

COVER

EXTROPY: The Journal of Transhumanist Thought is a journal of ideas, dedicated to discussing and developing themes in the following areas:

- Transhumanism and futurist philosophy
- Life extension, immortalism and biostasis
- Smart drugs and other intelligence intensifying technologies
- Artificial intelligence (AI) and personality uploading
- Nanocomputers and nanotechnology
- Memetics (ideas as replicating agents)
- Experimental free communities in space, on the oceans, and in cyberspace
- Effective thinking, information filtering, & life management
- Self-transformative psychology
- Spontaneous orders (free markets, neural networks, evolutionary processes, etc)
- Digital economy (privacy technologies, digital money and electronic markets)
- Critical analysis of extreme environmentalism
- Probing the ultimate limits of physics
- Artificial life

EDITORIAL COMMITTEE

EDITOR:	Max More, MA, Extropy Institute
Communications:	Russell E. Whitaker, ECFP Ventures, Ltd. (UK)
Computing, Linguistics:	Simon! D. Levy, MA, Linguistics Dept., University of Connecticut; Haskins Laboratories
Cryonics & Immortalism:	Michael R. Perry, Ph.D; Director, Society for Venturism
Law, Politics:	Tom Morrow, MA, JD, Extropy Institute
Memetics:	Keith Henson, Founder, L-5 Society
Nanotechnology:	J. Storrs Hall, Rutgers University, Laboratory for Computer Science Research
Philosophy, Politics:	Max More, MA, Extropy Institute
Physics:	Prof. Gregory Benford, Physics Dept., University of California, Irvine
Science:	David Krieger, MLS, American Information Exchange
Transhumanities:	Mark Plus, MS, Editor, <i>Venturist Monthly News</i>



Extropy Institute

EXTROPY (ISSN 1057-1035) is published twice per year (quarterly from January '94) by Extropy Institute (ExI), a nonprofit educational corporation, 11860 Magnolia Avenue, Suite R, Riverside, CA 92503. Phone: (909) 688-2323. E-mail to: more@usc.edu. Copyright ©1993 Extropy Institute.

Distributed nationally by Armadillo, Los Angeles, CA; Desert Moon Periodicals, Sante Fe, NM; FinePrint, Austin, TX; IPD, Solana Beach, CA; Ubiquity, Brooklyn, NY; & Tower Magazines, W. Sacramento, CA; Internationally by Sala Communications, Amsterdam, Netherlands; Gaia Media, Basel, Switzerland; and in the UK by Counter Productions, London, UK.

Manuscripts and letters submitted for publication must be typed or printed double-spaced, and accompanied by a stamped, self-addressed envelope.

Make checks payable to "Extropy Institute."
*All payments must be in US dollars, drawn on a US bank.

SUBSCRIPTIONS (4 issues):

USA:	\$18
Canada and Mexico*:	\$22
Overseas*:	\$32 (air); \$24 (surface)
Institutions:	USA: \$40, Canada, Mexico: \$45 Overseas: \$60/\$44 (air/surface)

BACK ISSUES: See p.33 for contents
#s 1, 2, 4, 5, 6: \$4 each. #s 7, 8, 9, 10: \$5 each.

EXTROPY INSTITUTE (See p.32 for details)

Membership (Includes one year subscription to *Extropy* and to 6-12 issues of *Exponent*, ExI's newsletter), plus discounts on conferences, tapes, T-shirts, software, and books, and invitations to festivals, and local meetings, and other gatherings:

USA:	\$30
Canada:	\$40
Overseas:	\$50/\$62 (<i>Extropy</i> by surface/full air)
Sustaining:	\$100
Benefactor:	\$300
Sponsor:	\$1000

EXTROPY #11

Summer/Fall 1993 (vol.5, no.1)

Uploading Consciousness	Ralph Merkle	05
Extropian Principles 2.5	Max More	09
Traversable Wormholes: Some Implications or, Contact! A Post-Singularity Phase Change	Michael Price Illustrated by Ralph Whelan	14
A Conversation with Mark Miller, Part 2 <i>The Day the Universe Stood Still</i>	David Krieger	24
"Bunkrapt": The Abstractions that Lead to Scares about Population and Resources	Julian L. Simon	34

THE TRANSHUMAN TASTE (reviews)

<i>Theories of Everything</i>	Peter McCluskey	4 2
<i>In Our Own Image: Building an Artificial Person</i>	Derek Zahn	4 3
<i>Mirror Worlds</i>	Harry Shapiro Hawk	4 4

Extropy Institute Update	3 2
Back Issues	3 3
Extropy 5th Birthday Party	3 3
Extro 1 Conference 1994: Call for papers	4 6
Contributors	4 7

EDITORIAL

Going Quarterly: From 1988-90, *Extropy* came out quarterly, the workload — then much lighter — divided between Tom Morrow and myself. On becoming sole editor and producer, I cut the frequency to twice yearly to ensure *Extropy's* continued appearance despite my graduate work and teaching duties. Happily, this journal will return to quarterly publication starting with next issue: *Extropy* #12 will appear in January, #13 in April, and so on. This move was encouraged by my quitting teaching to work full-time for Extropy Institute (thanks to those Extropy members who pulled together to hire me), and enabled by the increasing supply of appropriate writing.

If you bought this issue of *Extropy* at a newsstand or bookstore and intend to read future issues, please consider subscribing directly. Not only will you save money and receive your copy quickly and conveniently, you will be helping *Extropy* to survive and thrive. Distributors generally take 55% of the cover price (\$4.95 from next issue), leaving us well under half once shipping costs are paid. Due to the minimal advertising in these pages, this return makes it hard to cover costs. Current subscription information can be found on the inside front cover.

In this issue: The idea of uploading one's consciousness, personality, or self, leaving behind the biological human body for an intellectually and physically superior vehicle is an aspect of the Extropian outlook gathering much attention. The recent story in the British *GQ* ("Meet the Extropians") is a case in point. Although some of us expect a more gradual process of human-machine merging, the possibility of uploading (taken as a starting point in last issue's mindstretcher by roboticist Hans Moravec), merits serious analysis. In his article, Ralph Merkle — one of today's few professional nanotechnology researchers — calculates the goals to be achieved by our technology if we are to make this vision a reality.

The Extropian Principles 2.5 substantially refines version 2.0 from a year ago. For those of you who have not seen the Principles before, you should know that this manifesto is intended to be a concise, consistent, and comprehensive presentation of the Extropian philosophy. I welcome

Change of address: Please note that we have moved since last issue. We may move again before the next issue comes out in January '94, but mail will be forwarded.

your suggestions for further refinements; future versions are inevitable.

I'm delighted to present an ambitious examination of the uses of spacetime wormholes, based on current physics research, to realize the Extropian goal of boundless expansion throughout spacetime, civilizing the universe. I'm especially pleased to introduce readers to the author, Michael Price, with whom I first worked six years ago, on the UK cryonics newsletter, *Biostasis*. Ralph Whelan, who learned Aldus Freehand to produce the article's illustrations, deserves special thanks, both for the graphics and for helpful comments on the layout of the issue.

David Krieger concludes his conversation with Mark Miller, this second half even more stimulating, disturbing, and intriguing than the first. Miller identifies five variants of the libertarian political position, focusing especially on "nanarchy" — a possible system of the future designed to minimize coercion by removing the enforcement of rights from human control. Prepare to be both horrified and thrilled.

Economist Julian L. Simon, author of numerous books on the economics of population, immigration, and resources, and an unrelenting foe of the foolish kind of environmentalism, investigates why so many politicians, are enraptured by this bunk — why are they *bunkrap*? Three reviews and two Extropian event notices complete the issue.

Upward and Outward!

Max More

Extropy #12 (available in early January '94) will likely feature:

Boundless Constellations: The Emergence of Celestial Civilization

Ocean Colonization: A Practical Analysis
Neural-Computer Interfacing

Logical Languages: Artificial Language and Posthuman Rationality

Two Questions for Extropianism

Utility Fog (nanotech), Pt. 1

Interviewer David Krieger strikes again

More reviews (inc. Kosko's *Fuzzy Thinking: The New Science of Fuzzy Logic*)

Posthuman Sexuality

EXTROPY — a measure of intelligence, information, energy, life, experience, diversity, opportunity and capacity for growth. Extropianism is the philosophy that seeks to increase extropy. The Extropian Principles are: (1) Boundless Expansion; (2) Self-Transformation; (3) Intelligent Technology; (4) Spontaneous Order; (5) Dynamic Optimism.

TRANSHUMANISM — Philosophies of life (such as Extropianism) that seek to continue and accelerate the evolution of intelligent life beyond the limitations of the human form to a posthuman condition by means of science and technology, guided by life-furthering principles and values, while rejecting religious dogma and irrationalism. [See *Extropy* #6]

Uploading

Transferring Consciousness from Brain to Computer

Ralph C. Merkle

Xerox PARC 3333 Coyote Hill Road Palo Alto, CA 94304

Your brain is a material object. The behavior of material objects is described by the laws of physics. The laws of physics can be modeled on a computer. Therefore, the behavior of your brain can be modeled on a computer. Q.E.D.

So why haven't we done it already?

Well, we'd need a fairly big computer. And we'd have to get a very detailed description of your brain. The only ways we know of getting *that* detailed a description are destructive. That means we'd have to take your brain apart. Most people most of the time object to this. Even if *you* don't object, the legal system would. Destructive analysis of someone's brain is viewed dimly by the courts. These minor objections could be circumvented by waiting until you are legally dead. At that point, the courts wouldn't object if you didn't object. And although brain function has usually (though not always) stopped by the time you're declared legally dead, the information should still be there for a while (though you'd probably lose short term memory). When we power down the system we lose volatile memory, but non-volatile memory and the circuitry are still there.

Let's assume we've solved the legal hassles, and we're preparing to analyze your brain using the new, advanced Mark 7 Neural Analysis System. We've hooked up the Mark 7 to the Intel Pentadecium. The first question we might ask is: how much memory should we buy? How many bits does it take to describe your brain?

Your brain is made of atoms. Each atom has a location in three-space that we can represent with three coordinates: X, Y, and Z. Atoms are usually a few tenths of a nanometer apart. If we could record the position of each atom to within 0.01 nanometers, we would know its position accurately enough to know what chemi-

cal it was a part of, what bonds it had formed, and so on. The brain is roughly .1 meters across, so .01 nanometers is about 1 part in 10^{10} : we need to know the position in each coordinate to within one part in ten billion. A number of this size can be represented with about 33 bits. There are three coordinates, X, Y, and Z, so the position of an atom can be represented in 99 bits. An additional few bits are needed to store the type of the atom (whether hydrogen, oxygen, carbon, etc.), bringing the total to slightly over 100 bits.

With about 100 bits per atom we could certainly describe your brain as precisely as we'd need. (Purists might object that this does not take into account the positions of the electrons. While this is technically true, it's usually not hard in biological systems to infer the electronic structure if you have the coordinates of all the nuclei. We might wish to have a little more information, e.g., Na+, OH-, etc. With this additional ionization information our knowledge of the system would be essentially complete). Examining the published plots of the number of atoms required to store a bit of information as a function of the year, we find that somewhere between 2010 and 2020 we should be able to store one bit with one atom. If one atom in your brain is described by 100 bits, and each bit occupies one atom, then the memory required to hold a digital description of your brain accurate to the last atom would occupy about 100 times the size of your brain. The brain is somewhat over one liter, so it would require a computer memory with a volume of somewhat over one hundred liters to encode the location of each and every atom in the brain in a digital format. There are some-

what over 10^{26} atoms in the brain, so our storage system needs to hold about 10^{28} bits.

For those readers who might view the feasibility of such a memory system with some doubt, recall that DNA requires roughly 16 atoms to store a bit of information (not including the water in which it floats). Your body, with 10^{10} bits per cell stored in DNA and 10^{14} cells, stores almost 10^{24} bits of information (and it's unlikely that you're an optimal memory storage device). We're assuming only a modest improvement in storage technology over DNA; and as we'll see, we don't actually need as much storage as we've computed here.

How Many Bits to Describe a Molecule

While such a feat is remarkable, it is also much more than we need. Chemists usually think of atoms in groups — called molecules. For example, water is a molecule made of three atoms: an oxygen and two hydrogens. If we describe each atom separately, we will require 100 bits per atom, or 300 bits total. If, however, we give the position of the oxygen atom and give the orientation of the molecule, we need: 99 bits for the location of the oxygen atom plus perhaps 20 bits to describe the type of molecule ("water", in this case) and perhaps another 30 bits to give the orientation of the water molecule (10 bits for each of the three rotational axes). This means we can store the description of a water molecule in only 150 bits, instead of the 300 bits required to describe the three atoms separately. (The 20 bits used to describe the type of the molecule can describe up to 1,000,000 different molecules: more than are present in the brain).

As the molecule we are describing gets larger and larger, the savings in storage gets bigger and bigger. A whole protein molecule will still require only

150 bits to describe, even though it is made of thousands of atoms. The canonical position of every atom in the molecule is specified once the type of the molecule (which occupies a mere 20 bits) is given. A large molecule might adopt many configurations, so it might at first seem that we'd require many more bits to describe it. However, biological macromolecules typically assume one favored configuration rather than a random configuration, and it is this favored configuration that we will describe.

Describing the brain one atom at a time is much less compact than describing it one molecule at a time.

Do We Really Need to Describe Each Molecule?

While this reduces our storage requirements quite a bit, we could go much further. Instead of describing molecules, we could describe entire sub-cellular organelles. It seems excessive to describe a mitochondrion by describing each and every molecule in it. It would be sufficient simply to note the location and perhaps the size of the mitochondrion, for all mitochondria perform the same function: they produce energy for the cell. While there are indeed minor differences from mitochondrion to mitochondrion, these differences don't matter much and could reasonably be neglected.

If we're concerned about the behavior of the nervous system then worrying about the location of each mitochondrion seems excessive. We could describe an entire cell with only a general description of the function it performs: this nerve cell has synaptic connections of a certain type with that other cell, it has a certain shape, and so on. If we assume there are 10^{15} synapses, and if we need (very roughly) 100 bits per synapse, this brings us down to 10^{17} bits. We could be yet more economical of storage: a group of cells in the retina might perform a 'center surround' computation, so the entire group (including all their synapses and fine morphology) could be summarized in one succinct functional description.

How Many Bits Do We Really Need?

This kind of logic can be continued, but where does it stop? What is the most compact description which captures all the essential information? While many minor details of neural structure are irrelevant, our memories clearly matter. If we

can't fully describe long term memory we've gone too far.

How many bits does it take to hold human memory? Cherniak[6] said: "On the usual assumption that the synapse is the necessary substrate of memory, supposing very roughly that (given anatomical and physiological 'noise') each synapse encodes about one binary bit of information, and a thousand synapses per neuron are available for this task: 10^{10} cortical neurons $\times 10^3$ synapses = 10^{13} bits of arbitrary information (1.25 terabytes) that could be stored in the cerebral cortex." A problem with hardware-based estimates is that they have to make assumptions about how the information is stored. The brain is highly redundant and not completely understood: the mere fact that a great mass of synapses exists does not imply that they are in fact contributing to the memory capacity. This makes the work of Landauer[7] very interesting for he has entirely avoided this hardware guessing game by measuring the actual functional capacity of human memory directly.

A Functional Estimate of Human Long Term Memory Capacity

Landauer works at Bell Communications Research — closely affiliated with Bell Labs where the modern study of information theory was begun by C. E. Shannon to analyze the information carrying capacity of telephone lines (a subject of great interest to a telephone company). Landauer naturally used these tools by viewing human memory as a novel "telephone line" that carries information from the past to the future. The capacity of this "telephone line" can be determined by measuring the information that goes in and the information that comes out, allowing the great power of modern information theory to be applied.

Landauer reviewed and quantitatively analyzed experiments by himself and others in which people were asked to read text; look at pictures; hear words, short passages of music, sentences and nonsense syllables. After delays ranging from minutes to days or longer the subjects were then tested to determine how much they had retained. The tests were quite sensitive (they did not merely ask "What do you remember?") often using true/false or multiple choice questions, in which even a vague memory of the material would increase the chances of making the correct choice. Often, the differential abilities of a group that had been exposed to the material and another group that

had not been exposed to the material were used. The difference in the scores between the two groups was used to estimate the amount actually remembered (to control for the number of correct answers an intelligent human could guess without ever having seen the material). Because experiments by many different experimenters were summarized and analyzed, the results of the analysis are fairly robust; they are insensitive to fine details or specific conditions of one or another experiment. Finally, the amount remembered was divided by the time allotted to memorization to determine the number of bits remembered per second.

The remarkable result of this work was that human beings remembered very nearly two bits per second under *all* the experimental conditions. Visual, verbal, musical, or whatever — two bits per second. Continued over a lifetime, this rate of memorization would produce somewhat over 10^9 bits, or some hundreds of megabytes.

While this estimate is probably only accurate to within an order of magnitude, Landauer says

We need answers at this level of accuracy to think about such questions as: What sort of storage and retrieval capacities will computers need to mimic human performance? What sort of physical unit should we expect to constitute the elements of information storage in the brain: molecular parts, synaptic junctions, whole cells, or cell-circuits? What kinds of coding and storage methods are reasonable to postulate for the neural support of human capabilities? In modeling or mimicking human intelligence, what size of memory and what efficiencies of use should we imagine we are copying? How much would a robot need to know to match a person?

Landauer's estimate is interesting because of its small size. While Landauer doesn't measure everything (he did not measure, for example, the bit rate in learning to ride a bicycle nor does his estimate even consider the size of "working memory") his estimate of memory capacity suggests that the capabilities of the human brain are more approachable than we had thought.

How many bits do we need to satisfactorily describe your brain? We have quite a range: from 10^{28} to 10^9 . If we assume we have to describe every neuron and every synapse (and every nerve impulse traveling along every neuron), we're probably safe in estimating something like 10^{18} bits. Those who object to this approximation can buy the more expensive High Fidelity system which keeps

If the changes that have been introduced by the uploading process are smaller than the behavioral changes introduced by (say) a beer, a night's sleep or a cup of coffee, then it's getting rather difficult to argue that uploading has somehow destroyed the real you and substituted a "fake" you that just seems (by all objective measures) to be you.

track of each and every atom. If people will buy gold-plated Monster Speaker cables.....

How Much Computing Power?

Now that we have a rough idea of the information storage we'll need, how many operations per second will we need? How fast does the brain operate? While MIPS are appropriate for a PC, there are several measures we might use for the brain. We might count the number of synapses, estimate their average speed of operation, and so determine synapse operations per second. If there are roughly 10^{15} synapses operating at about 10 impulses/second[2], we get roughly 10^{16} synapse operations per second.

A second approach is to estimate the computational power of the retina, and then multiply this estimate by the ratio of brain size to retinal size. The retina is relatively well understood so we can make a reasonable estimate of its computational power. The output of the retina — carried by the optic nerve — is primarily from retinal ganglion cells that perform "center surround" computations (or related computations of roughly similar complexity). If we assume that a typical center surround computation requires about 100 analog adds and is done about 100 times per second[3], then computation of the output of each ganglion cell requires about 10,000 analog adds per second. There are about 1,000,000 axons in the optic nerve[5, p. 21], so the retina as a whole performs

about 10^{10} analog adds per second. There are about 10^8 nerve cells in the retina[5, p. 26], and between 10^{10} and 10^{12} nerve cells in the brain[5, p. 7], so the brain is roughly 100 to 10,000 times larger than the retina. By this logic, the brain should be able to do about 10^{12} to 10^{14} operations per second (in good agreement with the estimate of Moravec, who considers this approach in more detail[4, p. 57 & 163]).

A third approach is to measure the total energy used by the brain each second, and then determine the energy used for each "basic operation". Dividing the former by the latter gives the total number of basic operations per second. We need two pieces of information: the total energy consumed by the brain each second, and the energy used by a "basic operation".

The total energy consumption of the brain is about 25 watts[2]. Much of this is used either for "house keeping" or is wasted, perhaps 10 watts is used for "useful computation".

The Energy of a Nerve Impulse

Nerve impulses are carried by either myelinated or un-myelinated axons. Myelinated axons are wrapped in a fatty insulating myelin sheath, interrupted at intervals of about 1 millimeter exposing the axon. These interruptions are called "nodes of Ranvier". Propagation of a nerve impulse in a myelinated axon is from one node of Ranvier to the next — jumping over the insulated portion.

