the character of our political institutions is that popular involvement in them will potentially increase but in any case become more open. At present, political elites are cocooned in a clubby old-boy world of established personalities and their acolytes that reminds me of the small and tightly guarded world of celebrity authors, at least as it was until recently when print media began to buckle under the impact of new media and digital technology. At present, the main way for the masses to influence the political elites is by voting at a rate of a few bytes per lifetime. Similarly, in the pre-revolutionary world of trade books, the main way for the masses to influence the authors was by the simple binary act of buying or not buying their books. Both of those old worlds are changing.

In both cases, new technology makes the difference. New media like digital audio and video have opened up the closed literary world of book authors to include galaxies of more interesting media stars. Digital publishing now lets people outside the expensively educated elite print and market their own books. Blogs and chat forums let readers give authors feedback that can make a difference. Similarly, online voting portals with questionnaire gates of the sort I propose will open up politics.

Putting politics online offers wannabe politicians a much more natural and navigable transition from everyday life to political circles than via the traditional smoked-filled backrooms. In my proposed future, online tracking of voting, blogging, and posting will allow people with the right analytic tools and data access to verify a person's contribution to a cause far more thoroughly than a local candidate committee armed with piles of dog-eared papers ever could. In the world of Facebook and Twitter, we can leave street-corner agitation and door-to-door campaigns behind.

Healthy politics gets people involved, online

Politics will change dramatically as online life gets more pervasive and more universal. The machinery of democracy is destined either to develop in something like the ways I propose or to ossify and become a bad joke – two bytes per lifetime. Politics as most democracies worldwide still practice it is like book publishing before Amazon. Healthy politics gets people involved, and today that demands online interaction.

The question you may reasonably ask here is whether this changes the essence of democracy. What I suggest looks at first sight like a fuller and more modern realization of democracy. It doesn't sound at all, for example, like a move toward Platonic philosopher kings. But restricting voting rights to people who pass multichoice tests suggests a more radical agenda. We need to know who administers the tests and who decides what a candidate voter should know before below being allowed to vote on an issue.

Democracy as we have it now is spread widely but thinly. The principle of one vote for each adult citizen is good so far as it goes. But voters who don't have a clue have exactly the same formal democratic

power as voters who are fully briefed on all the issues and have solid reasons for their voting preferences. Our only defense for this practice is that any restrictions on voting are hard to defend against skeptics and probably open to abuse. We tend to fall back on the weak position that anyone who's really interested can become a party activist or donate money to a favored party, as if activism in local parties or payoffs to campaign managers were the only ways to move beyond the two-byte life. I say we can do better.

First, let's review the principle here. In any organization and in any organism, there has to be a feedback loop to connect sweet or bitter experience with the wellsprings of action. In a person, it's the boot loop that connects the synthetic unity of an experienced world with the sharp end of the acting subject. Everything comes together in my reasons for my action. Similarly, the action committee of a political party can summarize an exhausting few hours of passionate discussion with a three-point plan for the activists. Similarly, the political institutions of a democratic state can interpret the voting tallies in a general election as an instruction to begin the process of forming a new government. In each case, the feedback goes through a bottleneck before its fruit ripens in organized action.

Democracy is just a vote away from tyranny

To put the feedback loop principle in a wide perspective, it helps to recall the deep historical view. We can do so by briefly reviewing the four main forms of political organization known to the ancient Greeks: monarchy, oligarchy, democracy, and tyranny. This puts a conceptual frame around the question of how and why we can or should extend democracy.

Monarchy is an ancient and traditional form. The philosopher Hegel admired its incarnation in nineteenth-century Prussia. Many Britons admire its incarnation in the present-day United Kingdom. Mono-theists honor the form symbolically in their worship of a patriarch God. The feedback loop from perception to action is closed by the will of the monarch, who acts in the name of an established tradition.

Oligarchy is rule by a small group. Economists use the related term oligopoly to denote a group of companies that dominate a sector of industry. Corporations that lack a charismatic leader are effectively

oligarchies run by their executive boards. A state run by Plato's philosopher kings would be an oligarchy. The feedback loop is closed by agreement among the oligarchs.