A nerve cell has a "resting potential" — the outside of the nerve cell is 0 volts (by definition), while the inside is about -60 millivolts. When a nerve impulse passes by, the internal voltage briefly rises above 0 volts because of an inrush of Na^+ ions. The inrushing Na^+ goes through special protein pores in the nerve cell membrane called "voltage activated sodium channels". They are normally closed, but when

the nerve impulse comes by they open for about a millisecond and then spontaneously close again[2].

When a single voltage-activated sodium channel opens, it has a conductance of about 15 picosiemens [1]. (A siemen is the reciprocal of an ohm, and is also called a "mho"). In myelinated nerve cells there are roughly 60,000 channels at each node of Ranvier (and nowhere else). The total charge that crosses the membrane at one node in one millisecond can thus be computed: about 5.4×10^{-11} coulombs (over 300 million ions per node). The energy dissipated is just the charge times the voltage, or 3.2×10^{-12} joules. If we view this one millimeter jump as a "basic operation" then we can easily compute the maximum number of such "Ranvier ops" the brain can perform each second: 3.1×10^{12} .

Although the details differ for unmyelinated nerve cells, the energy cost of traveling one millimeter is about the same.

To translate "Ranvier ops" (1-millimeter jumps) into synapse operations we must know the average distance between synapses, which is not normally given in neuroscience texts. We can estimate it: a human can recognize an image in about 100 milliseconds, which can take at most 100 one-millisecond synapse delays. A single signal probably travels 100 millimeters in that time (from the eye to the back of the brain, and then some). If it passes 100 synapses in 100 millimeters then it passes one synapse every millimeter — which means one "synapse operation" is about one "Ranvier operation".

If propagating a nerve impulse a distance of 1 millimeter requires about 3.2×10^{-12} joules and the total energy dissipated by the brain is about 10 watts, then nerve impulses in your brain can collectively travel at most 3.1×10^{12} millimeters per second. By estimating the distance between synapses we can in turn estimate how many synapse operations per second your brain can do. This estimate is three to four orders of magnitude smaller than an estimate based simply on counting synapses and multiplying by the aver-

age firing rate, and similar to an estimate based on functional estimates of retinal computational power. It seems reasonable to conclude that the human brain has a "raw" computational power towards the low end of the range between 10^{12} and 10^{16} "operations" per second.

We'll use the upper end of this range, 10^{16} operations a second.

Our Model Isn't Perfect

We have been glossing over a point: a computational model of a physical system will fail to precisely predict the behavior of that system down to the motion of the last electron for two reasons: quantum mechanics is fundamentally random in nature, and any computational model has an inherent limit to its precision. The former implies that we can at best predict the probable future course of events, not the actual future course of events. The latter is even worse — we cannot precisely predict even the probable course of future events. A good example of this second point is the weather: weather prediction more than a week or two into the future might well be inherently impossible given any error in the initial conditions or computations. Any error at all (rounding off to a mere million digits of accuracy) will eventually result in gross errors between the actual events and the events predicted by the computational model. The model predicts sunshine next Tuesday, and we get rain. This kind of error cannot be avoided.

We have been simplifying our computations even further by not bothering to compute the state of every atom, or even of every molecule. We've been operating at the level of synapses or higher, which introduces another sort of "noise" into the computation.

It's safe to conclude that any computational model of your brain will almost certainly deviate from the behavior of the original — eventually in some gross and detectable fashion. If you decide that it doesn't matter which of two courses of action to follow and allow yourself to decide on whim, then it seems plausible that some slight influence might cause a computational model of your brain to select the opposite course. But is this difference "significant?" Given that our model is highly accurate for short periods of time and that any deviations are either random or represent the accumulation of slight errors, does it matter that the behavior of the model and of the original eventually deviate in some gross and obvious fashion?

We can view this another way: your brain, as a physical system, is already subject to a variety of outside and essentially random influences caused by

(among other things): temperature fluctuations; microwaves, light, and other electromagnetic radiation; cosmic rays; last night's dinner; a beer, etc. If the errors in our computational model are smaller than these influences do we really care about the difference? Is it "significant?" The human brain can and does continue to function reasonably well in the presence of gross perturbations (the death of many neurons, for example) yet this does not detract from our consciousness or life — I don't die even if tens of thousands of neurons do. In fact, I usually don't even notice the loss. A model of your brain that described the behavior of every synapse and nerve impulse, and did a reasonably accurate job at that level, would seem to capture everything that is essential to being "you."

Yet how can we tell? How will we judge the "accuracy" of our computational model? How can we say what is "significant" and what is "insignificant?" We might adopt a variation of the Turing test: if an external tester can't tell the difference, then there is no difference. But is the opinion of an external tester enough? How about your opinion? If you "feel" a difference, wouldn't this mean that the model was a "mere copy" and not really you?

Well, we could ask: "Hi! We've uploaded your brain into an Intel Pentadecium, how are you feeling?" "Absolutely top notch!" "Do you think you're not you?" "Nope, I'm me. And this simulated body is great!" "How's the orgy?" "Wonderful! Who worked on this software? I'd like to shake their hand, they've done a really great job! Uh, I hope you don't mind, but maybe I could talk with you a bit more after the party's over? I'm being distracted....."

The ultimate in experimental evidence: try it and see!

If everyone agrees that you're you, including you, and if behavioral tests can't show any difference, then is there any difference? Perhaps, but the grounds for objection are getting rather slim. If the changes that have been introduced by the uploading process are smaller than the behavioral changes introduced by (say) a beer, a night's sleep or a cup of coffee, then it's getting rather difficult to argue that uploading has somehow destroyed the real you and substituted a "fake" you that just seems (by all objective measures) to be you.

Summary

Roughly, uploading will need a computer with a memory of about 10^{18} bits, able to do around 10^{16} "operations" a second. A computer of this capacity should fit comfortably into a cubic centimeter in the

early 21st century.

It will also require the highly accurate analysis of your nervous system. This kind of analysis should also become feasible in the 21st century. There is already considerable interest in understanding the human brain: for example, the Brain Mapping Initiative has already been started[8]. Transmission electron microscopy has been used to do complete three-dimensional reconstructions of small volumes of neural tissue and this relatively primitive approach could be scaled up to much larger volumes[9]. The use of more advanced technology should make the complete and inexpensive analysis of the human brain feasible.

The biggest obstacle to uploading today is the primitive state of current technology and the unfortunate fact that our current hardware has an MTBF (Mean Time Between Failures) of 70 years (I've already used up 41, how about you?). Even worse, actual failures occur unpredictably and the failure mode is catastrophic, resulting in complete erasure of all software. Bummer.

But if you can bridge the gap (it's only a few decades) then you've got it made. All you have to do is freeze your system state if a crash occurs and wait for the crash recovery technology to be developed. Fortunately, cryonic suspension services are available today which quite literally let you freeze your state: call Alcor at 800-367-2228. Which means if you can't stay alive and healthy until the technology is developed (and approved by the FDA, don't forget the regulatory delays!) you can be suspended until you can be uploaded.

And then you'll get to find out exactly how good that Roman Orgy simulation package really is.

References:

1. *Ionic Channels of Excitable Membranes*, by Bertil Hille, Sinauer 1984.
2. *Principles of Neural Science*, by Eric R. Kandel, James H. Schwartz and Thomas M. Jessell, 3rd ed., Elsevier 1991.
3. Tom Binford, private communication.
4. *Mind Children*, by Hans Moravec, Harvard University Press, 1988.
5. *From Neuron to Brain*, second edition (1984) by Stephen W. Kuffler, John G. Nichols, and A. Robert Martin. Sinauer.
6. "The Bounded Brain: Toward Quantitative Neuroanatomy," by Christopher Chemiak, *Journal of Cognitive Neuroscience*, Volume 2, No. 1, pages 58-68.
7. "How Much Do People Remember? Some Estimates of the Quantity of Learned Information in Long-term Memory," by Thomas K. Landauer, in *Cognitive Science* 10, 477-493, 1986
8. *Mapping the Brain and its Functions*, edited by Constance Pechura and Joseph Martin, National Academy Press 1991.
9. "Large Scale Analysis of Neural Structures" by Ralph C. Merkle, Xerox PARC Technical Report CSL-89-10, November 1989.

EXTROPIAN PRINCIPLES v.2.5

Max More President, Extropy Institute

(July 1993)

EXTROPY — A measure of intelligence, information, energy, vitality, experience, diversity, opportunity, and capacity for growth.

EXTROPIANISM — The philosophy that seeks to increase extropy.

Extropianism is a *transhumanist* philosophy: Like humanism, transhumanism values reason and humanity and sees no grounds for belief in unknowable, supernatural forces externally controlling our destiny, but goes further in urging us to push beyond the merely human stage of evolution. As physicist Freeman Dyson has said: "Humanity looks to me like a magnificent beginning but not the final word." Religions traditionally have provided a sense of meaning and purpose in life, but have also suppressed intelligence and stifled progress. The Extropian philosophy provides an inspiring and uplifting meaning and direction to our lives, while remaining flexible and firmly founded in science, reason, and the boundless search for improvement.

1. Boundless Expansion — Seeking more intelligence, wisdom, and effectiveness, an unlimited lifespan, and the removal of political, cultural, biological, and psychological limits to self-actualization and self-realization. Perpetually overcoming constraints on our progress and possibilities. Expanding into the universe and advancing without end.

2. Self-Transformation — Affirming continual moral, intellectual, and physical self-improvement, through reason and critical thinking, personal responsibility, and experimentation. Seeking biological and neurological augmentation.

3. Dynamic Optimism — Fueling dynamic action with positive expectations. Adopting a rational, action-based optimism, shunning both blind faith and stagnant pessimism.

4. Intelligent Technology — Applying science and technology creatively to transcend "natural" limits imposed by our biological heritage, culture, and environment.

5. Spontaneous Order — Supporting decentralized, voluntaristic social coordination processes. Fostering tolerance, diversity, long-term thinking, personal responsibility, and individual liberty.

These principles are developed below. Deeper treatments can be found in various issues of *EXTROPY: The Journal of Transhumanist Thought* — Spontaneous Order in #7, Dynamic Optimism in #8, and Self-Transformation in #10.

1 BOUNDLESS EXPANSION

Extropians recognize the unique place of our species, and our opportunity to advance nature's evolution to new peaks. Beginning as mindless matter, parts of nature developed in a slow evolutionary ascendance, leading to progressively more powerful brains. Chemical reactions generated tropistic behavior, which was superseded by instinctual and Skinnerian stimulus-response behavior, and then by conscious learning and experimentation. With the advent of the conceptual awareness of humankind, the rate of advancement sharply accelerated as intelligence, technology, and the scientific method were applied to our condition. We seek to sustain and quicken this evolutionary process of expanding extropy, transcending biological and psychological limits into posthumanity.

In aspiring to posthumanity, we reject natural and traditional limitations on our possibilities. We champion the rational use of science and technology to eradicate constraints on lifespan, intelligence, personal vitality, freedom, and experience. We recognize the absurdity of meekly accepting "natural" limits to our lifespans. The future will bring a graduation from Earth — the cradle of human and transhuman intelligence — and the inhabitation of the cosmos.

Resource limits are not immutable. Extropians affirm a rational, market-mediated environmentalism aimed at sustaining and enhancing the conditions for our flourishing. We oppose apocalyptic environmentalism which hallucinates catastrophe, issues a stream of irresponsible doomsday predictions, and attempts to strangle our continued evolution. Intelligent management of resources and environment will be fostered by the Extropian goal of vastly extended lifespan. The market price system encourages conservation, substitution, and innovation, preventing any need for a brake on growth and progress. Migration into space will immensely enlarge the energy and resources accessible to our civilization. Ex-

tended lifespans will foster wisdom and foresight, while restraining recklessness and profligacy.

We seek to sustain and quicken this evolutionary process of expanding extropy, transcending biological and psychological limits into posthumanity.

No mysteries are sacrosanct, no limits unquestionable; the unknown will yield to the ingenious mind. We seek to understand the universe and to master reality up to and beyond any currently foreseeable limits.

2 SELF-TRANSFORMATION

Extropians affirm reason, critical inquiry, intellectual independence, and honesty. We reject blind faith and the passive, comfortable thinking that leads to dogma, mysticism, and conformity. Our commitment to positive self-transformation requires us to critically analyze our current beliefs, behaviors, and strategies. Extropians therefore feel proud by readily learning from error rather than by professing infallibility. We prefer analytical thought to fuzzy but comfortable delusion, empiricism to mysticism, and independent evaluation to conformity. We affirm a philosophy of life but distance ourselves from religious dogma because of its blind faith, debasement of human worth, and systematic irrationality.

We seek to become better than we are, while affirming our current worth. Perpetual self-improvement — physical, intellectual, psychological, and ethical — requires us to continually re-examine our lives. Self-esteem in the present cannot mean self-satisfaction, since a probing

mind can always envisage a superior self in the future. Extropians are committed to deepening their wisdom, honing their rationality, and augmenting their physical and intellectual capabilities. We choose challenge over comfort, innovation over emulation, transformation over torpor.

Extropians are neophiles and experimentalists who track new research for more efficient means of achieving goals and who are willing to explore novel technologies of self-transformation. In our quest to advance to a posthuman stage, we rely on our own judgment, seek our own path, and reject both blind conformity and mindless rebellion. Extropians frequently diverge from the mainstream because they refuse to be chained by any dogma, whether religious, political, or intellectual. Extropians choose their values and behavior reflectively, standing firm when required but responding flexibly to new conditions.

Personal responsibility and autonomy go hand-in-hand with self-experimentation. Extropians take responsibility for the consequences of their choices, refusing to blame others for the results of their own free actions. Experimentation and self-transformation require risks; we wish to be free to evaluate potential risks and benefits for ourselves, applying our own judgment, and assuming responsibility for the outcome. We seek neither to rule others nor to be ruled. We vigorously resist those who use the institutionalized coercion of the State to impose their judgments of the safety and effectiveness of various means of self-experimentation. Personal responsibility and self-determination are incompatible with authoritarian centralized control, which stifles the choices and spontaneous ordering of autonomous persons.

Coercion, whether for the purported "good of the whole" or for the paternalistic protection of the individual, is unacceptable to us. Compulsion breeds ignorance and weakens the connection between personal choice and personal outcome, thereby destroying personal responsibility. Extropians are rational individualists, living by their own judgment, making reflective, informed choices, profiting from both success and shortcoming.

As neophiles, Extropians study advanced, emerging, and future technologies for their self-transformative potential. We support biomedical research to understand and control the aging process. We examine any plausible means of conquering death, including interim measures like biostasis, and long-term possibilities such as migration of personality from biological bodies into superior embodiments

("uploading").

We practice and plan for biological and neurological augmentation through means such as neurochemical enhancers, computers and electronic networks, General Semantics, fuzzy logic, and other guides to effective thinking, meditation and visualization techniques, accelerated learning strategies, applied cognitive psychology, and soon neural-computer integration. Shrugging off the limits imposed on us by our natural heritage, we apply the evolutionary gift of our rational, empirical intelligence to surpass the confines of our humanity, crossing the threshold into the transhuman and posthuman stages that await us.

3 DYNAMIC OPTIMISM

Extropians espouse a positive, dynamic, empowering attitude. Seeing no rational support for belief in a non-physical "afterlife", we seek to realize our ideals in *this* world. Rather than enduring an unfulfilling life sustained by a desperate longing for an illusory heaven, we direct our energies enthusiastically into moving toward our ever-evolving vision.

Living vigorously, effectively, and joyfully, requires dismissing gloom, defeatism, and ingrained cultural negativism. Problems — technical, social, psychological, ecological — are to be acknowledged but not allowed to dominate our thinking and our direction. We respond to gloom and defeatism by exploring and exploiting new possibilities. Extropians hold an optimistic view of the future, foreseeing potent antidotes to many ancient human ailments, requiring only that we take charge and *create* that future. Dynamic optimism disallows passively waiting and wishing for tomorrow; it propels us exuberantly into immediate activity, confidently confronting today's challenges while generating more potent solutions for our future.

We question limits others take for granted. Observing accelerating scientific and technical learning, ascending standards of living, and evolving social and moral practices, we project continuing progress. Today there are more researchers studying aging, medicine, computers, biotechnology, nanotechnology, and other enabling disciplines than in all of history. Technological and social development continue to accelerate leading, in the eyes of some of us, to a Singularity — a time in the future when everything will be so radically different from today, and chang-

ing so fast, that we cannot accurately foresee life beyond that horizon. Extropians strive to maintain the pace of progress by encouraging support for crucial research, and pioneering the implementation of its results.

Where others see difficulties, we see challenges. Where others give up, we move forward.

Adopting dynamic optimism means focusing on possibilities and opportunities, being alert to solutions and potentialities. It means refusing to whine about what cannot be avoided, learning from mistakes rather than dwelling on them in a victimizing, punishing manner. Dynamic optimism requires us to take the initiative, to jump up and plough into our difficulties, our actions declaring that we *can* achieve our goals, rather than sitting back and submerging ourselves in defeatist thinking.

Our actions and words radiate dynamic optimism, inspiring others to excel. We are responsible for taking the initiative in spreading this invigorating optimism; sustaining and strengthening our own dynamism is more easily achieved in a mutually reinforcing environment. We stimulate optimism in others by communicating our Extropian ideas and by living our ideals.

Dynamic optimism and passive faith are incompatible. Faith in a better future is confidence that an external force, whether God, State, or extraterrestrials, will solve our problems. Faith, or the Pollyanna/Dr. Pangloss variety of optimism, breeds passivity by promising progress as a gift bestowed on us by superior forces. But, in return for the gift, faith requires a fixed belief in and supplication to external forces, thereby creating dogmatic beliefs and irrationally rigid behavior. Dynamic optimism fosters initiative and intelligence, assuring us that we are capable of improving life through our own efforts. Opportunities and possibilities are everywhere, calling to us to seize them and to build upon them. Attaining our goals requires only that we believe in ourselves, work diligently, and be willing to revise our strategies.

Where others see difficulties, we see challenges. Where others give up, we move forward. Where others say *enough is enough*, we say: *Forward! Upward! Out-*

ward! We espouse personal, social, and technological evolution into ever higher forms. Extropians see too far and change too rapidly to feel future shock. Let us advance the wave of evolutionary progress.

4 INTELLIGENT TECHNOLOGY

Extropians affirm the necessity and desirability of science and technology. We use practical methods to advance our goals of expanded intelligence, superior physical abilities, self-constitution, and immortality, rather than joining the well-trodden path of comfortable self-delusion, mysticism, and credulity. We regard science and technology as indispensable means to the evolution and achievement of our most noble values, ideals, and visions. We seek to foster these disciplined forms of intelligence, and to direct them toward eradicating the barriers to our Extropian objectives, radically transforming both the internal and external conditions of existence.

We will co-evolve with the products of our minds, integrating with them, finally merging with our intelligent technology in a posthuman synthesis, amplifying our abilities and extending our freedom.

Technology is a natural extension and expression of human intellect and will, of creativity, curiosity, and imagination. We foresee and encourage the development of ever more flexible, smart, responsive technology. We will co-evolve with the products of our minds, integrating with them, finally merging with our intelligent technology in a posthuman synthesis, amplifying our abilities and extending our freedom.

Profound technological innovation excites rather than frightens us. We welcome change, expanding our horizons, exploring new territory boldly and inventively. We favor careful and cautious development of powerful technologies, but will neither stifle evolutionary advancement nor cringe before the unfamiliar. Regard-

ing timidity and stagnation as unworthy of us, we choose to stride valiantly into the future. Extropians therefore favor surging ahead — delighting in future shock — rather than ignobly stagnating or reverting to primitivism. Intelligent use of biotechnology, nanotechnology, space and other technologies, in conjunction with a free market system, can remove resource constraints and discharge environmental pressures.

We are evolving away from tribalism, feudalism, authoritarianism, and democracy towards a polycentric system of distributed power shared among autonomous agents, their plans coordinated by the economic network.

We see the coming years and decades as a time of enormous changes, changes that will vastly expand our opportunities and abilities, transforming our lives for the better. This technological transformation will be accelerated by genetic engineering, life extending biosciences, intelligence intensifiers, smarter interfaces to swifter computers, neural-computer integration, virtual reality, enormous and interconnected databases, swift electronic communications, artificial intelligence, neuroscience, neural networks, artificial life, off-planet migration, and nanotechnology.