Democracy is what we in the Western world know and love. We think we've hit the jackpot with universal suffrage and elected chambers of government, forgetting that it took our systems centuries to melt down from something more like an oligarchy in which only established and prosperous citizens had the vote to the present free-for-all. The democratic feedback loop is closed by the voting results.

Tyranny is what traditionally follows democracy once a demagogue gets into power. The demagogue takes the popular will for granted and acts like a monarch, but without the grace inherited through the monarchic traditions. In the best cases, the tyrant survives long enough to establish a new monarchy. In the worst case, the tyrant goes down in flames and a military junta (an oligarchy) takes over. Since tyrants rule through arbitrary will, the feedback loop is their own boot loop.

The take-away message of this history lesson is that democracy is just a vote away from tyranny. This is the danger that commentators see in the Islamic belt, where democratic states with Islamist parties run the risk that Islamists, once voted into power, will immediately establish a theocratic dictatorship. As the former U.S. diplomat Edward Djerejian famously put it: "One man, one vote, one time."

To defend democracy, we need to ensure that the voting procedure that closes the feedback loop still makes sense. We don't want a sham democracy with two-byte citizens serving as political cover for a stealth plutarchy (rule by the rich) and we don't want a media circus where a demagogue can burst forth on a tsunami of single-issue passion and seize the reins of power. We want responsible citizens who take a serious interest in the issues of the day to punch at a higher weight in the political league than the two-byte crowd.

What I propose is that we take a measured step forward in the world of online services and establish voting portals protected by questionnaire gates. Naturally, the Q-gates threaten to become a political hot potato in this proposal, but we have to bite down on it and take the pain. The alternative is the looming threat of *Ein Volk, Ein Reich, Ein Führer*.

We take our cue from educational software

So how do we do it? In short, we take our cue from educational software. In a good science program, you only get to enjoy the advanced topics after you've passed muster in boot camp. Each session ends with a few questions to check your understanding of what you just learned. If you get a passing grade, you can move on. If not, you go back and try again. The experts who write the programs and the questions know full well what the student needs to know and the student knows enough to trust the experts.

That's where the analogy with political Q-gates begins to wobble. The political questions that most need the tiebreaker of a democratic vote are questions where experts are nonexistent. Sure, people will pop up and say they're experts, but precisely those people may be the ones to steer clear of in any drafting committee for Q-gate content provision.

But I think we can overdo this line of objection. We all know what sort of people can be trusted to draft the relatively technical background briefs and related comprehension questions for issues where a popular vote can be helpful to the powers that be. Civil services are full of them. The burocracy of the European Union is full of them. The danger is only that the stuff they draft will be dry and incomprehensible. We need the active and continuous engagement of outsiders in the process. In my view, the outsiders best suited for the role are journalists and media pundits. They understand how to communicate with members of the general public, they have a feel for the popular issues of the day, and they need a job in the new world.

We have a plan of action to get started

So there we have it. We have a plan of action to get started, to move our democracy out of its thralldom to the fetish of one vote per citizen per season to a more activist and dynamic democracy where people are supported to drill as deep into the establishment as they wish, so long as they submit to the discipline of Q-gates. We'll get a more informed citizenry, a more responsive body politic, and more flexible capability to meet new challenges. More to the point of my exercise in

this book, we'll get an institutional infrastructure that can serve as a basis for the higher functions of Globorg.

In the longer term, what we'll see is increasing automation of the voting portal and its Q-gates. With good frontends, people will find it natural to vote on a few initiatives every day before breakfast. The global statistics for the big issues will be as interesting to follow as the Wall Street Big Board or the other stock indexes or football or baseball statistics are for some people today. People will get more involved because they'll feel their own influence on things. People might even rack up their own voting scores as eagerly as they now rack up friends on Facebook or admirers on Twitter. Once avatar software works as advertised, people will delegate the task of slogging through the Q-gates and preparing a voting recommendation to their avatars. Then the avatars will take over the job entirely. Then the entire political process will have migrated to the cloud and we'll have our dominion.