5 SPONTANEOUS ORDER

Extropians emphasize self-generating, organic, spontaneous orders over centrally planned, imposed orders. Both types of order have their place, but the underappreciated spontaneous variety are crucial for our social interactions. Spontaneous orders have properties that make them especially conducive to Extropian goals and values; we see spontaneously ordering processes in many contexts, including biological evolution, the self-regulation of ecosystems, artificial life studies, memetics (the study of replicating information patterns), agoric open systems (market-like allocation of computational

resources), brain function and neurocomputation.

The principle of spontaneous order is embodied in the free market system — a system that does not yet exist in a pure form. We are evolving away from tribalism, feudalism, authoritarianism, and democracy towards a polycentric system of distributed power shared among autonomous agents, their plans coordinated by the economic network. The free market allows complex institutions to develop, encourages innovation, rewards individual initiative, cultivates personal responsibility, fosters diversity, and decentralizes power. Market economies spur the technological and social progress essential to the Extropian philosophy. We have no use for the technocratic idea of central control by self-proclaimed experts. No group of experts can understand and control the endless complexity of an economy and society. Expert knowledge is best harnessed and transmitted through the superbly efficient mediation of the free market's price signals — signals that embody more information than any person or organization could ever gather.

Sustained progress and effective, rational decision-making require the diverse sources of information and differing perspectives that evolve in spontaneous orders. Centralized command of behavior constrains exploration, diversity, and dissenting opinion. Respecting spontaneous order means supporting voluntaristic, autonomy-maximizing institutions as opposed to rigidly hierarchical, authoritarian groupings with their bureaucratic structure, suppression of innovation and dissent, and smothering of individual incentives. Our understanding of spontaneous orders grounds our opposition to self-proclaimed and involuntarily imposed "authorities", and makes us skeptical of political solutions, unquestioning obedience to leaders, and inflexible hierarchies.

Making effective use of a spontaneously ordering social system requires a degree of tolerance and self-restraint, allowing others to pursue their lives as they choose, just as we wish to be free to go our own way. Mutual progress and fulfillment will result from a cooperative and benevolent attitude towards all those who respect our rights. Tolerating diversity and disagreement requires us to maintain control of the impulses built into the human organism, and to uphold demanding standards of rational personal behavior. Extropians are guided in their actions by studying the fields of strategy, decision theory, game theory, and ethology. These reveal to us the benefits of coopera-

tion, and encourage the long-term thinking appropriate to persons seeking an unlimited lifespan.

We who have become transhuman will be primed to transform ourselves into posthumans — persons of unprecedented physical, intellectual, and psychological capacity, self-programming, potentially immortal, unlimited individuals.

CONCLUSION

These are principles not only of belief but of *action*. We become transhuman only when we have fully integrated these values into our lives, when we have consciously transformed ourselves ready for the future, rising above outmoded human beliefs and behaviors. When technology allows us to reconstitute ourselves physiologically, genetically, and neurologically, we who have become transhuman will be primed to transform ourselves into posthumans — persons of unprecedented physical, intellectual, and psychological capacity, self-programming, potentially immortal, unlimited individuals.

As posthumans we will both embody extropy and generate more — more intelligence, information, energy, vitality, experience, diversity, opportunity, and growth. As we progress from human to transhuman to posthuman, our understanding and application of these Principles will evolve with us. The Extropian Principles are a new operating system for our selves; always seeking to improve upon them, we will avoid dogmatizing them. The Principles derive their value by guiding us to our true goal: the maximization in our lives of **extropy**.

BEST DO IT SO!

Boundless
Expansion
Self-
Transformation
Dynamic
Optimism
Intelligent
Technology
Spontaneous
Order

READINGS

These books are listed because they express Extropian ideas. However, appearance on this list should not be taken to imply full agreement of a book or its author with the Extropian principles, or vice versa. Reading just the first ten books listed will illuminate many components of the evolving Extropian worldview.

Paul M. Churchland:	<i>Matter and Consciousness</i>
Richard Dawkins:	<i>The Selfish Gene</i>
Eric Drexler:	<i>Engines of Creation</i>
David Friedman:	<i>The Machinery of Freedom (2nd Ed.)</i>
Hans Moravec:	<i>Mind Children: The Future of Robot and Human Intelligence</i>
Ed Regis:	<i>Great Mambo Chicken and the Transhuman Condition</i>
Julian L. Simon:	<i>The Ultimate Resource</i>
Robert Anton Wilson:	<i>Prometheus Rising</i>
Ayn Rand:	<i>Atlas Shrugged</i> (fiction)
Marc Stiegler:	<i>The Gentle Seduction</i> (fiction)
Harry Browne:	<i>How I Found Freedom in An Unfree World</i>
Paul M. Churchland:	<i>A Neurocomputational Perspective</i>
Stephen R. Covey:	<i>The 7 Habits of Highly Effective People</i>
Mike Darwin & Brian Wowk:	<i>Cryonics: Reaching For Tomorrow</i>
Ward Dean & John Morgenthaler:	<i>Smart Drugs and Nutrients</i>
Daniel C. Dennett:	<i>Elbow Room: The Varieties of Free Will Worth Wanting</i>
Eric Drexler:	<i>Nanosystems: Molecular Machinery, Manufacturing, and Computation</i>
Eric Drexler, C. Peterson with Gayle Pergamit:	<i>Unbounding the Future: The Nanotechnology Revolution</i>
Freeman Dyson:	<i>Infinite in All Directions</i>

F.M. Esfandiary:	<i>Optimism One</i> <i>Up-Wingers</i> <i>Telespheres</i>
Robert Ettinger:	<i>The Prospect of Immortality</i> <i>Man Into Superman</i> <i>Are You A Transhuman?</i>
FM-2030:	<i>Morals By Agreement</i>
David Gauthier:	<i>The Immortalist</i>
Alan Harrington:	<i>Info-Psychology</i>
Timothy Leary:	<i>The Miracle of Theism</i>
J.L. Mackie:	<i>The Libertarian Idea</i>
Jan Narveson:	<i>A Step Farther Out</i>
Jerry Pournelle:	<i>Order Out of Chaos</i>
Ilya Prigogine and Isabelle Stengers:	<i>Markets, Entrepreneurs and Liberty</i>
W. Duncan Reekie:	<i>Prolongevity II</i>
Albert Rosenfeld:	<i>The Resourceful Earth</i>
Julian Simon and Herman Kahn (eds):	<i>Powershift</i>
Alvin Toffler:	<i>The New Inquisition</i>
Robert Anton Wilson:	

Fiction:

Roger MacBride Allen:	<i>The Modular Man</i>
Greg Egan:	<i>Quarantine</i>
Robert Heinlein:	<i>Methusalem's Children</i> <i>Time Enough for Love</i> <i>Voyage To Yesteryear</i> <i>Inherit the Stars</i>
James P. Hogan:	<i>The Silicon Man</i>
Charles Platt:	<i>The Great Explosion</i>
Eric Frank Russell:	<i>Illuminatus! (3 vols.)</i>
Robert Shea and Robert Anton Wilson:	<i>The Probability Broach</i>
L. Neil Smith:	<i>Schismatrix</i>
Bruce Sterling:	<i>True Names</i>
Vernor Vinge:	"The Ungoverned" in <i>Across Realtime</i>

ACKNOWLEDGMENT

My thanks to all those who have commented on the numerous drafts of the revised Principles, especially Jamie Dinkelacker, Derek Ryan, and Ralph Whelan.

COPYRIGHT POLICY

The Extropian Principles 2.5 may be reproduced in any publication, private or public, physical or electronic, without need for further authorization, so long as they appear unedited, in their entirety and with this notice. Notification of publication or distribution would be appreciated. The Extropian Principles 2.5 are copyright ©1993 by Max More, c/o Extropy Institute, 11860 Magnolia Avenue, Suite R, Riverside, CA 92503.

TRAVERSABLE WORMHOLES SOME IMPLICATIONS *or* CONTACT! A POST-SINGULARITY PHASE CHANGE

Michael Clive Price

© June 1993

Illustrations by Ralph Whelan

That is often the way it is in physics — our mistake is not that we take our theories too seriously, but that we do not take them seriously enough.

— Steven Weinberg

Everything will be accomplished that does not violate known fundamental laws of science.

— Gerald Feinberg

You must follow me carefully. I shall have to controvert one or two ideas that are almost universally accepted.

— Opening words of the Time Traveler, from *The Time Machine*, by H.G. Wells.

Summary: Traversable wormholes permit faster-than-light travel, within general relativity, but not time travel and associated acausal paradoxes. This article explores some of the implications traversable wormholes have on the expansion of civilizations through the universe. In particular it is found each civilization, or empire, imposes a local, accessible, region of simultaneity, or empire-time, which differs from the more natural timeframe cosmologists use. Distant regions of the universe, and alien civilizations if they exist, can be reached in short periods of empire-time. Expanding empire-time zones fuse, on contact with each other, forming an absolute, but artificial, universal time frame. Finally, some information-processing limitations of Euclidean space are contrasted with wormhole connected non-Euclidean space.

0. INTRODUCTION

To establish an interstellar trading civilization we need faster-than-light (FTL) travel or communication, which the recently proposed traversable wormholes provide. This article is a “what-if”, and, in the words of Weinberg, takes the idea and its implications seriously. In the spirit of Feinberg I assume that the ultimate limits of technology are best suggested by the laws of physics [1].

The article is structured thus:

0. INTRODUCTION: You’re reading it.
1. SLOWER THAN LIGHT: Problems and frustrations of living in universe without faster than light travel, exacerbated by the adoption of nanotechnology.
2. FASTER THAN LIGHT: Other proposals for breaking the light barrier.
3. TRAVERSABLE WORMHOLES: The latest candidate for FTL, and some of its properties.
4. EXPLORING THE UNIVERSE: How to explore the universe with a modified

Bussard ramscoop and on-board wormhole.

5. TIME TRAVEL: Why traversable wormholes do not permit time travel, but allow FTL, and remain compatible with relativity.

6. EMPIRE-TIME: The differences between the local, or empire, time frame an expanding civilization imposes on its surroundings and the more conventional conception of time.

7. ALIENS: Contacting aliens. In particular it examines how local empire-time zones fuse together, forming...

8. UNIVERSAL TIME: ... a universal simultaneity, creating a post-Singularity cosmological phase change, Contact.

9. BEYOND THE OBSERVABLE UNIVERSE: Implications of exploring beyond the edge of the observable universe.

10. BASEMENT UNIVERSES: Some pros and cons of Euclidean space against wormhole-linked arrays of basement universes.

11. CONCLUSION

12. ACKNOWLEDGMENTS

13. REFERENCES

14. FURTHER READING

1. SLOWER THAN LIGHT


We can colonize the universe at sub-light velocities [2, 3], but the colonies remain separated from each other by the vastness of interstellar space. In the past trading empires have coped with time delays on commerce routes of the order of a few years, or decades at most. This suggests that integrated, interstellar economic and cultural zones are limited, at most, to only a few star systems.

Nanotechnology [4] only exacerbates the situation. We expect full nanotech, uploading, AIs and other self-transformative technology to arrive (over a period of some few years, often dubbed the Singularity) before interstellar travel becomes practical. Assume, for illustrative purposes, that we keep the same dimensions for our brains as at the moment. Once we are uploaded onto, and redesigned on, a decent nanotech platform our mental speeds can be expected to exceed our present rates by the ratio of the speed of electrical impulses to neurochemical impulses — about a million-fold speed-up. Subjective time, in the information world Hans Moravec has called cyberspace [5], speeds up by this factor. Perhaps we can't expect an ultimately materials-based economy (which even cyberspace is, with its need for raw processing power) to speed up by this amount. Economic speed-up of a factor of a thousand, as the geometric mean of one and a million, might be more reasonable and I shall adopt this factor for illustrative purposes. Even so, the doubling time for the economy is reduced from decades to weeks. Trade across more than light weeks is much less economically significant due to the growth and change in markets during a doubling. Although individual stellar systems can form single economic zones, they remain in economic isolation from even their nearest neighbors, including their surrounding Oort cloud or cometary halo.

With full nanotech and nuclear transmutation there is little need to transfer matter. Trade in the distant future is likely to consist of mostly information. Design plans for new products, assembled on receipt. Patterns of uploaded consciousness of intrepid travelers. Gossip and news. But, with communication delays to Alpha Centauri of the order of millions of subjective years, two-way dialogues are difficult to imagine — even when we are enjoying unlimited life spans. Old news is no news.

Interstellar communication and exploration, without FTL, is a one-way process. If you had a yen to travel to the Alpha Centauri system you could. Squirt

The Einstein-Rosen Bridge connects two distant regions of space via a comparatively short spacetime tunnel. In this depiction, spacetime is represented two-dimensionally as a plane, artificially curved to aid visualization of the Bridge. Einstein-Rosen Bridges appear to be extremely short-lived, so much so that their use as passage-ways for light or matter is precluded.



your encoded engrams down an interstellar modem and decode at Alpha Centauri (assuming the receiving station hasn't shut down in the intervening millions of years of subjective cultural change and economic transformation). You could leave a copy of your consciousness behind as redundancy or if you wanted to explore both regions, but I suspect many of us will not find this completely satisfactory. The speed of light barrier would limit and cramp our style much more than it does at present.

Trade routes, we have seen, are unlikely to spread beyond single star systems, at least until after the economy has plateaued (maybe never). Information-based cultures are unlikely to spread beyond single planets before time delays cause social fragmentation. Mars, at its closest to Earth, is 4 light-minutes away. After nanotech speed-up the effective communication distance to Mars increases to several subjective-light-years. Other planets become as distant to nanotech-based societies as the stars are to us. And stars become as distant as present-day galaxies.

2. FASTER THAN LIGHT

Life, on the galactic scale, becomes incredibly dull without FTL. In science fiction a standard plot device is to invent some faster-than-light mechanism, to make stories interesting. As you might expect, there have been a number of efforts to circumvent the light speed barrier in science-fact as well as fiction.

What stops faster-than-light travel? According to relativity, as an object accelerates toward the light-speed barrier its mass increases asymptotically, slowing

its acceleration (with constant thrust). Ship time also slows down, which also reduces thrust (e.g. for a photon drive the frequency of the photon beam red-shifts, reducing apparent thrust to an off-ship, stationary observer). Both effects make the speed of light an insurmountable barrier.

Since the advent of relativity there have been a number of approaches to traveling faster than light:

1) Tachyons: Tachyons are posited FTL particles, compatible with relativity. They never cross the lightspeed barrier, which is all that relativity forbids, being superluminal from emission to absorption. Unfortunately there are serious doubts about whether they could be used for transmitting information [6]. Moreover, no tachyons have been detected, so things look bleak either way.

2) Superluminal quantum effects: Einstein-Podolsky-Rosen & quantum 'teleportation' [7]. This relies on an accompanying classical sublight signal, so no FTL. Other quantum schemes (e.g. pure EPR signaling) rely on transmitting information via the posited collapse of the wavefunction, on which no general consensus exists. Until this is settled we can't expect too much here. No quantum superluminal effect has been demonstrated in the laboratory, either.

3) Spinning black-holes: Things looked hopeful for a while that large spinning or charged black-holes might permit travel into other regions of somewhere. Later work showed that the passage of anything through a black-hole sets off a gravitational feedback process that crushes the


traveler to death. Also infalling radiation blue-shifts to infinity [8], frying the traveler, if tidal forces, which are similarly inflated (no matter how big the black-hole), don't shred her first.

4) Einstein-Rosen bridges: An Einstein-Rosen bridge connects two otherwise widely separated regions of space, with a bridge, throat or tunnel of space, whose length is independent of the conventional separation. Unfortunately the throat is very short-lived, pinching off so quickly that only tachyons (if they existed) could travel through them and get out the other end, [9]. But if you could travel faster than light you wouldn't need a wormhole — Catch-22! Einstein-Rosen bridges are non-traversable wormholes.

As each attempt has failed the conventional wisdom has strengthened that FTL travel is the 20th century's analog of the alchemist's dream of transmuting lead into gold. Or flying to the Moon. Or living forever. They seemed impossible dreams at the time....

3. TRAVERSABLE WORMHOLES

In 1985 Carl Sagan appealed to theoretical physicists for plausible methods of FTL travel to include in his forthcoming book, *Contact*. Stimulated by this request, amongst others, were Kip Thorne and his graduate students at Caltech. Instead of looking at how different forms of matter distort space they turned the problem around and asked, what states of matter are required to hold a wormhole open permanently, so no pinch off occurs? The answer is 'exotic' states — highly stressed states, with enormous tensile strengths. The tension or pressure of 'exotic' states exceeds the local energy density. We have no familiarity with substantial 'exotic' states today, but they existed under conditions of extraordinary pressure in the early universe and exists in very tenuous forms today. Carl Sagan published *Contact* in 1985 [10], incorporating the Caltech team's early work on traversable wormholes in the novel. Thorne *et al.* published



Particles entering one "end" of a wormhole are expelled from the other, in another part of the Universe. Such a voyage would be instantaneous, enabling faster-than-light transmission of matter and data. Generally, it should be more economical to send data (i.e., light) rather than matter, as in the above depiction of the transmittal of the contents of a diskette (in this case, bylaws for a new Chapter of Ex!).

(Note: This diagram takes liberties to aid visualization in that the transmitter and receiver appear "outside" the fabric of normal spacetime, which in this depiction is two-dimensional.)

their conclusions in 1988 [11], including a recommendation for students to read *Contact* as a light introduction to traversable wormholes and 'exotic' states!

In 1989 Matt Visser showed how more general traversable wormholes could be constructed [12] or, more precisely, the material requirements for wormhole stability. A Visser-style wormhole requires 'exotic' states confined to the edges of a three-dimensional volume, for example the edges of a cube. Although there is only one cube of material, it appears at two locations to the external observer. The cube links the two 'ends' of a wormhole together. The cube has no interior, but merely facilitates passage from 'one' cube to the 'other'. Each face of the cube, instead of showing the interior of the cube, opens onto the view from the corresponding face of the other cube. A traveler, passing between the edges of 'one' cube, emerges from between the edges of the 'other' cube, unaware of anything special about the journey.

The 'exotic' nature of the edge material requires negative energy density and tension/pressure. But the laws of physics do not forbid such materials. The energy density of the vacuum may be negative, as is the Casimir field generated in the empty space between two plate conductors or in the particle-creating region around a black-hole. Negative pressure fields, according to standard astrophysics, drove the expansion of the universe during its 'infla-

tionary' phase. Cosmic string has negative tension and also tries to exhibit acausal behavior. Clearly 'exotic' states are not barred by physics.

The negative energy of a wormhole has equal magnitude to the energy of a black-hole, where the wormhole throat radius equals the black-hole Schwarzschild radius. A traversable wormhole can be thought of as the negative energy counterpart to a black-hole. The energy of a traversable wormhole, like a black-hole, scales with its linear dimensions. A one meter cube entrance requires a negative mass of roughly 10^{27} kg. A Planck-scale wormhole, throat diameter of 10^{-33} m, has a negative mass of 10^8 kg.

Negative energies, though they exist in nature, have so far only been seen in association with other positive energies, yielding systems with total positive mass. The negative Casimir energies observed are confined between metal conductors whose mass gives the total system of conductor plus vacuum a positive, overall energy. Similarly the particle creating region of an event horizon is energetically dwarfed by the associated black-hole mass. Being conservative in my induction, I'll assume that the total mass of a wormhole is positive, of the same order as the negative energy, which is suggested by some other recent work [13], although only a conjecture.

Construction of 'exotic' cubes is, of

course, far, far beyond our present day engineering capabilities. I would seriously doubt the possibility of achieving such capability were it not for the self-transformative technologies mentioned earlier. With AIs and nanotech combined we expect the limits on intelligences to be governed by physics, not biology [1], [4]. Our brain's processing capacity is conventionally assessed between 10^{15} and 10^{18} bit/sec. A comparably sized nanoelectronic brain would have processing power of 10^{32} to 10^{36} bit/sec [14]. The 6 orders of magnitude absorbed by nanotech speed-up, mentioned in the opening paragraphs, still leaves 8 - 15 orders of magnitude expansion for complexity, or depth of thought, of our brains as we switch from biology to nanotechnology. So we should not blithely assume construction and manipulation of the exotic states required will long remain beyond the grasp of future, post-Singularity civilizations, populated by such super-intelligences, or cyberminds [5], unless prohibited by physical law [1]. The remainder of the article will assume the mass production of wormholes is economically achievable by future civilizations.

Leaving aside the problems of construction, let's look at the properties of wormholes. A wormhole collapses, or throat pinches off, when the amount of mass passing through its throat's vicinity approaches the same order as the amount of negative mass confined to its edges, threatening to form a black-hole. Surprisingly, the maximum rate of mass flow through a wormhole is independent of size. As the diameter of the throat expands so does the time taken to pass into and beyond the hole's Schwarzschild radius, giving a maximum rate of mass flow through the hole of $c^3/2G$, or approximately 2.10^{35} kg/s, where G is Newton's constant, c the speed of light.

Wormholes can be viewed as communication channels with enormous potential bandwidth. According to Shannon [15] and others [14], [16], information has a minimum energy of $kT \log_2$ per bit associated with it, where T is the absolute ambient temperature. The gravitational field of the hole will impose a size-dependent lower bound on the Hawking temperature of the wormhole, giving a channel capacity that scales with hole size, of 10^{52} bits/sec \times mass (in kg). This suggests it will usually be more economic to squirt the design of an object down a wormhole channel rather than the object itself. This bandwidth, or channel capacity, is the upper limit possible through a hole, but doesn't, in itself, give any clues as to how to achieve it.

Destination	Distance & Trip time at various gees		
	(light-years)	1-g	1000-g
Alpha Centauri	4.3	2.3 years	3.3 days
Center of Milky Way	30,000	11 years	6.5 days
Andromeda Galaxy	2,250,000	15 years	8 days
Nearest Alien Civilization?	100 M	19 years	9.5 days
Edge of observable universe	10,000 M	24 years	11 days
Edge of inflationary bubble?	10^{30}	70 years	28 days

Table 1: Probe Journey Times

These two properties of wormholes, fixed matter-throughput versus bandwidth scaling with mass or radius, suggest that large, cold wormholes will be used primarily for communications, rather than matter transference. Some exceptions might be that the object is unusually information-rich or can't be reduced to classical information (e.g. a quantum correlated EPR state [7]), without destroying the object. Another class of objects that will need direct physical transference, rather than being transmitted as information, are wormholes themselves. Having laboriously dragged one end of a wormhole somewhere, later wormholes are transferred via the first, to increase the connections between the two distant regions.

An object swallowed by the mouth of a wormhole leaves its electric charge, momentum and mass associated with the mouth, in an analogous manner with the no-hair theorem for black-holes. The no-hair theorem for black-holes says that a black-hole only remembers the total charge and mass (and angular momentum) of objects swallowed. Correspondingly, when an object is disgorged from a wormhole the mass and charge of the wormhole end is reduced, by the disgorged object's mass and charge. Matter and charge flows through a wormhole have to be balanced in either direction to prevent gravitational and electric flux lines being trapped and distorting the hole. To the external observer, who may not know a wormhole is involved, mass and charge appear locally conserved. Over the long term the wormhole is forced to act as a matter exchange, rather than a source or sink for matter. I'll return to this point when discussing the Bussard ramscoop idea.

4. EXPLORING THE UNIVERSE

Wormholes enable travel from one mouth to the other. To travel to distant parts of the universe one wormhole end stays at home and the other is carted away, at sublight velocities, to the destination.

To sustain high accelerations a space probe with an on-board, small, light, wormhole could be powered from base. The fuel (perhaps antimatter, in the form of super-heavy anti-particles) is uploaded through the base end of the wormhole to the on-board end of the wormhole, powering a photon drive. A corresponding mass (ballast) has to be exchanged to maintain the two-way mass balance, as I mentioned earlier. This matter has to be collected by the probe from its environment, which naturally leads to the suggestion that the probe should be a Bussard ramscoop [17], collecting ballast/fuel from interstellar gas with a magnetic 'trawl'. Half the collected matter is exchanged for antimatter via the wormhole, which is combined with the remaining matter to power the photon drive. A Bussard ramscoop gains in thrust as it reaches higher and higher relativistic speeds (the Lorentz-Fitzgerald contraction increases the density of oncoming interstellar plasma). To protect against relativistic dust impact damage, some of the extra energy and mass could be used for the construction of a heat shield (whose mass would partially off-set the gain in thrust with speed). At different velocities, different designs are optimal, so the probe would have to effect in-flight redesign.

At the relativistic speeds time dilation becomes a major factor. Time dilation reduces trip times for relativistic travelers. A probe accelerating at one-gee approaches lightspeed within a year. As it speeds up, probe time dilates more and more. I have given flight times assuming 1-gee acceleration, after the original plans [18], based on a hydrogen fusion motor. I've also included a higher 1000-gee flight time plot, based on the greater accelerations a nanotech ramscoop construction could withstand, and an antimatter drive could deliver. See Table 1 for probe or journey time to various locations (not allowing for slow-down).

The probe remains in communication with the home base, throughout the trip. As a drop point approaches, another wormhole plus deceleration rig is uploaded

through, detaching itself from the mother craft. Deceleration is quicker and less expensive than acceleration: the daughter craft brakes itself against interstellar/galactic gas, dust and magnetic fields, or even reflects the oncoming gas forwards to double the braking force. Transfer of colonists begins when deceleration is complete. The colonists transfer through the daughter hole, whilst the main probe continues its outward voyage. One of the first tasks of colonists is to secure the connections with home by increasing the local wormhole presence, transporting more wormholes from base via existing wormholes. Initial supplies, plant and machinery are transported as needed from base. Transport of manufacturing plants continues until local nanotech factories become more competitive than transport of finished product and local industries reach critical mass. After this, the wormholes become increasingly used for communications rather than materials transport.

An analogy with the cloud chamber springs to mind here. Charged particles are tracked through cloud chambers. Each particle is invisible, but its presence is revealed by the expanding wake of droplets left behind. Similarly the space probe is all but invisible, lost in the immensity of deep space. The burgeoning colonies left behind mark its passage. The colonies send out further wormhole probes. From a distance the whole affair resembles a growing 3-D snowflake, with Earth at the center. The tips of the snowflake indicate the positions of colony-probes.

Road, sea, and air routes allow the creation and operation of global markets. With the growth of transportation, once isolated economic zones are now forming more tightly integrated global trading blocs. Similarly, wormhole connections enable galactic and intergalactic economic blocs or zones to form.

5. TIME TRAVEL

As we have seen, wormholes are constrained by relativity to travel at sublight speeds, being time-dilated as normal. Clocks placed at the two mouths of a wormhole always remain in synchronization with each other [19]. If I look through one end of a wormhole and compare the near clock with the far clock, they will always agree, even if one end of the wormhole is traveling at relativistic speeds, many light-years away. We observe the two clocks keeping time with each other, yet relativity says the 'distant', traveling clock, is running slowly. How do we reconcile this? Only by concluding that the receding clock is being displaced in space *and* time [19]. A wormhole connects different regions of space and time.

If a wormhole enables someone to travel from Alpha Centauri 2996 to Sol 2993, and vice versa, then no paradox results because she can't travel back to Alpha Centauri (through conventional space, a distance of about 4.3 light-years) and arrive before she left (to cause a paradox).

Paradoxes result if a wormhole connects, say, Alpha Centauri 3000 to Sol 2993. Now a traveler can travel, through the wormhole, from Alpha Centauri 3000 to Sol 2993 and then make the return journey, through normal space within 5 years, at sublight speeds, arriving before her own departure. This is a problem because we can always time-dilate one end of a wormhole and not the other, either by placing one end in a gravitational field or transporting it with great speed. Wormholes, it would seem, can be always transformed into time machines.

As each attempt has failed, the conventional wisdom has strengthened that FTL travel is the 20th century's analog of the alchemist's dream of transmuting lead into gold. Or flying to the moon. Or living forever. They seemed impossible dreams at the time....


Problems begin when the wormhole ends move towards each other, and the time-shifted traveler is able to return, by traveling through conventional space, to visit herself before departure. If a traveler can visit an earlier part of her worldline then the possibility of acausal paradoxes is opened. This conclusion was realized soon after the first articles on traversable wormholes were published [11]. Depending on your view of the plausibility of time travel this is either, if you believe time travel possible, very exciting or, if you scoff at time travel, proof that traversable wormholes can't exist. No general consensus emerged in the pages of various physics journals as the subject was batted back and forth. Elaborate and very interesting papers [19, 20] reconciled time travel with quantum theory, whilst Hawking proposed, and gave plausibility arguments for, a Chronological Protection Conjecture, CPC, which says the Universe

Shall Not Allow Time Travel [21].

One of the time travel skeptics is Matt Visser. Early in 1993 he showed that wormholes do *not* enable time travel [22], by proposing physical mechanisms that enforce CPC. Visser showed that the mouths of a wormhole, with an induced clock difference, could not be brought close enough together (one wormhole end inside the light cone of the other) to permit causality violation. Quantum field and gravitational effects build up as the two ends of a wormhole approach each other and either collapse the wormhole or induce a mutual repulsion. Visser's work is not complete, but it seems swarms of virtual particles disrupt the region around a time machine, just before it would otherwise become operational. The virtual particle fluxes around a nearly chronologically violating region are able, via the uncertainty principle, to form closed spacelike (superluminal) loops and borrow energy off themselves, becoming more virulent than usual. As traversable wormholes approach being time machines, the energy of the virtual spacelike particle loops pinch off the throats, preventing formation of paradoxical, real closed timelike loops. This mechanism still works even if more than one pair of wormholes is involved. One end of a wormhole is excluded from the light cone of the other end, even if the light cone is transmitted via another wormhole. For the purposes of this article I'll adopt Visser's conclusion that the CPC mechanism is generic and blocks all forms of time travel via wormholes, but permits the operation of wormholes for the purpose of FTL travel.


6. EMPIRE-TIME

Wormholes do have one major trick up their sleeves. We have seen that wormholes don't permit time travel. But they do exhibit some very strange effects. Consider the journey from Earth to Andromeda of a 1-gee exploration probe (with the obligatory on-board wormhole), from the probe's perspective. At launch from Earth, in say the year 3000, the probe's view of Earth matches the view of Earth through the on-board wormhole — both show Earth 3000. After 15 years probe-time travel, at constant 1-gee acceleration, the ship reaches Andromeda. The view of Earth through the wormhole now shows Earth 3015. But the probe can calculate trip duration, using Minkowskian geometry, relative to the stationary, Earth-bound observer. This time works out to be 2,250,001 years. So the probe knows that it is 'really' year 2,253,001. We have to conclude that wormholes not only con-



The above diagram depicts the "light cone" approach to visualizing events and their spheres of causality. In the language of light cones, normal three-dimensional space is depicted with only two spatial dimensions, with time representing the third. Relativity tells us that any event (like event "E" above) can only affect items that fall within the event's light cone, which is expanding at precisely the speed of light. Therefore, if one imagines a spherical area in space of radius 1 lightyear, one is in effect imagining the sphere of causality for an event that occurred at the center of the sphere one year ago. In a light cone diagram, that sphere of radius 1 lightyear appears in "collapsed" form as a two-dimensional circle of radius 1 lightyear, like SLICE B above, which looks at a "slice" of timespace precisely 1 year after the event captured in SLICE A. Notice that an observer situated at point P cannot be affected by (or even aware of) event E until more than two years after the event occurs.

Unless a wormhole is involved. In the diagram below, the event captured by SLICE A is the launch at 1g of a large series of wormhole-bearing ships directed away from the Earth in various directions. The ships are at relativistic speeds by the 1-year point (SLICE B), and since the wormholes allow instantaneous travel but NOT time travel, anything passed through the wormhole gates at, say, 3 years after launch Earth time, emerges instantly on the ship at 3 years after launch SHIP time (SLICE D). And since ship time is lagging relative to Earth time, the object or information passed through the wormhole gate emerges MORE than 3 lightyears from Earth. In this case, then, an observer at point P with a telescope pointed at Earth could be encountered by one of the wormhole probes years prior to seeing its launch!



nect widely separated regions, but also different times, as we said earlier. In this example Earth 3015 is connected with Andromeda 2,253,001.

Using the wormhole, a traveler can move between Earth 3015 and Andromeda 2,253,001. (Note that CPC prevents anyone trying to create a paradox. Creating an *additional* return wormhole connecting Andromeda 2,253,016 with Earth 4,503,002, say, would enable someone from Earth 4,503,002, to travel to Earth 3030, via Andromeda 2,253,016 to disrupt her own past. But the closed spacelike loops form, via the CPC mechanism, and block the arrangement.) Whilst a worm-

hole bridgehead is established, CPC prevents any connections to different times, within the future light cone, even indirectly via other wormhole connections. Because of this strict chronological enforcement it makes sense to define a local time, which I call empire-time, for use within the regions linked up. In this example, Earth time is the standard by which clocks can be defined.

The time frame being defined by the expansion of wormholes, which I've dubbed empire-time, is not coincident with the cosmological time frame. The cosmological spacetime is the spacetime frame in which the average background

distribution of matter is stationary. The cosmological frame, or co-moving frame, expands with the Hubble expansion of the universe. At each point in cosmological time the averaged distribution of matter is even, allowing the easiest calculation of dynamics of the expansion of the universe. Relativity says all reference frames are relative, but in truth most astronomers think of the cosmological frame as a natural choice, or 'Schelling frame', to adopt, even though we are drifting with respect to it.

Wormholes sent to Andromeda, in our example, at near light speeds, arrive in approximately year 2,253,001 cosmological time, but in year 3,015 empire-time. Assuming that once wormhole technology is developed we expand at near light speeds, then the surface of constant empire-time forms an inverted cone in cosmological spacetime, with Earth at the lower apex. (I use the language of cones to describe what is really a sphere, but this is conventional in relativity texts, because it lends itself to greater ease of visualization — think of time forming the vertical scale and the spatial dimensions contributing to the horizontal co-ordinates. Later times form surfaces stacked on top of earlier times.) At any particular moment in empire-time the entire surface of the empire-time cone is accessible to the wormhole traveler. Traveling along the wormhole highways away from Earth takes you into the far future in cosmological time, but not in empire-time. Later empire-time zones form inverted cones, open base uppermost, stacked on top of each other.

Empire-time is the time imposed by the wormholes throughout the region they connect up. This region I'll call an empire, although no central authority is implied but is allowed. Clocks within the empire can be synchronized with each other, provided they are close to a wormhole. A traveler within the empire could always set their clock by empire-time, because the wormholes provide a common reference frame, or a background, against which to define position and velocity. Because this reference frame is common to all occupants, the empire-time defined can be used to catalog events in a time-ordered fashion. Attempts to redefine the empire-time already laid down by the wormhole structure are firmly resisted by CPC. To redefine empire-time you have to repopulate a region with holes traveling at a vastly different speed than the original colonists. The CPC mechanism says, in empire-time terminology, two holes disturb each other as they approach closer than their empire-time difference times the speed of light, e.g., two holes

with an empire-time difference of a year can't approach closer than a light-year without being both violently disrupted and destroyed [23].

Once the empire-time frame has been defined, it becomes increasingly difficult to change it. As the population and economy of a region grow, the numbers of holes increases. Once established, to change the relationship between cosmological time and empire-time requires the complete upheaval of the local economy and denizens. Economic growth breeds chronological stability.

Questions about the distant cosmological future of our universe are answered directly by travel. How quickly is the Hubble expansion slowing? Would the natural universe expand forever or recollapse? Is the universe spatially closed? Send out a probe at one-gee. From the above table we see that within a century of empire-time it is reporting back from almost inconceivable distances and futurities, answering the questions about the fate of the natural universe. If you wish you can visit the end of the universe, and come back. "Go see the end of the universe" might be a catchy travel company's jingle. (Actually this is only possible in an open universe. In a closed universe there is a limit to how far you travel before CPC prevents you.)

7. ALIENS

Enrico Fermi said "if aliens existed they would be here" [24], reflecting the increasingly common view that circumstantial evidence indicates alien civilizations are very few and far flung in the universe. The easiest way to explore and colonize the universe is to send out self-replicating space probes, as Tipler has cogently argued [2], [3], which almost any civilization will do at some stage in its evolution. Within a cosmologically short period (i.e. millions of years) we could colonize the Milky Way and the rest of the Local Group. The arrival of a colony probe at a star system precludes and supersedes local biological evolution. This hasn't happened to us, otherwise we wouldn't be here. Since life on Earth has evolved over billions of years then we can't expect (statistically speaking) to find civilizations within our local group or, perhaps, anywhere in the universe. This is the Fermi Paradox.

A statistical elaboration of this argument [25] gives grounds for believing that the nearest aliens are currently over 100 million light-years distant. For illustrative purposes I'll assume the nearest alien civilization is 100 million light-years distant. In the cosmological frame, without

wormholes, we won't make contact with them for over 100 million years, making their existence an object of theoretical speculation, which can't be resolved for millions of years.

With relativistic probes and on-board wormholes, though, we can reach alien colonized regions within decades of empire-time, no matter (almost) how far away they are, although no probe can penetrate into an alien empire. Each empire defines its own empire-time, in conflict with the empire-time of the other. A probe from Earth flying into an alien zone not only crosses alien space, but also alien empire-time zones. As it approaches the alien home world it passes increasingly into the alien empire-time future. CPC forbids such travel by destroying lone wormholes that attempt to interpenetrate each others' empires.

This opens up the possibility of different expansion scenarios.

With relativistic probes and on-board wormholes, though, we can reach alien colonized regions within decades of empire-time, no matter (almost) how far away they are, although no probe can penetrate into an alien empire. Each empire defines its own empire-time, in conflict with the empire-time of the other

A well coordinated, centrally controlled species might halt expansion at the boundary of their home galaxy (say) for a few thousand empire-years, building up numbers, armaments etc. When their technology seems to have plateaued they resume expansion, relying on technology and numbers to overwhelm aliens. Such a strategy is technology dependent. If wormholes can be booby-trapped to explode on tampering or hostile attack, such a strategy fails. Consider what happens as they invade a neighboring, occupied galaxy. At the first sign of attack the defenders destroy their wormholes in the invasion zone and retreat in a scorched earth policy. The structure of their respective empire-times operates to favor the defenders. The attackers penetrate deeply towards the galactic core and home world

within a few years of their empire-time. 'Meanwhile' the defenders retreat, abandoning rim worlds one-by-one, over a period of tens of thousands of years of *their* empire-time. Each light-year crossed and the defenders' technology and economic power advances by a year (likely to be a large gain with nanotech growth rates), whilst the invaders' technology is in relative stasis. Eventually science, technology and weight of numbers tells and the balance of attack shifts in favor of the defenders. Unless an invader overwhelmed the defenders in some surprise, sneak attack, the attack fails. Wars have to be fought on a more subtle level. Enough material here to keep military strategists busy for a while.

A more likely scenario is: Contact is signaled by our leading wormhole probes failing in the overlap of our sphere of influence with the alien empire's sphere, a kind of neutral zone. Finding each other's probes is non-trivial. It might be easier to find the colonies than the original exploration vessels. To push the analogy with a particle zipping through a cloud chamber: search for the tell-tale droplets, rather than the elusive particle. The easiest way of doing this, at the point where the relativistic wormholes are destroyed, is to send out sub-light, mildly-relativistic survey probes (with on-board wormholes), from the nearest drop points, to establish diplomatic relations. If both sides explore each other with non- or mildly-relativistic probes (relative to the cosmological frame) then their empire-times will realign themselves, over the locale of the neutral zone, although this may take years, permitting diplomatic contact and, assuming no wars, eventual exchanges of wormholes.

Empire-times merge as empires merge. Clocks in one empire are synchronized with the clocks in the other. Initially to travel from one empire to another involves wormhole travel to the neutral zone and hopping over to a nearby alien hole, before entering into the alien's wormhole network. As wormholes are exchanged, direct travel becomes possible. The wormhole networks merge as more and more direct connections open up. The spheres of colonization are now available to each other and the two empire-times merge to form a double conical structure. If the alien empire began expansion before us, in cosmological time terms, then traveling to the alien home world would take us back to an era of cosmological time prior to the present.