With all that as buildup, then, the Globorg dominion is the state where the political life of our species is fully reflected in the cloud. We live full lives, we buy and sell services, we keep up with the news and vote whenever the mood takes us, and we track the global stats for all the issues that bug us. All the while, the engines of Globorg are there in the background, tracking us, supporting us, leading us back to the straight and narrow path of virtue if we stray too far, and generally making organized life tolerable for human apes in a machine world that's slowly learning to cope without our input.

We shall learn to repeat, "I'm Globorg and Globorg is me"

How will we react to all this? I suggest that we'll react much as we do now to the infrastructure all around us. We'll take it for granted. Globorg won't be a Big Brother to us. It'll mean far more than that. We won't be viable without it. Like electricity now, Globorg will be an essential part of what makes civilized life bearable. The global network of cloud hosts will feel like our extended body just as the cloud will feel like our extended mind. Each of us will learn to repeat the mantra, "I'm Globorg and Globorg is me."

In time, the mantra will pervade our entire psychology. In our best and most exalted moments, when our spirit feels most godlike, we'll touch the face of Globorg and feel at one with the great mover and

shaker of life on Earth. At those moments, we'll realize that Globorg is the greatest incarnation of our hopes for transcendence that planet Earth can offer.

At last we come to the "D" in dominion. The secular dominion erected in the name of Globorg may be called the GO Dominion and hence be given the acronym GOD. This will suit many former rabbis, priests, or imams who seek new ways to proselytize for traditional monotheist psychodynamics. We can rely on them to describe planet Earth as the temple of GOD and prescribe faith in GOD and generally do all they can to promote the continuity of the cult of GOD with the earlier faiths in the God of Abraham, Moses, Jesus, and Mohammed. In fact, we can encourage them to do so.

It may seem odd for a secular philosopher to encourage religious faith. Modern humanist critiques of religion make a devastating case against the crass absurdities and idiocies of naïve faith. Yet rationalist attacks on religion have failed for centuries to make much of a dent in the ranks of the faithful. Even in the most secular social orders, such as Nazi Germany or Soviet Russia, people of faith hung on and flourished anew after those orders had perished. Any anthropologist not blinded by doctrine will agree that religions must be tapping into something deep in otherwise healthy human beings. My quite pragmatic take here is to accept that if people have faith anyway, we can at least help them to focus their faith on GOD.

There may be a personal wellspring for this pragmatism. Secular humanism as we know it is focused on individuals and invites selfishness bordering on social suicide. The philosophy of the *me* generation turned to dust as it dried out. People who were encouraged to live in stripped-down core loops learned to travel light. Even kids were too heavy for the more radical feminists of my youth, who shied away from settling down and partied on until their beauty faded. Men and women who were too posh to push remained childless. My own reaction was to develop the ideas in this book instead of replicating. Pragmatically, encouraging a new cult of GOD can help people to keep breeding in a world that needs the extra neurons.

For an autobiologist, in my new kind of science, the real truth runs deeper. The global self beyond our personal selves is an outgrowth of the specific forms that hovered as self-centered attractors above the evolution of species by natural selection in the history of life on Earth.

All life forms grow toward an ideal form. In each case, the form is a structure in an abstract space of possible forms (like the recursive "Buddhabrot" images I mentioned earlier, sitting in the nirvana-like spaces of fractal geometry). In each case, to use the jargon of the biologists, the genotype hovers somewhere beyond the phenotype. Our genetic potential drives us to dream beyond our stunted incarnations in the world of stubborn facts. In that gap between dream and reality, our ancestors revelated their gods.

In historical fact, the divine revelation that attracted the Biblical patriarch Abraham some four thousand years ago proved fertile. The story of the God of Abraham turned out to be a success story without parallel in human history. The biological reasons are clear enough for anyone familiar with the gene-driven solidarity of tribes and families. In the modern world, as I've argued, what survives of this tradition is the idea that the self-loop, "I am," is somehow beyond all its packaged forms in people or things. The packages, the people and things – the prophets and their books – carried the magic of the loop through the generations, yet it transcended them.