Given the expansion rates quoted, once the first aliens are contacted, the second, third etc., follow soon after. In addition to directly contacting alien em-

pires we'd also make contact indirectly. To begin with we'd make contact with alien empires that had not met very many other aliens — just starting out, so to speak, as we were. This would soon change. As our probes reach further and further into the distant cosmological future we contact larger and larger alien empires, who, in turn, have met more and more other aliens. The crucial point is reached when the average number of civilizations a typical civilization is in direct contact with reaches three, or thereabouts. In our 1-gee flight scenario this point is reached about 4-5 months after first contact is established, i.e., in under 20 years exploration, plus time to establish diplomatic relations. If we plot the number of aliens contacted, directly *and indirectly*, against empire-time, we get an asymptote, bounded only by the total number of alien species in the universe, at this point.

8. UNIVERSAL TIME

This is a symmetrical situation. Not only will we be meeting aliens within an historically short period, but they will be meeting us shortly after their expansions begin. Consequently, all the space-faring species of the universe will be connecting up at about the same stage in their development. This gives us all shared interests and markets in common. We might expect each civilization to go through two future phase changes. The first phase change, the Singularity, is the adoption of full-blown nanotechnology and the consequent uploading from a biological to a nanotech platform. The second phase change, which I'll call Contact, occurs when each civilization, more or less simultaneously, links up with the rest of the universe, tapping the benefits of the near-infinite economies of scale this brings.

After Contact all the local empire-times have merged to form a universal time or simultaneity surface. On a very large scale the sheet of universal time conforms with the cosmological average. On closer inspection (i.e., scales of billions of years and light-years) the universal time surface reveals conical pit-like indentations, marking where each civilization arose and stamped its own chronological footprint on the surrounding spacetime topology, before merging with their neighbors' zones. By saying that the universal time surface is indented I reveal my own cosmological time prejudices. From the vantage of point of a future cybermind, post-Contact, it is surfaces of equal, cosmological time that appear bumpy, relative to the planes of constant universal time. To them, civilization birth points appear as the *summits* of cones in

The first phase change, the Singularity, is the adoption of full-blown nanotechnology and the consequent uploading from a biological to a nanotech platform. The second phase change, which I'll call Contact, occurs when each civilization, more or less simultaneously, links up with the rest of the universe, tapping the benefits of the near-infinite economies of scale this brings.

the cosmological time surface, relative to the flat universal time surface. Universal time would be the preferred time for discussing life, history, politics etc. — everything except prehistory before Contact.

Universal time has many similarities with absolute time, as Newton conceived of it [26]. Newton viewed absolute time as deriving from God's immanence, or presence throughout the universe. The universal time frame defined by wormholes is created by the civilizations within the universe, which is a much more satisfactory state of affairs to the modern scientific paradigm. In universal time, Contact is year zero.

Roughly half the civilizations we meet are likely to have been around, in cosmological terms, hundreds of millions or even billions of years before us. Gaining access to their empire-time zones will enable our astronomers to observe the expansion of the universe in the distant past (although always further away from here in space than cosmological time). The occurrence of the first civilization in the universe is the limit before which we could not travel, in cosmological time.

9. BEYOND THE OBSERVABLE UNIVERSE

The expansion of the universe is defined by a parameter called Hubble's constant, which relates the distance of a far galaxy with its velocity of recession. Beyond a certain distance the recession velocity exceeds the speed of light. Objects beyond this are red-shifted to infinity and are unobservable. This distance defines the edge of our observable universe, an event horizon, and lies approximately (subject to experimental error) 15-30 billion light-years away. This is the limit of the astronomers' universe. What lies beyond is left to cosmology to ponder on. Cosmological theories expounded over the last decade (in particular inflationary theories) indicate that the observable universe is just an infinitesimal speck in a greater

post-inflationary bubble that extends over distances of 10^{30} light-years or more, looking pretty much everywhere as it does here.

Inflationary theories differ about what lies beyond the inflationary bubble. Because these regions are inflating at huge rates, an event horizon prevents any substantial exploration outside the 'bubble'. Unless we make Contact we will never directly observe this since these regions will have changed greatly in the century or two of empire-time ($>10^{30}$ years of cosmological time) it takes to reach them. One possibility is that naturally occurring wormholes, relics of the inflationary period, and inflated to astronomical dimensions [27], may link our post-inflationary bubble with others, forming an infinitely large chaotic, fractal structure [28, 29].

A couple of paragraphs back I mentioned the phase change, Contact, associated with linking up with the rest of the universe and gaining the benefits of near-infinite economies of scale, access to huge information markets, etc. The present scope of Internet, the electronic global communications network, pales into utter insignificance before the size of the pan-universal internet that will form, post-Contact. It's worthwhile stopping for a moment and considering what this might do to our perception of ourselves and our place in the universe. At the moment we are the only civilization we know, unique and conceited. If civilizations lie scattered at distances of 100 million light-years, in a universe of radius 10^{30} light-years, this still yields over 10^{60} alien mother cultures. It is unlikely anyone could ever catalog all the civilizations and cultures, even if they did have a nanoelectronic brain! No single historian could encompass the sweep of history, no biologist catalog the species. We would have returned to the medieval world, surrounded by legends of distant lands populated by mythical and fantastic creatures. Construction of a single universal map and travel guide would be impossible. The culture shock of absorbing all the extra data would likely keep us

occupied for not far short of eternity.

10. BASEMENT UNIVERSES

Initially, no doubt, wormhole connections would supplement existing architectures, connecting together points in the existing locally Euclidean universe. The next logical step would be to start constructing extensions to the existing topology. The technologies involved in generating artificial inflation to expand the interiors of wormholes into basement, or baby, universes are of the same order of magnitude as creating traversable wormholes. A basement universe is a traversable wormhole with only one end and an inflated interior (rather than two ends and no interior). Rather like the Tardis, in concept, bigger on the inside than the outside. Computer-simulated basement universe formation has already been discussed in the literature [30, 31, 32]. The technology to construct traversable wormholes implies the ability to construct basement universes.

We have already mentioned that we expect speed-up of subjective time rates of a million or so with the adoption of full nanotech. If just a factor of a thousand translates into GDP and population growth rates, then doubling times drop from decades to weeks. I don't know if these growth rates are sustainable, even in empire-time, but they indicate that any limited resource is likely to be at a premium, within years of empire-time. Since the amount of natural space per civilization is likely to be limited to roughly 10^{24} cubic light-years, space will ultimately be at a premium. The need for living space dictates that eventually wormholes will be used to provide links to artificial basement universes. Or perhaps the possibility of wormhole wars, mentioned earlier, will tempt societies to move wholesale into basement universes for security.

In a sense exponential growth and Euclidean space are natural enemies. The volume enclosed by a Euclidean 3-sphere only increases with the cube of the radius. With exponential growth pressures driving expansion all civilizations confined to Euclidean space will rapidly (in historical terms) hit technological limitations or each other. Wormholes and associated basement universes offer the long-term prospect of escaping from this dilemma. An array of basement universes connected by wormholes has the useful property that the volume of habitable space accessible grows exponentially with distance from origin. A civilization driven by volumetric exponential growth need only grow radially at a constant rate through base-

ment universe space, unlike in Euclidean space, where it must expand radially exponentially.

This might seem somewhat like a subtle and obtuse piece of mathematics, but it's just restating that a tree with continually branching twigs eventually strangles itself, in Euclidean space, whereas it could grow forever through a tangled array of wormholes and basement universes, without the crowding out effect choking off growth.

A related limitation of Euclidean space is the amount of information a volume can contain. This limitation, the Bekenstein bound [33], [34], implies that to achieve unlimited information storage a system must spread itself increasingly thinly and operate more slowly [35], in the limit to zero, or else collapse into a black-hole. No such limitation applies to a space of connected basement universes. Each basement universe is shielded from the positive energy contribution of its neighbors, allowing infinitely complex, extended, networked structures to form.

11. CONCLUSION

We have seen that, whilst the construction of wormholes is technically very difficult, the long-term payoffs are very great. A civilization can expand through the universe, stamping its own chronology on its locality, at a speed only limited by its energy resources. At the very least, problems of construction, theoretical and practical, will exercise the advanced intelligences of the future considerably. In the longer term the possibility of opened-ended, perhaps even infinite, information processing lie before the civilizations which solve the problem of wormhole construction and transport. Without wormholes a civilization faces certain fragmentation as it expands. With wormholes it can remain integrated. Whether this centralizing power is used for good or ill is another question.

From a more detached point of view it is interesting that the universal time frame permits a return to the Newtonian conception of an absolute time and simultaneity, previously thought to be incompatible with general relativity. It is especially pleasing that the shape of the universal time surface is a function of the birth place-times of civilizations, rather than divine choice or blind, insensate cosmological processes.

12. ACKNOWLEDGMENTS

My thanks to all the Extropians and Cryonauts for their feedback, including Gregory Benford, Andrew Clifford, Ray

Cromwell, Dani Eder, Carl Feynman and Timothy Freeman for filtering out some of my worst errors and most obtuse wording. But most particular thanks to Robin Hanson, for jointly starting and working with me on this project, and specifically for pointing out how empire-time follows from time-dilation. Thanks in advance to Ralph Whelan for the diagrams. Needless to say none of the above share any responsibility for, or necessarily agree with, some of my conclusions, nor any of my errors.

13. REFERENCES

- [1] Gerald Feinberg. Physics and Life Prolongation. *Physics Today* 19(11) 45 (1966). The full quote is: "A good approximation for such predictions is to assume that everything will be accomplished that does not violate known fundamental laws of science as well as many things that do violate these laws."
- [2] Frank J. Tipler. Extraterrestrial Intelligent Beings Do Not Exist. *Quarterly Journal of the Royal Astronomical Society* 21(3) 267 (1980). Additional Remarks on Extraterrestrial Intelligence. *Quarterly Journal of the Royal Astronomical Society* 22(3) 279 (1981). See also Chapter 9 [3]
- [3] Frank J. Tipler and John D. Barrow. *The Anthropic Cosmological Principle*, Clarendon Press, Oxford (1986) ISBN 0198519494 Provides lots of evidence against the notion that the universe is crawling with aliens, Chapter 9.
- [4] K. Eric Drexler. *Engines of Creation*, Anchor Press/Doubleday, (1986) ISBN 0385199724. *Nanosystems: Molecular Machinery, Manufacturing, and Computation*, Wiley-Interscience, (1992) ISBN 0-471-57518-6. See also *Extropy* #10 4(2) 44 (1993) for a review by J. Storrs Hall.
- [5] Hans Moravec. Pigs in Cyberspace, *Extropy* #10 4(2) 5 (1993).
- [6] Gerald Feinberg. Possibility of Faster-Than-Light Particles. *Physical Review* 159(5) 1089 (1967) Tachyons can always be re-interpreted as moving forwards in time, with positive energy, by all observers.
- [7] Charles Bennett. Teleporting an Unknown Quantum State via Dual Classical and Einstein-Podolsky-Rosen Channels. *Physical Review Letters* 70, 1895 (1993) How to use a split EPR experiment to make a perfect quantum copy of a system — but the original is destroyed during classical signal assembly.
- [8] Werner Israel & Eric Poisson. Internal Structure of Black-Holes. *Physical Review D* 41(6) 1796 (1990) Infinite blue-shifting of infalling radiation (including gravitational) at the internal Cauchy horizon forms an infinite, hidden energy barrier (called blue-sheets) to travelers.
- [9] Charles W. Misner, Kip S. Thorne & John A. Wheeler. *Gravitation* W.H. Freeman (1973) ISBN 0716703440. Simply the best book on general relativity. See section 31.6 for Einstein-Rosen bridges.
- [10] Carl Sagan. *Contact, Arrow and Legend*, (1985) ISBN 0099469502 The novel that, amongst other things, features a fictional account of travel through a traversable wormhole. Quite good story, so I'm told.
- [11] Michael S. Morris & Kip S. Thorne. Worm-

holes in Spacetime and Their Use for Interstellar Travel. *American Journal of Physics* 56(5), 395 (1988) The article that started all the fuss. Good historical overview of the status of various energy conjectures. Outlines the requirements for traversable wormholes: violation of the averaged weak energy condition and the presence of exotic states.

[12] Matt Visser. Traversable Wormholes: Some Simple Examples. *Physical Review D* 39(10), 3182, (1989) Non-spherically symmetric traversable wormholes are shown to be superior to spherically symmetric ones.

[13] Ian H. Redmount & Wai-Mo Suen. Is Quantum Spacetime Foam Unstable? *Physical Review D* 47(6), 2163, 15-March-1993. The authors show that we'd expect small, Planck scale, negative-energy wormholes to inflate to observable dimensions. Since we don't observe this instability, perhaps, (i) the Wheeler quantum foam, with Planck-scale wormholes popping into and out of existence, doesn't exist, or (ii) wormholes are constrained by as-yet-unknown energy conditions — for instance if they only had net positive energies.

[14] Thomas Schneider. Channel Capacity of Molecular Machines. *Journal of Theoretical Biology*, 148(1) 83 (1991) Draws the parallels between an molecular assembly and Shannon's concept of channel capacity.

[15] Claude L. Shannon. Communication in the Presence of Noise. *Proceedings of the IRE* (now the IEEE), 37, 10 (1949)

[16] Thomas Schneider. Energy Dissipation from Molecular Machines. *Journal of Theoretical Biology*, 148(1) 125 (1991) Demonstrates an alternative derivation of the thermodynamic limits of computation.

[17] Iain Nicolson *The Road to the Stars*, Westbridge Books (1978), ISBN 0-7153-7618-7 Has some nice artistic impressions of Bussard ramscoops and more references for the space-faring enthusiast.

[18] R.W. Bussard. Galactic Matter and Interstellar Flight. *Astronautica Acta* 6 179 (1960) Contains plans for a fusion-powered, 1000-tonne, 1-gee ramscoop with 10^5 - $3 \cdot 10^6$ meter radius magnetic scoops.

[19] Fernando Echeverria, John Friedman, Gunnar Klinkhammar, Michael S. Morris, & Ulvi Yurtsever. Cauchy Problem in Spacetimes with Closed Timelike Curves. *Physical Review D* 42(6) 1915 (1990) Shows how the Feynman path integral, sum-over-histories approach to quantum field theory might remove time-travel causality paradoxes, an approach pioneered by Igor Novikov. See especially figure 11 for how wormhole clock synchronization works.

[20] E.V. Mikheeva and Igor D. Novikov. Inelastic Billiard Ball in a Spacetime with a Time Machine. *Physical Review D* 47(4) 1432 15-Feb-1993 An extension of [19] which covered the elastic collisions of a billiard ball with itself, to include inelastic collisions.

[21] Stephen W. Hawking. Chronology Protection Conjecture. *Physical Review D* 46(2) 603 (1992) Provides plausibility 'proof' of CPC. The back-reaction of the gravitational metric opposes conditions which permit causality violation.

[22] Matt Visser. From Wormholes to Time Machines: Remarks on Hawking's Chronology Protection Conjecture. *Physical Review D* 47(2),

554. 15-Jan-1993 Hawking's CPC is expanded upon and fleshed out. Wormholes are shown to enable FTL travel, but not time travel. Note that this work, though impressive, is not complete.

[23] Gary W Gibbons & Stephen W Hawking. Selection Rules for Topology Change. *Communications of Mathematical Physics* 148(2) 345 (1992) Wormholes have to be created/destroyed in pairs.

[24] Enrico Fermi, quoted by Carl Sagan. *Planetary and Space Science* 11(5) 495 (1963), quoted by Tipler & Barrow, 578 [3]

[25] Dani Eder (private communication) models the occurrence time of galactic and intergalactic civilizations on a normal distribution, bell curve. The time between the first and second civilizations, in a large volume, is an appreciable fraction of the standard deviation or sigma, which is related to biological and galactic formation timescales of many billions of years. The time between the arising of the first and second civilizations in an arbitrary volume scales with the standard deviation, but is very insensitive to other factors. Assuming the galaxy could support millions or billions of civilizations (as is conventional, via the Drake equation), but that sigma is of the order of billions of years, gives the statistical time between the first and second civilizations as hundreds of millions of years. This may be a gross underestimate, if life is much rarer (although contact time is not much altered, in empire-time, unless we are completely alone).

[26] Sir Isaac Newton. On the Gravity and Equilibrium of Fluids (1688+) Translated in *Unpublished Papers of Isaac Newton* ed. A.R. and Marie Boas Hall (1962). Newton preferred the term Universal Ruler to describe God. The divine time-frame was the Universal Ruler's view of its creation. So central was this to Newton's theology that he equated the relativism espoused by Leibniz as tantamount to atheism.

[27] Thomas Roman. Inflating Lorentzian Wormholes. *Physical Review D* 47(4), 1370 15-Feb-1993 Speculation on inflating Planckian wormholes to macroscopic and astronomical dimensions. Also that naturally occurring wormholes from the inflationary era, may link our habitable bubble (of which the observable universe is a subset) to other bubbles.

[28] Andrei D. Linde. An Eternally Self-Reproducing Cosmos? *Scientific American*, 268(5) 10 (1993) Computer simulation of some eternal inflation models.

[29] Dalia S. Goldwirth & Tsvi Piran. Inflation — an Alternative to the Singular Big Bang. *General Relativity and Gravitation* 23(1)7 (1991) More detailed presentation of the theory of [28], based on generic, dimensional arguments. Habitable post-inflationary bubbles float within an eternally inflating, expanding sea.

[30] Matt Visser. Wormholes, Baby Universes and Causality. *Physical Review D* 41(4), 1116 (1990) Banning topological transformations preserves causality. Baby universes created by inflation, for instance, are constrained to be linked with the parent universe by a wormhole. Superseded by [22], to some extent, which introduces other mechanisms for preserving causality. The possibility of topology changes to spacetime (which the creation of wormholes require) is addressed in [36].

[31] Katherine A. Holcomb, Seok Jae Park, & Ethan T. Vishniac. Formation of a "child" uni-

verse in an inflationary cosmological model. *Physical Review D* 39(4) 1058 15-Feb-1989 Base-ment universe formation modeled by number-crunching the Einstein equations.

[32] Alan H. Guth, S.K. Blau, & E.I. Guendelman. Dynamics of False Vacuum Bubbles. *Physical Review D* 35(6) 1747 (1987) Localized inflation is modeled, leading to the creation of baby universes.

[33] Jacob D. Bekenstein & Marcelo Schiffer. Proof of the Quantum Bound on Specific Entropy for Free Fields. *Physical Review D* 39(4) 1109, (1989) Do Zero-Frequency Modes contribute to the Entropy? *Physical Review D* 42(10) 3598 (1990), an extension to the above article.

[34] Jacob D. Bekenstein. Entropy Content and Information flow in Systems with Limited Energy. *Physical Review D* 30(8) 1669, (1984). Entropy Bounds and the 2nd Law for Black-Holes. *Physical Review D* 27(10) 2262 (1983). See also *Physical Review D* 23 287 (1981). The Bekenstein upper bound for information, in a system, is $I = 4 \pi^2 E R / h c \ln 2$, where E = energy, R = Euclidean radius Note: also valid for black-holes.

[35] Freeman J. Dyson. Time Without End: Physics and Biology in an Open Universe. *Reviews of Modern Physics*, 51(3), 447 (1979). Suggests ways in which intelligences can survive indefinitely in an open, expanding universe. See also Chapter 10 [3]

[36] Gary T. Horowitz. Topology Change in Classical and Quantum-Gravity. *Classical and Quantum Gravity* 8(4) 587 (1991) How topology transformations can be naturally handled by classical general relativity. The singularities imposed by [37] need not be physically significant, since no physical quantity becomes infinite.

[37] Robert P. Geroch. Topology in General Relativity. *Journal of Mathematical Physics* 8(4) 782 (1968) Creation of a wormhole requires either CTLs or a "mild" [36] singularity, according to classical general relativity. If we exclude CTLs, by Visser's logic, then wormhole creation implies a mild, physically permissible, singularity.

14. FURTHER READING

Apart from the references given, popular accounts are available. None, yet, incorporate Visser's work on the chronology protection conjecture.

Paul Halpern. *Cosmic Wormholes: The Search for Interstellar Shortcuts*. Dutton Press (1992) ISBN 0525934774. Very good presentation of the arguments against time-travel and FTL through black-holes.

John Gribbin. *In Search of the Edge of Time*. Black Swan (1992) ISBN 0552994626. Covers a lot of the relativity background material assumed here, with useful diagrams.

Thomas Donaldson. *The Holes of Spacetime. Analog Science Fiction and Fact*, July 1993, ISSN 10592113. Covers most of the wormhole basics, but not any of the CPC material and implications — i.e., empire and universal time.

A Conversation With Mark S. Miller

by David Krieger

Part Two:

The Day the Universe Stood Still

Mark S. Miller is one of the system architects of the Xanadu project, the electronic hypertext system conceived by Ted Nelson as the future of publishing. Miller is also co-author, with K. Eric Drexler, of the *Agorics papers* (published in *The Ecology of Computation*, B. A. Huberman, ed.; New York: Elsevier-North Holland, 1988) which first presented the idea of agoric open computing systems — computer operating systems in which system resources such as memory and processing cycles are traded by programs and processes on an internal open market. I spoke with Mark in Palo Alto, California, in September, 1992. In Part One of this interview (see Extropy #10), we discussed “creole physics” and the credit theory of identity. In this, the conclusion of the interview, we talk about the five flavors of libertarianism, including “nanarchy,” and Jim Bennett’s “Reverse Polish Moon Treaty”:

Let’s talk about libertarianism. In *Literary Machines*¹, Ted Nelson refers to your complaints about there being too much gravity, and your detestation of governments. How did you become a libertarian?

Well, I started out as a socialist. I was a socialist for the reason I believe many people are socialist, which is for actually much the same reason many people are libertarians, which is a deep desire to see people being more free and, in the case of when I was a socialist, a misunderstanding of under what kind of system people would in fact be more free. A strong anti-authoritarianism, a strong sense that the current structures and institutions in society are in need of much change and reform, and that something much better is possible.

So I was a socialist for many years. I would argue with my cousin Neil. Finally he handed me a copy of *The Moon Is A Harsh Mistress*.² When I put that book down, I was certainly not convinced of the case for free markets or for libertarianism — it doesn’t even use the word *libertarianism* — but I knew that my socialism was wrong. When I put that book down I was no longer a socialist. I knew that my socialism had been naive, that there were a lot of things I hadn’t thought through, and that I needed to reassess things. It got me started thinking about all of these matters again from a fresh perspective.

I went through the sequence many libertarians go through, which is getting into libertarianism wholeheartedly by becoming a Randroid, by getting in through Objectivism. So I read all of Rand’s novels, I *inhaled* them, I loved them. This was the Truth, the Beauty, the Light — it was such an incredible experience that I did become an obsessive, obnoxious Randroid

for a while, and many of the people from the XanAMIX community remember me from those days.

The next stage that many of us go through, which I went through, is to go from there to a kind of Nozick-style first-principles axiomatic libertarianism. From there, you kind of mellow out, if you will, into more of a Hayekian, Austrian, evolutionary libertarian.

Spontaneous orders and emergent systems, rather than a rigid “These are the rules, these are the results” kind of thing.

I have an interesting observation about libertarianism, which I think will be especially interesting to Extropians. There are, as far as I can see, five internally-consistent libertarian political philosophies, and they’re on a spectrum, from least to most statist, where even the most statist I am content to call libertarian.

There is David Friedman’s anarcho-capitalism³. I specifically say “David Friedman’s” anarcho-capitalism, because no one else, and that means Rothbard, has constructed an anarcho-capitalist story that deals with all the meta-issues by actually dealing with them, as opposed to defining them away. There’s the basic “Who will watch the watchers?” problem.

Next to anarcho-capitalism you have something which I’m going to label “nanarchy” and will be getting back to. That’s sort of a non-conventional one that I regard as actually the most important.

Next to nanarchy is minarchy. There are an infinite number of gradations of minarchy. Nozick has done a very good job of defining all the different gradations there⁴, but basically, a minarchy is a government that uses force both to preserve

its monopoly on force as well as to enforce people’s rights. So, centralized enforcement mechanisms, and perhaps a centralized judicial system... and also, most dangerous of all, a centralized legislative system. And obviously you can have all sorts of gradations of minarchy that leave some of these in and some out.

Next over is Epstein’s libertarianism. I don’t think Epstein ever calls himself a libertarian, but his book *Takings*⁵ lays out a political philosophy that to me is clearly libertarian. The fascinating thing about it is that it’s a libertarianism that doesn’t get stuck on externality problems. Essentially, in Epstein’s system, the role of the state, by derivation from the takings clause, is to use force to solve externalities. Epstein does a very good job of grounding his interpretation in the history of English common law, which is where it comes from, in order to make the case that this is the proper historical understanding of what they meant, and therefore the interpretation that the Supreme Court should currently be using, which it’s not, unfortunately.

Epstein’s interpretation is that a “taking” in any action by the government by coercive force, that decreases your property in some way. It might be that it’s decreasing the amount of property you own; it might be that it’s a decrease in the set of rights that you hold with respect to that property... so zoning would be a decrease in your rights with respect to that property; it is a taking equally well as losing title to the land...

In reducing the utility of the land.

Right. And, that just compensation is to be assessed by the courts at somewhat above fair market value for the good. It has to be assessed by the courts because the fact that a taking happened meant that you don’t necessarily have a market to give you a price. Going back to the whole history of tort law, one could do some kind of adequate job of assessing such things perhaps.

The reason that the government can only use force to solve externality problems is, if the value of what it has to give in compensation is greater than the value of what it took, then it can only proceed if the net effect of using the force is a net creation of wealth, so that there’s a sur-

plus of value to be returned to the people whose property was taken. The particular externality problem which David Friedman acknowledges as the hardest problem for libertarians, in *Machinery of Freedom*, is national defense, which is easily solved under Epstein's takings system, as follows: The compensation itself doesn't have to be in money, so, in the case of national defense, there is a takings on the tax dollars, and there is compensation in national defense, with the right of anyone to challenge in court whether the value of the national defense that they got—

Was worth what was taken from them.
Exactly.

And there's one more.

Right. The fifth one is something that is called "libertarian socialism." It's proposed by David Miller — by the way, not the David Miller who's a friend of Phil Salin's; it's a different David Miller — in an article in *Critical Review*, in the special issue they had on *The Fatal Conceit*.⁶ The article was written in response to Hayek. The article really was, to me, an explication of libertarian socialism.

David Miller buys, and has really internalized, all of the Austrian arguments, all of the libertarian arguments, all of the free-market economist arguments about the ineffectiveness of centralized planning, about the value of spontaneous order, and about the ability of an unconstrained market to produce more wealth than any centralized interference with that market — that is, the inability of a centralized interference to increase the market's wealth-creating capacity. However, the part of the Austrian story he doesn't buy is Hayek's demolishing of "social-justice." The system he proposes is basically absolute free-market minarchy, plus negative income-tax transfer payments. There's exactly this one program.

He also, by the way, buys all of the libertarian criticism of how welfare bureaucracies and food stamps and all of these interventions in poor people's lives are demeaning and destructive of the lives of poor people, and that you need to get the hands of these bureaucrats out of the lives of those people. However, he doesn't buy the arguments against the transfer payments, his basic argument for the transfer payments being that a dollar is worth more to a poor person than to a rich person. So, therefore, in a sort of "greatest good to greatest number," or maximizing net total utils, that a dollar is worth more utils to a poor person than to a rich person, and therefore if you have this net transfer of wealth from rich to poor, you've increased the overall good of society. There's that one way in which poor people get additional money. They're free to spend that money on whatever the hell they want, and they get no targeted help.

Definitions of Terms

Lotus Marketplace — A product proposed by Lotus Development Corporation in 1990. Lotus Marketplace would have been a CD-ROM database with demographic information about millions of individuals and households, including estimated income. There was vocal opposition to the product from consumer and privacy advocacy groups like the Computer Professionals for Social Responsibility, and Lotus discontinued the product without ever releasing it commercially.

Program proving — Techniques that allow one to prove that a program will operate as advertised, without actually executing the program. Program proving is complicated by the **halting problem**, the proof that it is possible to construct programs for which it cannot be determined in advance whether or not they will ever finish executing.

Util — A notional unit for measuring the usefulness, or utility, of a particular outcome; employed in economics and game theory.

Zero-knowledge proofs — Cryptographic techniques that allow the members of a group to, for example, present conclusive proof that one of them has taken a particular action, without it being possible to prove which of them took the action. The canonical application of zero-knowledge proofs is the **dining cryptographers problem**, in which three notional cryptographers wish to demonstrate that one of them has paid for their dinner, without revealing which of the three it was. Zero-knowledge proofs have been generalized to provide services like anonymous remailers and "DC nets" for electronic mail, which can generate e-mail messages that cannot be traced to their author.

I think, by the way, that *that* libertarian socialism is something that libertarians should do more to promote, whether they believe in it or not, because it's a wonderful step in an argument. It's also a good system compared to all the systems we have in the world right now. It's a worthwhile thing to move to, and the reason it's so worthwhile to promote, is that in terms of the stated goals of the left-liberals, the most important stated goals are the goals which you can't easily demolish by argumentation. All the kinds of elitist intervention in people's lives, I think are easy targets, comparatively, but in terms of what those folks really care about, this really gives it to them, and lets us really get what we care about, as well. Really, it's a brilliant compromise. Whether you think of it as a compromise, or if you, like David Miller, think that it's the best solution, at least there's a brilliant compromise that deserves to be promoted.

You wanted to return to nanarchy.

The basic problem with the libertarian anarchy/minarchy debate is that both sides are trying to solve a problem that's incredibly difficult to solve, which is the old "Who will watch the watchers?" problem. Free markets are a wonderful mechanism of evolutionary interaction within a framework of rules of lack of coercion. However, the enforcement of those rules itself relies on the ability to use coercion...

so the enforcers are operating in a biological framework, not a market framework, in which force is possible, and they need to operate in that framework in order to create a setting in which people are operating as if force was not possible.

David Friedman especially does, I think, a fairly brilliant job of trying to wrap the whole thing into a circle in a way that's completely self-consistent, of trying to turn market forces in on the users of force, but the paradox is that it's only by the proper activities of those users of force that there is a market framework to turn back on them. It may be the case that anarcho-capitalism is much better at dealing with this paradox than any centralized system, but it's still not very good. In the film *The Day The Earth Stood Still*⁷ —

I see where you're going —

— it was presented to the Earth-folk, that the aliens, the people who lived in the rest of the solar system, were operating under a system of very loose non-aggression rules that operate only between planets, that were enforced by a "race of robots" — a strange choice of term, but that's what they said. The result was that the alien who came to earth was clearly not threatening the Earthfolk himself, because he had no choice in whether the robots would smash the Earth for engaging in aggression. He was simply informing them so that they would know before they stepped

into it that there was this additional constraint that simply was there, it was enforced by the robots. Neither the Earthfolk nor the aliens could do anything about it, and it was as if another law of physics had been added to the Universe. There was simply this additional constraint, which was not corruptible, not subject to new legislation, not subject to amendment, not subject to any kind of corruption, not subject to overthrow — it was as if suddenly you found yourself living in a Universe where the laws of physics had an additional constraint added.

Blue goo.

Blue goo. By the way, the historical derivation of blue goo: Drexler and I were discussing the gray goo problem, and talking about nanarchy as a possible way out. Roger Gregory, who was as aghast as the rest of us at the idea of nanarchy, and we were all fairly aghast at it — I'm for it, because I think all the other alternatives are worse, not that I think there's anything particularly good about this one —

*Sort of like what Churchill said about democracy.*⁸

Some initial wave from human civilization, from human technology, is going to explode out into the universe, expanding out into the universe at close to the speed of light. And Drexler has pointed out that the nature of that wave, of what is on the frontier of that wave, will determine the long-term nature of the universe, and will determine what universe it is that everything that follows that wave has to work with.

The system that was presented in *The Day The Earth Stood Still* is a literal implementation of the rhetorical ideal of classical liberalism. These robots were executing a program where the program embodied the laws. This was in fact the first time that it's ever been presented, as far as I know, in any work of fiction or in any other work — a presentation of "a government of laws and not men."

Applications of nanotechnology are left to the reader as an exercise?
Essentially.

I haven't seen The Day The Earth Stood Still. I'd like to ask you, is there any indication in the film how the creators of the robots were removed from control?

The creators of the robots were not removed from control. It was clearly the case, as I would advocate for our near future, that the creators of the robots created them to be autonomous, and the creators purposely denied themselves the ability to control the robots, because had they retained that ability, the system of robots would have been corruptible.

Essentially, the builders of the robots themselves needed to be somewhat incorruptible in order to not put backdoors and trapdoors and so forth into the coding that drives the robots. That was certainly not covered in the movie. The issue of us building a dispersed system of communicating nano-Gorts —

Exactly. Roger Gregory expressed his revulsion at nanarchy very elegantly when he termed it "bluegoo," and I think perhaps stated, "I prefer graygoo." I'm not sure about that last clause; you can check it out with Roger. *[I asked Roger about this at the First General Conference on Nanotechnology in November, 1992, and he said this sounds like something he would say.— dk]*

Back to the issue of corruptibility and trusting the creators of the system. The kind of thing that we will need to engineer will be extraordinarily difficult to begin with: a mutually-constraining developmental process for designing a secure mechanism and a secure software system to run on that mechanism, such that, by virtue of the nature of the social process by which it was created, and the way in which the pieces of that social process were mutually constraining, we can be confident that the system as a whole does not have any trapdoors in it.

Let's say you've got N groups cooperating. If you've arranged the social process such that to successfully insert a trapdoor would require a coordinated conspiracy among at least one half of those N groups, then by going through the appropriate procedure for picking those N groups, and keeping them out of unmonitored communication with each other, you can be confident that there are no large conspiracies.

As precedent for that, [an acquaintance of ours] was once a missile com-

mander — he was the guy in the bunker of a missile silo that you see portrayed at the beginning of *War Games*⁹, with his finger on the button.

The other fellow sits over in the next chair, and they've got the keys. If one of them goes berserk, the other's supposed to shoot him.

Exactly. I don't know how accurate that was; you can ask [him], he'll tell you. That was his position for a while. So he actually knows a lot about the procedure by which a launch decision is made, reported, and carried out, and what the interlocks are along the way. It turns out they've gamed this thing out to see what it would take to do a Dr. Strangelove¹⁰ — to have an unauthorized launch. The government did a hell of a lot of engineering of that particular system of people and machines and authorization; they did an incredible amount of engineering and gaming of the system. Nevertheless, under the best system they designed — I won't say the best they could do, but the best they did do — it was still the case that three particular people, none of whom were the President, Vice-President, or Speaker of the House, or the Chairman of the Joint Chiefs of Staff — none of whom was in any of those privileged positions, just three particular people involved in the carrying-through of the launch authorization or the reporting of it — had they formed a conspiracy, could have successfully caused a false launch. And, what they depended on — which I think was a wise decision, something we can reliably depend on — is, you pick these people from large pools of candidates that generally don't know each other. By picking them out of large pools it's very hard to have foreknowledge of who's going to be picked, so it's very hard to have pre-arranged a conspiracy.

And you rotate them rapidly so they're always working with different people.

Actually, I don't know if they do that. And, they're careful to make sure that the particular people they've picked have not been in contact. And then, once they're picked, they're very careful not to let them be in unmonitored contact. And it's incredibly hard to form a conspiracy, or to have a conspiracy pre-arranged, under those constraints.

Something else which that reminds me of is the existence of things like zero-knowledge proofs and digital cash, where, presumably, even if all of the other participants are conspiring, you can't reconstruct the connection between a given person and a given digital pseudonym. So I would assume it would be possible to create a software-engineering environment where every piece of code, in a sense, goes through a zero-knowledge proof so that even if N-1 conspirators are trying to put in a trapdoor, the Nth person is still able to check

and prevent a conspiracy of that type. That's exactly the kind of thing to look for, to try to engineer this process so that it can succeed. Finally, it's that progress on zero-knowledge proofs, and similar kinds of cryptographic technologies, as well as the progress on program proving — which is easy to disparage but it's actually been good progress — those are the things that give me hope that we could actually carry this thing out, and have, not absolute confidence in the result, but sufficiently high confidence in the result that, given that we're in this real-time situation and to *not* turn the system on has all sorts of other dangers, this system wins if we have sufficient confidence that it's clearly less dangerous to turn it on than not to turn it on. It doesn't have to be 100%.


This brings us around to — what's the name of the holiday involved in the Reverse Polish Moon Treaty?

Inheritance Day¹¹. The false dichotomy that people keep raising is they imagine that if we don't do something like nanarchy, that somehow it'll be this nice spontaneous-order free-for-all of people dispersing in all sorts of different directions, and "Let a million flowers bloom," and a diversity happening in the colonization of the universe as different people go in indifferent directions, and a continuation of diversity on into the future. If I believed that could happen without the proper seed having been arranged through a difficult coordinated social process ahead of time, I would think that's clearly the way to go, because I hate central planning, and what I'm advocating here is central planning, and central authority —

Of a non-human agency.

Of a non-human agency, creating an additional constraint which will forever be imposed on everybody. That's pretty goddamn offensive. So, the question is, what else can be bad enough to think that that alternative is the best chance we've got?

Ralph Merkle has a nice image for this issue, which is that some initial wave from human civilization, from human technology, is going to explode out into the universe, expanding out into the universe at close to the speed of light¹². And Drexler has pointed out that the nature of that wave, of what is on the frontier of that wave, will determine the long-term nature of the universe, and will determine what universe it is that everything that follows that wave has to work with. And we have no choice about whether that initial wave happens, and whether it happens soon, other than self-destruction. We can all kill ourselves, destroy civilization and prevent the wave, but then we're no better off; probably the one scenario in which we're worse off. The only choice we've got is what the wave will be.



In the absence of a coordinated nanarchy development effort, if you go with the homestead model, and in the presence of the possibility soon of self-replicating, space-faring machines that are able to arrange for their own military defense and able to use the resources that they're acquiring by spreading to engage in that defense, what results is a terrible winner-take-all race, where "take all" in this case is "take the entire Universe." You've got this race where whoever gets there first takes the entire Universe, and the rest of us are left with essentially nothing. Alternatively, if no one power gets out there first in a defensible way, you might end up with several powers getting out there first and spreading in somewhat different directions —

That's still a very extreme oligarchy.

It's a very extreme oligarchy, and it's also not necessarily stable because of the logic of military power in a system where whoever expands outward fastest or expands in the direction of more available mass-energy most effectively, gets to have more mass-energy at his disposal to beat on the other guys. There's a positive feedback in there that probably still ends up with one winner taking all.

So, what we can choose to do instead, very carefully, with incredible dangers, is to have the first wave that explodes out there into the universe be the minimum framework of enforced rules such that, within that framework of enforced rules, that kind of military instability cannot happen.

Then there's the open question — and it's important to emphasize that the question is open — of what is a decent and minimal framework of enforced rules that is sufficient to give us more confidence with those rules than without, that we will have a universe of ongoing diversity? My proposed set of rules is a division of

property according to the Reverse Polish interpretation of Inheritance Day, the Reverse Polish Moon Treaty, and an enforcement of non-aggression between the property whose title is recognized according to the Reverse Polish treaty. Beyond that, then one would like to say, "Anything goes." Beyond that, the only rules that are enforced are the laws of physics. Anything further, people can arrange to have enforced within their own property, since it's theirs, but they can't inflict those conditions by force on someone else's property, since that violates the framework. Have Extropians been adequately introduced to the Reverse Polish Moon Treaty?

I don't think so.

The person to really talk to about this is Jim Bennett.

I was there the night that you were talking about it with him and you drew your utility curve.

I drew the utility curve by asking the people in the room, including you, "Let's take the following two extreme points." One is, I don't get any of the universe beyond a little bit of the Earth and the Solar System, which is really essentially nothing of the Universe. That might be vast wealth compared to what anybody has right now, if you have nanotechnology with which to deal with those resources, but essentially I get nothing of the Universe. Let's call that one util.

At the other extreme, I get the entire Universe as my own personal playground — I own the Universe, everybody recognizes my title, or I'm able to militarily defend it, but somehow, in an actual real sense, it's all mine, and I can have my will over the entire Universe. Let's call that 100 utils.

Now, how many utils would each of you assign to getting one ten-billionth of

the universe? Typical answers are, 5%, 10%, 50% even. Let's take the lowest that I've heard, which is like 5%.

The odds of being the one person to get everything —

Right. The odds of being the one person to get everything is, let's say, one in a billion, because you're enough of an elitist to say you're above average — you're in the top billion. Let's do an expected-value calculation. If we go with the homesteading model, your expected value is essentially one plus epsilon. You're taking the 100 utils and multiplying it by the probability that you will be the one who got to homestead the universe. Contrast that with the Inheritance Day model, in which you have basically a certainty of getting one ten-billionth of the value of the universe, that gets you five utils.