For all its faults, the monotheist social order created a rich and fertile heritage for our species. There's a powerful idea from biology that can help us understand this fact. Ironically, the idea stems from Richard Dawkins, whose role as one of the "four horsemen" of the New Atheist apocalypse attracted my earlier comment. In his 1982 book *The Extended Phenotype*, Dawkins drew from the fact that genes often determine animal behavior the conclusion that human social and technical works and achievements are part of our self-expression as a species. Now, years later, it seems that religion can be understood as part of the phenotype emerging from the human gene pool. To cap this emergence, through the purified monotheist concept of self, we can even see how life on Earth will grow beyond the human species.

Globorg will be the superself of life on Earth

Globorg will be the expression of the superself of all the life on Earth. Above and beyond the animal and personal selves in our world, the self of Globorg will bind us all together. With the Globorg Dominion, we make the leap from human vanity to scientific husbandry of the entire lifenet on our planet.

To round off the story, I have an idea for celebrating the religious roots of the new GOD.

Twenty years ago, I invented something called the Weatherdome. This was a dome built around the main attraction, which was a glass hemisphere about 13 meters in diameter, lit up with a billion pixels. Around the inside of the outer dome was a spiral ramp to serve as a viewing gallery. The glass hemisphere was the upper half of a giant Globall (an electronic globe, as you may recall) and it was fed continuously with live data from orbiting weather satellites to show the sunlit half of planet Earth. The image had a scale of one to one million, so each kilometer on the Earth's surface mapped to a millimeter on the glass surface.

Looking forward, I see a good location to build the Weatherdome. To underscore the importance of the Earth for the secular faith that replaces Abrahamic monotheism, we should build it in the one city on Earth that symbolizes all three monotheistic faiths, namely Jerusalem. The proposal may seem eccentric in the current political climate, but once the secular dominion of Globorg is well established it will seem no odder than any other architectural proposal.

Above the main entrance to the Weatherdome will be written:

GOD IS GREAT

Notes

The notes at the back of a book go mostly unread. So they're a waste of paper. They belong online. To read them, go to the book's web page:

www.andyross.net/gig.htm

References

For offline readers, the following short list of references is a minimal placeholder for the pedant's paradise I hope to build up online.

Adams, D. (1979-1992), The Hitchhiker's Guide to the Galaxy. Ballantine
Arntz, W. et al. (directors, 2004): What the Bleep Do We Know. Captured Light
Attenborough, R. (director, 1982): Gandhi. Columbia
Blackmore, S. (1999): The Meme Machine. Oxford UP
Bradley, F. H. (1876): Ethical Studies. Oxford UP
Cameron, J, (director, 1984): The Terminator. Hemdale
Cameron, J. (director, 1991): Terminator 2: Judgment Day. Carolco
Cameron, J. (director, 2009): Avatar. Twentieth Century Fox
Chalmers, D. J. (1996): The Conscious Mind. Oxford UP
Churchill, W. S. (1948-1954): The Second World War. Cassell, 6 vols
Clark, A. (2008): Supersizing the Mind. Oxford UP
Clausewitz, C. von (1832): Vom Kriege. Gutenberg
Darwin, C. (1859): On the Origin of Species. London
Dawkins, R. (1976): The Selfish Gene. Oxford UP
Dawkins, R. (1982): The Extended Phenotype. Oxford UP
Dawkins, R. (2006): The God Delusion. Bantam
Dawkins, R. et al. (2007): The Four Horsemen. DVD, RDFRS
Dennett, D. C. (1991): Consciousness Explained. Allen Lane
Dennett, D. C. (1995): Darwin's Dangerous Idea. Simon & Schuster
Dennett, D. C. (2006): Breaking the Spell. Allen Lane
Descartes, R. (1637): Discourse on Method. Literature.org
Dunbar, R. (1996): The Trouble with Science. Harvard UP
Einstein, A. (1905): The Great Works. AIP Center for History of Physics
Emmerich, R. (director, 1996): Independence Day. Centropolis
Ferguson, N. (2008): The Ascent of Money. Penguin
Feynman, R. P. (1959): There's plenty of room at the bottom. APS, Caltech
Fogel, R. (2010): Why China's economy will grow to $123 trillion by 2040. Foreign Policy
Fukuyama, F. (1992): The End of History and the Last Man. Harper
Gates, W. H. (1995): The Road Ahead. Viking
Gibbon, E. (1776-1789): History of the Decline and Fall of the Roman Empire. Gutenberg, 6 vols