I think it's clear that now, in terms of choosing which of these two regimes of assignment of property right in the currently unowned Universe, which one is in the best interest of each of us? I think that it's clearly the Inheritance Day model.

The Inheritance Day model says: the whole Universe is currently unowned; value is value to people; it's important not to do a continual re-allocation, because that only creates a tragedy-of-the-commons problem which leads to universal starvation. It'd be a tragedy to take this incredible amount of wealth which could make all of us incredibly wealthy, and instead make all of us incredibly poor by engaging in a system of continual redistribution... which is, by the way, the reason that, even though libertarian socialism is great in the short term, it would still be fatal applied to the Universe. So, if we go around advocating libertarian socialism, we need to be clear —

That it has a time limit.

It has a time limit, or it has a resource limit. Inheritance Day says, only those resources covered by the Moon Treaty, which starts at somewhere between the Earth and the Moon's orbit, and proceeds outward from the Moon. You could have a libertarian-socialist system for redistributing Earth wealth, then a one-time equal distribution of wealth-in-the-rest-of-the-Universe, never to be re-distributed. I suspect that, when it's clear how much wealth we're talking about giving to everybody now, that any motivation of compassion people have will clearly not be an issue for arguing for continual redistribution.

"But you've got a billion star systems and I only have a million." It would be difficult to make plea like that credibly in such a situation.

One very interesting additional point about Inheritance Day: Inheritance Day, once it is carried out, actually leads, in a

way that is more effective than anything else that I know of that is politically realistic, to a net redistribution of current wealth from richer to poorer, without any coercion, in a completely voluntary, non-offensive way. It lets people who are miserably poor now, benefit *now*, before moving out into the universe, from the wealth that rich people have now. It's the only non-coercive redistribution scheme that I've ever heard of that can work, and the way it works is simply by giving everyone a title to all of these resources whose future value has some net present value, and allowing them to trade. Somebody who's on the edge of starvation naturally has a very short time horizon.

So they're going to trade Some of their distant real estate to increase the probability that they will survive to reach it, by trading it to someone who's presently richer.

Right, who's presently richer and can therefore afford to have a longer-term time horizon. Therefore, speculators in the future value of these resources will actually be paying people who are currently miserably poor *money now* — to someone who's currently starving in Somalia, that could very well save their life. And I think that it's very interesting to point out that it's a completely free-market solution that dumps money into relief for starving people, in a way that actually helps them out, and that doesn't create any conflict between the rich and poor. So it really gives all sides of the political debate what they want.

You might say a little bit more about the present Moon Treaty, and why the Inheritance Day model is called the Reverse Polish Moon Treaty.

The Moon Treaty states that the moon and all extraterrestrial resources shall be considered to be the common heritage of mankind, but the treaty itself doesn't specify an implementation protocol. The reason that the L5 organization, and all freedom-loving organizations that take the future seriously, either fought or should have fought the Moon Treaty when it first came up, is that it was before the fall of Communism, in a situation where a lot of the countries voting in the U.N. thought central planning was the natural way to do things, and in which the language of the Moon Treaty was modeled on the Law of the Sea Treaty.

The Law of the Sea Treaty was a terrible precedent for a central-planning way of allocating rights and resources, in fact a non-property right. "The common heritage of all mankind" meant that no one could own any property out there, which meant that really it would be owned by the bureaucracy that got to say what would be done with it — the common heritage of a small elitist U.N. bureaucracy which got to use it as a permanent

source of power over the rest of humanity. The understanding of what would be the implementation protocol, back then, was incredibly awful.

Now that the inability of central planning to function is universally acknowledged —

Except over in Berkeley.

Except in Berkeley. The proposal is that the U.S. in fact should now be encouraged to ratify the Moon Treaty, and Jim is drafting an implementation protocol that basically says, "In these days of the universally-acknowledged failure of central planning, it's clear how to correctly interpret the common-heritage clause, which is, if it's the common heritage of all mankind, give everyone their piece." It's *ours*, right?

Where the Polish part comes in is, what does it mean to give everyone their piece? What does it mean to take the whole Universe and divide it equally? Because Poland had very much the same problem, which is, they had all of these businesses, all of these large factories, which were all run "in the name of the people." They were in fact run by a small elitist bureaucracy who ran it for their own aggrandizement, but once that bureaucracy was out of power, there was now a government interested in freedom and free markets, but faced with all this state property and trying to figure out what to do with it. Well... it was all run in the name of the people, let's give it to the people.

How do you give it to the people? Milton Friedman, I believe, had a proposal. It was the right conceptual model, but was impractical in practice, which is, make these companies publicly-traded companies, with stock in them, then give every member of the population an equal share of stock in each of the companies, then allow them to trade on an open stock market. It's the perfect model — it succeeds at dividing things equally without having to make any judgments ahead of time as to what the relative valuations of different things are; and, by allowing trade from there, it allows people to determine the values and to individually own certain things by buying it up from everybody else.

The reason it's impractical is that the transaction costs are too high, because, let's say, just to pick numbers out of the air, you have 10,000 companies, you have a population of a million, then each member of the population would have a one-millionth share in each of 10,000 companies. So each individual stock certificate is worth so little that you end up not keeping track of it, not looking for trading opportunities, and not looking for opportunities to buy from someone else because what you're buying from them is so much less than the effort of finding them and

convincing them that they should sell it to you.

So what they did instead is something which I think is quite brilliant— it's the first practical alternative to "I cut, you choose," and it's something that game theorists should really take seriously because it solves the "I cut, you choose" model for a large multi-person game, whereas the standard multi-person generalization of "I cut, you choose" is an N-squared algorithm.

What they did was to create a set of mutual funds. Once again, these are not the real numbers, let's say they created ten mutual funds. Each of the mutual funds has an equal share in each of the 10,000 companies, and then each member of the population has an equal share in each of the ten mutual funds. Now, every individual bit of ownership is itself individually sufficiently valuable to make a market at it, and after t_0 , everything's tradable on the open market —

The managers of the mutual funds can trade in shares of the companies, and the individual members of the population can trade in shares of the mutual funds.

But in addition, individuals can trade in the shares of companies, the companies can trade in the shares of the mutual funds —

The mutual funds can buy shares of each other.

Right. So the whole thing's on the open market, no constraints. Now, unfortunately, Poland did not succeed at implementing this scheme in this neutral way because it got bogged down in a political process of special interests. But it was the right model and it's exactly the same problem: Poland faced the common heritage problem. Those companies were the common heritage of the Polish population; they were declared to be so by the Communists. The successor government decided to actually implement what the Communists had been saying, through a real system that actually implemented common heritage through private property, and succeeded in coming up with a way to do so. I think that there's a lot of details to be worked out, because the Universe is not publicly-traded companies, but I think that that's the first viable structure I've heard for Inheritance Day.

It deals with the fact that the Universe in different directions has very different value, so you can't do a spatial division. You largely don't know what's out there yet, so you can't do a division based on having mapped things out. And you don't have a market in it yet, so you can't use any notion of market prices to use that as your basis for division. You need this in order to create that system.

Now, taking it back to nanarchy. The system of property rights that we legislate

At the conclusion of *The Day the Earth Stood Still*, the alien, Klaatu (played by Michael Rennie), having been murdered by the American army, then rescued and restored to life by the robot policeman Gort, delivers this speech to a gathering of scientists from every Earth nation:

I am leaving soon, and you will forgive me if I speak bluntly. The Universe grows smaller every day, and the threat of aggression by any group, anywhere, can no longer be tolerated. There must be security for all, or no one is secure. This does not mean giving up any freedom, except the freedom to act irresponsibly. Your ancestors knew this, when they made laws to govern themselves, and hired policemen to enforce them. We of the other planets have long accepted this principle. We have an organization for the mutual protection of all planets, and for the complete elimination of aggression. The test of any such higher authority is of course the police force that supports it. For our policemen, we created a race of robots. Their function is to patrol the planets, in spaceships like this one, and preserve the peace. In matters of aggression, we have given them absolute power over us. This power cannot be revoked. At the first sign of violence, they act automatically against the aggressor. The penalty for provoking their action is too terrible to risk. The result is, we live in peace, without arms or armies, secure in the knowledge that we are free from aggression and war... free to pursue more profitable enterprises. We do not pretend to have achieved perfection... but we do have a system, and it works. I came here to give you these facts. It is no concern of ours how you run your own planet, but if you threaten to extend your violence, this Earth of yours will be reduced to a burned-out cinder. Your choice is simple: join us, and live in peace, or pursue your present course, and face obliteration. We shall be waiting for your answer. The decision rests with you.

by signing on to the Moon Treaty, as written, with this attached implementation protocol — that's the proposed system of property rights for a future nanarchic system to either directly impose or, if we can get away with something more minimal (which would be good, because something that ambitious has a lot to it), to impose a more minimal set of constraints, on top of which we can bootstrap a system of enforcement mechanisms that are capable of enforcing such a system of property rights, but in which that system that's layered on top is not an irrevocable system.

This is very much like operating system design. The kernel of the operating system is something that you can't escape from. It sets the foundational rules. There's a methodology in software engineering called "mechanism/policy separation," where you want to embody in the irrevocable kernel as few policies about how to do things as possible... basically, just the enabling mechanisms that allow all of those policies to be built in a diverse way by different users of the operating system to serve different needs, and enabling experimentation.

There's also a good analogy to the essence of constitutional systems. An "unamendable constitution" is really sort

of the kernel of your operating system. What you'd like is that the kernel of that nanarchy system (which we've spread into the universe in an irrevocable way) is just sufficient that we can convince ourselves that systems layered on top of it to enforce complex systems of property rights are themselves things that we can experiment with and change and modify and amend, as we learn more and as we evolve, and that the underlying nanarchic framework is just sufficient to let that process of amending that system of enforced property rights be one which is not destabilized by runaway military replicators.

By the way, one enormous piece of credit that I'm amazed that I haven't been saying repeatedly over and over again in this is Eric Drexler. A lot of these ideas, perhaps most of these ideas, about nanotechnology and military stability in the future come from conversations with him, as part of the conversational process. I would say I did more of the bouncing and he did more of the thinking for a lot of this. In all of the years in which Drexler and I have collaborated and engaged in intense conversations out of which wonderful ideas came up, there's only one incident where we each thought we were the proper author of the idea and the other

one was mistaken. Sometimes, when we think we're mistaken, it's when we each think the other one's the proper author, not himself.

I'd like to make another point here: the issue of "What is property rights? What is aggression and coercion?" so that we can have something for this framework to be enforcing. The great problem there is, humankind has been arguing political philosophy with each other for many many centuries, and we're a long way from universal global consensus. And, we're also about to step into a world in which, by virtue of these problems of identity, the discreteness of "moral agency" goes away. "Who is a moral agent?" becomes problematic.

If I build a fully sapient AI that then goes out and murders someone, am I responsible?

Right. The second issue is, "What is a good reconstruction of political philosophy in a world of ambiguous moral agency?" This question is fascinating to me, and is much of the reason why I pursued the work that Eric Drexler and I did on agoric open systems.

Now, libertarians have often said that the only things which should be considered moral agents, and therefore the only things which should be covered by non-aggression rules, are things which are sentient. That gets you into the infinite morass of "What the hell is 'sentient'?" And it's unresolvable: it's an observer question, not an intrinsic question of the thing that you're labeling.

Agoric open systems are a system of rules that are extraordinarily similar to the system of rules that libertarians of all five stripes have been proposing: it's a system of rules extraordinarily similar to that, as the operating rules for a computer system — an operating system or a programming-language system. Both programming languages and operating systems are foundational computational systems on top of which you have a lot of computational entities interacting with each other, and different operating system designs and different language designs establish different frameworks for the interaction of agents. Another way to think of it is that, every operating system and every programming language is essentially a different set of laws of physics, in which the agents are interacting.

The thing that's fascinating to me is that we did not need to impose any 'sentient' constraints. In fact the issues brought up by the whole process of thinking about agoric systems made clear that you want to assign rights to lots of little things that are clearly very far from sentient. A little mathematical server, a sort of equation-solver object in your computer system, is something you want the agoric framework to treat as having full property rights. Similarly, any other little computational

server object that other objects make requests of and that itself uses computational resources writes subcontracts out to other objects. It's a system of property rights governing the interaction of all these computational objects. Even though none of them are sentient, it's clearly the case that the system as a whole does much better by giving to each of these non-sentient objects full "libertarian" rights, because that's a better framework of rules for governing their interactions than any other.

Let's examine the issue you brought up. You create an AI, and it goes out and kills somebody; are you responsible? Now, the interesting thing is, under nanarchy, and under agoric open systems, that issue doesn't come up, and the reason it doesn't come up is the other difference between this and governmental systems, which is that current governmental systems of rights enforcement don't actually enforce rights, they punish violators of rights. The idea with both nanarchy and agoric open systems is that we've added this additional constraint: It's impossible for you or an agent you create to murder that person. When the coercion is attempted, it is prevented. That would be the case in a nanarchy — the thing that is monitoring for certain inter-boundary activities that may be coercive, stops any that fall within the possibility of coercion. They just don't succeed in happening. There's also no issue of whom to punish — what does execution mean to someone with backup copies?

And, in agoric open systems, you can no more steal from someone than you or I could go faster than the speed of light. You just don't have the language to express coercion; you don't have the tools available in which you could create a concept of coercion. So, we can talk about the distinction between pre-enforcement and post-enforcement. Post-enforcement depends on punishment creating an incentive to not commit a crime, and that gets trashed by post-Singularity confusions of identity, so what we need to do is transfer to a system of pre-enforcement.

A lot of these ideas about the strategies of the long-term future of the Universe are actually ideas that are many years old; I feel like I have had very few new ideas with respect to that in the last few years. The thing that's really changed is that now there's a community to say these ideas to. Back when these ideas were being hatched, there was no such community. And I want to express my deep gratitude to the whole phenomenon of Extropianism for creating a community of people that can share and bounce around ideas like this.

I think you've put your finger on something that everybody in the Extropian community feels; I felt the same way when I first discov-

ered that there was an Extropian mailing list— I discovered the mailing list before the magazine. That there were other people out there — "Yeah, they have all these same ideas, and they consider them mutually consistent —

The XanAMIX community, the Extropian group, and the Assembler Multitude, I consider to be really all one community; it's just that different names are for different foci of the community. We need to get more integration between the Agorics folk at George Mason University and Extropianism. Extropianism is much more technologically literate and much more technologically informed than they are. Obviously, going the other way, they're much more economically inclined than the Extropians are; the non-obvious but much more important thing is they're much more epistemologically well-informed than the Extropians are. It seems to me that epistemological issues are the great missing piece in Extropian philosophy right now.

I think Max More would be surprised to hear you say that, and I know that he would want to launch into a long and involved conversation on that topic.

Very, very late at the Too Many Erics party, I actually did say this to Max More, and he wrote down my recommendation, which I will repeat here. I have certainly not read every book on epistemology, but in my opinion, the *very best book on epistemology*, for what that means from somebody who hasn't read all of them is, *The Retreat to Commitment*¹³, by Bill Bartley. My one-paragraph summary of the history of epistemology is:

Hume said, "We can't really know anything," but nobody believed him, including Hume. Then Popper came along and said, "Hume was right, but here's what you can do instead." And then Bartley came along and had one hell of a debugging session.

Bartley's evolutionary epistemology is essentially Popper's work debugged, and it's the only debugged epistemology that I've ever encountered.

While the tape was off just now, you said you find it weird that you find yourself so often going back to a scene from a movie as an illustration for discussions about post-Singularity life. That reminded me of something that Jim Bennett said the other night at the Assembler Multitude. The topic for discussion was privacy and technology, and he used the example of the Sicilian organized crime families in the United States and the tongs in China, who would develop these idiomatic languages that were very rich in allusions to tales and stories that were not familiar to the outside, so that they could use that as a communication code, in order to conceal meaning from outsiders while communicating with insiders. So there's an Extropian core of

common movies and stories, and I can refer to "The Ungoverned" and you automatically have an indication of what topics that might lead to.

Yes, and even though you haven't actually seen "The Day The Earth Stood Still", nevertheless my reference to it had a lot of content for you.

Other references to it in our environment are rich enough that I got the gist of it.

Another big issue about libertarianism and technology in the future, is that libertarians are of confused and contradictory minds about the issue of privacy, and what a right to privacy is.

That became apparent at the Assembler Multitude session that I was describing. We went from "Gee, it would be nice if we could keep an eye on all the people in high places," to "Gee, it would be nice if no one could keep an eye on anyone."

People who thought they were defending the cause of freedom thought they scored a big victory in preventing Lotus Marketplace from being distributed, and I think that they're rather badly mistaken. I think that something that I have to say about you to a third person is not your information, it's my information. It's something that I know; it's in my head, it happens to be about you but it's my information about you, and I am not violating your privacy by speaking that information to somebody else, unless I have an agreement with you not to do so — unless either I heard that information in a way that contractually constrained me from revealing it further, or I acquired the information itself through illegitimate means such as by breaking into your house.

If I am, let's say, witness to a transaction that took place — I watched you engage in a transaction with a shopkeeper, and I saw you buy that stick of bubble gum for five cents. I saw you do it, you and the shopkeeper didn't do what you may have needed to do to prevent me from being able to see it. The fact that the transaction occurred is my information, and I don't violate anyone's rights by reporting it further.

Credit reporting agencies, and Lotus Marketplace, are essentially a large-scale form of that kind of reporting of information that no one was contractually obligated not to report.

It's my understanding that a large part of Lotus Marketplace was census information, in the sense that it broke down average incomes by nine-digit ZIP code, to the extent that you would know that the average income on the block where So-and-so lives is X dollars.

One of the things that's going to happen with large-scale online databases and online media where people post things that other people can read, is that in some

sense gossip will travel faster and be more persistent. Spontaneous order admirers, particularly libertarians and Extropians, should understand, first of all, that gossip is speech that should be protected by free speech — what Lotus was doing was in some sense engaging in a very large-scale piece of gossip — and spontaneous order admirers should appreciate and understand the positive values of gossip.

In some sense I'm being provocative and hurting my cause by using the word *gossip* for this because it's a word that has negative connotations, but in fact a lot of the dynamics by which people in small towns were constrained to be decent to each other, and a dynamic which, to a significant extent, has been breaking down in large cities where so many interactions are anonymous and untraceable, is the fact that actions have consequences, in that they would be reported back and forth through the gossip mill — people realized the cost of a negative reputation, as well as value of a positive reputation, through the mechanism of people feeling free to talk to each other about other people.

With larger-scale societies, verbal communication between people just doesn't have the spreading power to keep the dynamic of that reputation system intact. However, computer-based media that violate people's intuitions about privacy are simply the old small-town reputation mechanisms, scaled-up by technology to a scale that can deal with the current scale of society that the old gossip mechanisms can no longer deal with.

I can see how through something like Netnews, I can have an opinion about the trustworthiness or the judgment of someone on the opposite coast whom I've never met, based on the evaluation of their behaviors, particularly verbal behaviors they've emitted over the net, by other people.

By the way, I should point out that much of the credit for this thinking comes from Gayle Pergamit, via conversations I had with Phil Salin.

I think that due to the public online media, combined with cryptographic technology, we're going to find that information which is not out of the bag is able to be kept in the bag; people are able to keep certain information private, much more effectively than they ever have in the past, and private from eyes that are prying with millions of dollars worth of resources behind the prying effort —

And acres of computing power.

Right. And we've never been able before to be securely private against those resources when they're concentrated on anybody; cryptography gives us that. So, in some sense, things that successfully get defended in their privacy are really defended much better than they have been in the past, but on the other hand, once it's

out of the bag, once it's public, it becomes much more visibly and persistently public. I think that it's very important to recognize that what we have in this transition is a system that's much more equitable with respect to information because, previously, information that got out of the bag was accessible to elites and not accessible in general.

This is why the Lotus Marketplace decision should be incredibly offensive to lovers of freedom. It's not that the information on that CD-ROM is otherwise inaccessible; it is simply otherwise accessible only to elites. It's not a question of people's privacy being violated, it's a question of accessibility to information that is claimed to already be used to "violate people's privacy" —

It would be more equitable to have it more widely distributed, to at least give the non-elites the same access to the same information. Right. By the way, I have no idea whether it's even conceivable that you could work any of this into a single article.