Greene, B. (2004): The Fabric of the Cosmos. Knopf
Hegel, G. W. F. (1807): Phänomenologie des Geistes. Gutenberg
Huntington, S. P. (1996): The Clash of Civilizations. Simon & Schuster
Huxley, A. (1932): Brave New World. Chatto & Windus
James, W. (1902): The Varieties of Religious Experience. Longman
Joyce, J. (1922): Ulysses. Penguin
Kant, I. (1781): Kritik der reinen Vernunft. Gutenberg
Keidel, A. (2008): China's Economic Rise. Carnegie
Kurzweil, R. (2005): The Singularity Is Near. Viking
Kurzweil, R.; Grossman, T. (2004): Fantastic Voyage. Plume
Kurzweil, R.; Grossman, T. (2009): Transcend. Rodale
Lewontin, R. C. (2010): Not So Natural Selection. NYRB, LVII, 9
Llinás, R. R. (2001): I of the Vortex. MIT Press
Lucas, G. (1977): Star Wars. Lucasfilm
Mao Z. (1964): Quotations from Chairman Mao. PR China
Martin, J. (2007): The Meaning of the 21st Century. Riverhead
Marx, K.; Engels, F. (1848): The Communist Manifesto. London
Metzinger, T. (2003): Being No One. MIT Press
Metzinger, T. (2009): The Ego Tunnel. Basic Books
Murphy, R. R.; Woods, D. D. (2009): Beyond Asimov: The three laws of responsible robotics. IEEE Intelligent Systems, 24(4):14-20
Newton, I. (1687): Philosophiae Naturalis Principia Mathematica. London
Nietzsche, F. (1891): Also sprach Zarathustra. Gutenberg
Noë, A. (2009): Out of Our Heads. Hill and Wang
Orwell, G. (1949): Nineteen Eighty-Four. Warburg
Plato (c. 360 BCE): The Republic. Gutenberg
Proyas, A. (director, 2004): I, Robot. Twentieth Century Fox
Roddenberry, E. W. (screenwriter, 1964): Star Trek. Desilu
Ross, J. A. (1991): The Globall Hyperatlas. The Visual Computer, 8:1-7
Ross, J. A. (2009a): Godblogs. The Ross Blog
Ross, J. A. (2009b): SAP NetWeaver BI Accelerator. Galileo
Ross, J. A. (2009c): Mindworlds. Imprint Academic
Ross, J. A. (2010): Avatars 'R' Us. The Ross Blog
Russell, B. A. W.; Whitehead, A. N. (1910-1913): Principia Mathematica. Cambridge UP, 3 vols
Scott, R. (director, 1982): Blade Runner. Ladd Company
Searle, J. R. (1980): Minds, brains and programs. Behavioral and Brain Sciences, 3:417-457
Shakespeare, W. (c. 1600): Hamlet. Oxford UP
Singer, P. W. (2009): Wired for War. Penguin
Smith, A. (1776): The Wealth of Nations. London

Spielberg, S. (director, 2001): Artificial Intelligence. Warner Brothers
Taylor, C. (1975): Hegel. Cambridge UP
Tolle, E. (1999): The Power of Now. New World Library
Tolstoy, L. (1869): War and Peace. Knopf
Verhoeven, P. (director, 1987): RoboCop. Orion
Vinge, V. (2007): Rainbows End. Tor
Wachowski, A. and L. (directors, 1999): The Matrix. Groucho II
Wittgenstein, L. (1922): Tractatus Logico-Philosophicus. Routledge
Wittgenstein, L. (1958): Philosophical Investigations. Blackwell
Wolfram, S. (2002): A New Kind of Science. Wolfram Research

Index