What I'm envisioning now is a series of these interviews... and I'd like to thank you for getting things off to a great start.

I pride myself on being one of the few people — Max More and Eric Drexler also certainly included — to take seriously both hard issues of philosophy and a post-Singularity future, to really take both seriously, and I think that that, rather than any skill at philosophizing, is really what's responsible for my being able to make interesting progress in the philosophy for a post-Singularity world.

Notes

¹Nelson, Ted. *Literary Machines*, Edition 87.1. Published by the author.

²Heinlein, Robert. *The Moon is a Harsh Mistress*. New York: Berkeley, 1966.

³Friedman, David. *The Machinery of Freedom*, 2d Edition. La Salle, Ill.: Open Court, 1989.

⁴Nozick, Robert. *Anarchy, State, and Utopia*. New York: Basic Books, 1974.

⁵Epstein, Richard A. *Takings*. Cambridge, Mass.: Harvard University Press, 1985.

⁶The bibliographic citation for this article is not available.

⁷*The Day the Earth Stood Still*. Twentieth-Century Fox, 1951. Available on CBS-Fox Video.

⁸"Democracy is the worst form of government — except for all the others."

⁹*War Games*. United Artists, 1983. Available on CBS-Fox Video.

¹⁰*Doctor Strangelove: Or, How I Stopped Worrying and Learned to Love the Bomb*. Hawk Films, 1963.

¹¹Inheritance Day was proposed by Eric Drexler in *Engines of Creation* (New York: Doubleday, 1986).

¹²Also in *Engines of Creation*, Drexler presents an evolutionary argument explaining why near-lightspeed expansion is almost inevitable.

¹³Bartley, William W., III. *The Retreat to Commitment*. La Salle, Ill.: Open Court, 1984.

Extropy Institute

While *Extropy* celebrates the commencement of its sixth year (see facing page for information), Extropy Institute completed its first in May. The year was incredibly busy, challenging, and encouraging. As an *Extropy* reader, you may want to hear what the journal's umbrella organization has accomplished, and what ExI's goals are for Year Two.

What Have We Done? Starting from zero in May 1992, ExI membership has grown steadily, passing 100 in December, and 200 by the time you read this. This response, coming despite painfully limited funds for advertising our existence, encourages us and sustains our efforts.

Extropy's print run, between last summer and now, has ballooned from 750 to 3,200 copies. Alas, this figure has not been matched by a proportionate improvement in finances, given minimal advertising revenue, since most copies are sold through distributors who take more than half the cover price. Happily, direct subscriptions are also growing, though more slowly.

Awareness of our existence has been pumped up by increasingly frequent media attention. The widely circulated, pro-tech *Wired* gave *Extropy* a mention in their second issue, sparking off a prolonged flurry of information requests. *Factsheet Five*, a thriving review of alternative publications, gives *Extropy* especially favorable reviews. A major feature in the British *GQ* ("Meet the Extropians") is drawing further media interest. As a result, ExI director Russell Whitaker has been busy giving interviews in England. Here in the U.S.,

I've appeared on ideas chat shows like *Breakthroughs: A TransCentury Update* (hosted by futurist Nancie Clark), and *Electric Coffee*, plus a few documentaries. Recently, I was delighted to be flown to Amsterdam by organizer Luc Sala to speak at the New Edge Conference — a tremendously enjoyable event where I met numerous creative and future-oriented people.

For details of membership dues, see p.2, lower right box.

Extropy Institute

11860 Magnolia Avenue, Suite R
Riverside, CA 92503
909-688-2323

ExI Directors

Max More, President, Editor of *Extropy*. more@usc.edu
Tom Morrow, Vice President. tOmorrow@aol.com
Simon! D. Levy, Editor of *Exponent*. levy@yalehask.bitnet
Tanya Jones, Treasurer. c/o alcor@cup.portal.com
Ralph Whelan, Secretary. c/o alcor@cup.portal.com
David Krieger. dkrieger@netcom.com
Russell E. Whitaker. whitaker@eternity.demon.co.uk

The Extropian cultural and intellectual network has grown and deepened. Parties spring up frequently, often when an Extropian is visiting from afar. The Bay Area Extropians now gather weekly for lunch and intellectual stimulation.

A major locus of Extropian culture is the Extropians e-mail list, still getting stronger and more sophisticated after two years. The evolution of interaction in cyberspace is being pushed along by voluntary social experiments on the list, and by advanced list software, planned and implemented by Harry S. Hawk (ExI's Electronic Communication's Officer), and Ray Cromwell, aided by people like Perry Metzger (original founder of the list). These and other developments are reported regularly in ExI's newsletter, *Exponent*. Thanks to Editor Simon! D. Levy, *Exponent* appeared every two months.

What Will We Do in Year 2? Some of our most certain goals include: (1) Double *Extropy's* frequency, publishing quarterly, while continuing to improve it, by spreading the creative input wider. Improve finances by building the direct subscriber base, and attracting advertising. (2) Publish *Exponent* more frequently (9-12 per year). (3) Spur formation of local discussion groups, as well as the Nexus Network (domiciles shared by Extropians with supportive values and lifestyles). (4) Produce an ExI information booklet (given to all new members), explaining our principles, history, activities, and goals, plus a guide to terms and ideas common among Extropians but unfamiliar to most others. (5) Boost membership by information mailings to sympathetic groups. (6) Develop ExI's Advisory Council. (7) Start gathering material for an Extropy book. (8) Hold Extro 1 conference (see p.46), and present Extropy Awards. (9) Build up book service.

If you share our values and goals, we urge you to join us in the core of the Extropian community, as an Extropy Institute member. As a member you will receive detailed analysis in *Extropy*, topical news and essays in *Exponent*, and discounts on various items, including conference fees.

As a member, you will also be sustaining our efforts to build an extropic culture. Scientific humanism and optimism were strong in the intellectual atmosphere of the nineteenth century. Today's prevailing grey climate of gloom, credulity, and timidity is waiting to be dispelled. Join us in transhumanising the culture, taking charge of our lives, and creating a future fit for posthumans.

Upward and Outward!

Max More
President, Extropy Institute



Extropy Institute

Extropy 5th Birthday Party Extropaganza

August 28 1993, from 2.0pm

Extropy: The Journal of Transhumanist Thought — the premier intellectual futurist publication — will be celebrating its fifth birthday on Saturday August 28 1993. The first issue of *Extropy* — then a 24-page, small-format magazine with a print run around 50 — was collected from the printers by editors and publishers Max More and Tom Morrow on Saturday September 8 1988. To avoid clashing with other major events near the end of the summer, we're celebrating a few days early.

Evolving drastically throughout its five years and 11 issues, *Extropy* is now read by over 3,000 people and is distributed internationally. From issue #12 in January 1994, *Extropy* will return to its original quarterly publication schedule. The energy and enthusiasm generated by and attracted to *Extropy* led to the formation of Extropy Institute (ExI), under whose umbrella it is now published, and to spinoffs such as the Extropians e-mail list, and ExI's newsletter, *Exponent*. All subscribers, readers, and interested parties are invited to join in the birthday celebration, for a day of fun and intellectual stimulation.

The festivities will be opened at 2.0pm by Max More and Tom Morrow, with a comparison of the early issues and the current issue.

Long-time Extropian Mark Desilets has kindly made available his home for the event. Attendees should bring drinks and

food with them, so as not to overburden our host. Also bring along appropriate toys, gadgets, and a playful attitude. Mark's house does have a hot tub, so come prepared; please note that some clothing will be required in the tub, so as not to shock the neighbors with the sight of our transhuman physiqués!

Here's what you need to know:

Date: Saturday August 28 1993, starting 2pm.
Location: 580 Burnside Bend, Boulder Creek, CA 95006
(H): (408) 338-0636 (W): (408) 437-5122

Mark says:

My house is in Boulder Creek, which is on Hwy 9, which runs from Saratoga (in the San Jose area) to Santa Cruz. The closest quasi-major town is Scott's Valley, which is on Hwy 17, running from Oakland through San Jose to Santa Cruz. There will be (floor, alas) crash space for about as many as need it, and there is plenty of local hotelery for those so inclined and so financially endowed. Also, there is Big Basin Redwoods State Park just a few miles away, for anyone who wishes to camp. What more can I tell you?

Questions? Contact Mark as above, or Max More at Extropy Institute: 909-688-2323

Back Issues

#1, 2, 4, 5, 6: \$4 each; #7, 8, 9, 10: \$5 each.
Available from Extropy Institute (address, p.2)

#10, Vol.4 No.2 (Winter/Spring 93): Pigs in Cyberspace, by Hans Moravec; Protecting Privacy with Electronic Cash, by Hal Finney; Technological Self-Transformation, by Max More; Mark Miller interview, by David Krieger, Pt. 1: Creole Physics & the Credit Theory of Identity; Nanocomputers: 21st Century Hypercomputing, by J. Storrs Hall; The Transhuman Taste (Reviews); Two books on Ayn Rand & Objectivism; *Nanosystems*; *Genius*.

#9, Vol.4 No.1 (Summer 1992): The Extropian Principles, 2.0, by Max More; Extropy Institute Launches, by Max More; Persons, Programs, and Uploading Consciousness, by David Ross; Nanotechnology and Faith, by J. Storrs Hall; The Making of a Small World (fiction), by R. Michael Perry; Genetic Algorithms, by Simon! D. Levy; Time Travel and Computing, by Hans Moravec; Futique Neologisms 3; Exercise and Longevity, by Fran Finney; The Transhuman Taste (Reviews); *The Anthropoc Cosmological Principle*, *The Blind Watchmaker*, *The Ultimate Resource*, *Population Matters*, *The Resourceful Earth*, *Bionomics*.

#8 Vol.3 No.2 (Winter 1991-92): Idea Futures: Encouraging an Honest Consensus, by Robin Hanson; Dynamic Optimism, by Max More; Neurocomputing 5: Artificial Life, by Simon! D. Levy; Futique Neologisms (futurist lexicon); Extropia: A Home for Our Hopes, by Tom Morrow; Human-Transhuman-Posthuman, by Max More; Reviews of: Stiegler's *David's Sling*, Drexler's *Unbounding the Future*,

Platt's *The Silicon Man*; News of scientific advances and movement news; Reviews of zines.

#7 Vol.3 No.1 (Spring 1991): A Memetic Approach to 'Selling' Cryonics, H. Keith Henson & Arel Lucas; Privately Produced Law, Tom Morrow; Order Without Orderers, Max More; Futique Neologisms; Neurocomputing 4: Self-Organization in Artificial Neural Networks, by Simon! D. Levy; Forum on Transhumanism; Reviews of *Smart Pills*, *Surely You're Joking Mr. Feynman*, *Great Mambo Chicken* and *the Transhuman Condition*; and more...

#6 (Summer 1990): Transhumanism: Towards a Futurist Philosophy, by Max More; The Thermodynamics of Death, Michael C. Price; The Opening of the Transhuman Mind, by Mark Plus; The Extropian Principles, by Max More; Neurocomputing Part 3, by Simon! D. Levy; Forum on Arch-Anarchy and Deep Anarchy; Reviews: *Order Out of Chaos*, *The Emperor's New Mind*, *A Neurocomputational Perspective*, *Loompanics Greatest Hits*, *The Machinery of Freedom*; Extropian Resources, and more.

#5 (Winter 1990): Forum: Art and Communication; Leaping the Abyss, by Gregory Benford; Arch-Anarchy, by A; Deep Anarchy, by Max O'Connor; I am a Child, by Fred Chamberlain; Perceptrons (Neurocomputing 2), by Simon! D. Levy; On Competition and Species Loss, by Max O'Connor; A Review of Intoxication, by Rob Michels; Intelligence at Work, by Max O'Connor and

Simon D. Levy; Extropian Resources, by Max O'Connor and Tom W. Bell; The Extropian Declaration, by Tom W. Bell and Max O'Connor; Our Enemy, The State, by Max O'Connor and Tom W. Bell.

#4 (Summer 1989): Forum; In Praise of the Devil, by Max O'Connor; Neurocomputing, by Simon D. Levy; Why Monogamy? by Tom W. Bell; What's Wrong With Death? by Max O'Connor; Reviews: Are You a Transhuman? Postscript to "Morality or Reality" by Max O'Connor; Efficient Aesthetics, by Tom W. Bell; Intelligence at Work: Advances in Science by Max O'Connor.

#3 (Spring 1989) is out of print.

#2 (Winter 1989): Review of *Mind Children*, by Max O'Connor; Darwin's Difficulty, by H. Keith Henson and Arel Lucas; A Truly Instant Breakfast, by Steven B. Harris M.D.; Wisdomism, by Tom W. Bell; Nanotechnology News, by Max O'Connor; Weirdness Watch, by Mark E. Potts.

#1 (Fall 1988): A brief overview of extropian philosophy and an introduction to some of the topics we plan to address: AI, Intelligence Increase Technologies, Immortalism, Nanotechnology, Spontaneous Orders, Psychochemicals, Extropic Psychology, Morality, Mindfucking, Space Colonization, Libertarian Economics and Politics, Memetics, and Aesthetics; "Morality or Reality," by Max O'Connor.

BUNKRAPT:

THE ABSTRACTIONS THAT LEAD TO SCARES ABOUT RESOURCES AND POPULATION GROWTH

or

WHY DO POLITICIANS THINK GREEN?

Julian L. Simon

INTRODUCTION

Immediately after you show someone that all trends pertaining to human welfare — health, wealth, education, leisure, availability of natural resources, cleanliness of our air and water, you name it — have been improving rather than deteriorating, the question arises: Why, then, do our political leaders tell us the opposite — life is more dangerous, our planet is “plundered” and “in crisis”, we are running out of resources, pollution is increasing — that is, that things are getting worse when they are really getting better? Why do the politicians say that there is need to “save the planet”?

A “commitment” to environmentalist ideals has become a commonplace, an article of faith for just about every community and organization. The extent of consensus that we are in trouble because things are getting worse, and therefore the government should “do something” about it, may be seen in many evidences — for example:

Item: At the Rio Summit (or “Rio Nadir”) of 1992 more heads of state convened than had ever convened elsewhere, and signed a “treaty” concerning the treatment of plants, animals, natural resources, and the environment.

Item: U.S. Vice-President Albert Gore’s book, *Earth in the Balance* (1992), recites the entire litany of environmental laments and is a best-seller. The United Nations Fund for Population Activities 1991 book, *Population, Resources, and the Environment* opens with this quote: “The ‘triad’ of excessive population growth, environmental degradation and poverty threaten us and our planet as never before” (p. 2). And the authors of *Limits to Growth* (Meadows et. al., 1972) are back with a reprise, *Beyond the Limits* (Meadows et. al., 1992).

Item: The “official” *Global 2000 Report* to President Carter of 1980 pitched the U.S. into a fever of environmental activism and became the basic document for the environmental movement, selling over a million copies. Other countries undertook parallel reports such as *Britain 2010*, *China 21st Century*, *Japan’s The Year 2000*, and others for Portugal, Mexico, Taiwan, Norway, Canada, Indonesia, Mauritius, Iceland, Thailand, and Ireland.

Item: In 1987 the Brundlandt Commission, a group of household names chaired by the Prime Minister of Norway, published *Our Common Future* on these same environmental issues.

Item: In November of 1991, the nation’s Roman Catholic Bishops “acknowledged that overpopulation drains world resources”. The bishops asked Catholics “to examine our lifestyles, behaviors and policies, to see how we contribute to the destruction or neglect of the environment”. Even the Pope issued a 1988 encyclical “*In Sollicitudo Rei Socialis*” and a 1990 New Year’s message on this theme of environmental “crisis” and “plundering of natural resources,” and “the reality of an innumerable multitude of people.” The Pope apparently has “gotten religion” and changed his message since then, however.

Item: In 1992 there took place in Washington a “Consultation on the Environment and Jewish Life”, intended as “a Jewish communal response to the world environmental crisis”. The invitation letter said: “We appreciate the many important issues on the Jewish communal agenda. But the threat of ecological catastrophe is so frightening and universal that we believe we must mobilize our community’s considerable intellectual and organizational resources as soon as possible”. The signers of the invitation included just about every important figure

in the organized Jewish community in the United States — religious leaders, communal leaders, Jewish Senators and other elected officials.

The tapestry of explanation for this mass belief is surely very complex. A few of the threads include:

1. Institutional self-interest of the media in trying to catch big audiences with scary bad-news stories.

2. Ostensible environmental problems justifying the funding of scientific research and the budgets of organizations such as the UN’s Food and Agricultural Organization (FAO) and Fund for Population Activities (UNFPA) and the U.S. Council on Environmental Quality whose reason for existence is fear of resource scarcity and population growth.

3. The preservationists who prefer

bucolic surroundings to resource development.

4. Psychological propensities deep in our psyche that predispose us to warnings of doom.

5. The marvelously evocative inflammatory rhetoric that has been created to arouse fear — “population bomb”, “empty pumps”, “save the children”, “end of the world as we know it”, and “end of the age of affluence”, for example. (See my 1981 book, Chapter 22 of 2nd edition, 1993), for a long discussion of this rhetoric.)

6. Simple racism, especially with respect to population growth in other parts of the world, and with respect to immigrants of various shades and ethnicities entering the United States.

7. An attitude toward the factual truth that induces people to exaggerate and even lie when convinced that the eleventh-hour danger to the public justifies such dishonest practices. Joining the environmental movement is seen by many as a last chance to do good, just as joining the Communist Party in the 1930's seemed an opportunity for social contribution by many generous-minded people. Once having joined the movement, foul means are deemed acceptable by many if the end is thought to be beneficial.

8. This essay leaves aside the seven previous elements of the explanation, however, and focuses only on the *ideas* that undergird the newspaper and television stories, the intellectual infrastructure which give these stories credibility. The ideas to be discussed fall into two categories: misunderstandings of the nature of resource creation and population economics, and misunderstandings of the nature of a modern complex social-economic system. The essay attempts to explain why so many people are enraptured with this kind of bunk — that is, bankrupt.

MISUNDERSTANDING THE NATURE OF RESOURCE CREATION AND POPULATION SIZE

The Seductiveness of the Malthusian Logic

Beneath the Malthusian notion of diminishing returns, we find an inter-related set of fundamental ideas that we may call a “vision.” The vision underlying the thinking of today's conventional writers about resources and population is the concept of fixity or finiteness of resources in the relevant system of discourse. This idea is found in Malthus, of course. But the idea probably has always been a staple of hu-

man thinking, because so much of our situation must sensibly be regarded as fixed in the short run — the bottles of beer in the refrigerator, our paychecks, and the amount of energy that parents have to play basketball with their kids.

In contrast, the vision underlying my thinking about resources — also what is now a consensus of other students of these subjects (NAS, 1986), with such predecessors as William Petty, Friedrich Engels, Simon Kuznets, Friedrich Hayek, and the main developer of the idea, Harold

A key difference between the thinking of those who worry about impending doom, and those who see the prospects of a better life for more people in the future, apparently is whether one thinks in closed-system or open-system terms.

Barnett — is that the relevant system of discourse has a long enough horizon that it makes sense to treat the system operationally as not fixed rather than finite. We view the resource system as being as unlimited as the number of thoughts a person might have, or the number of variations that might ultimately be produced by biological evolution. That is, a key difference between the thinking of those who worry about impending doom, and those who see the prospects of a better life for more people in the future, apparently is whether one thinks in closed-system or open-system terms. For example, those who worry that the second law of thermodynamics dooms us to eventual decline necessarily see our world as a closed system with respect to energy and entropy; those who view the relevant universe as unbounded view the second law of thermodynamics as irrelevant to this discussion. I am among those who view the relevant part of the physical and social universe as open for most purposes.

Which vision is better in the context of resources and population is not subject to scientific test. Yet the choice profoundly affects our thinking. I believe that here lies the root of the matter.

Academics are particularly susceptible to the notion of Malthusian diminishing returns, perhaps because academics are more likely than are laymen to believe in abstract theories. (Academics

properly spend much of their lives battling to persuade others that abstract theorizing has importance and is not just an “ivory tower” recreation). In my experience, journalists and businessmen are less likely to be taken with the simple Malthusian abstraction, perhaps because they have no professional stake in this idea (in contrast to many biologists and some economists) and perhaps because journalists are more attuned to reaching judgments and making decisions in light of the full richness of a situation — on their “intuition” — rather than upon the logical relationships in a simple model. (More generally, businessmen and newsmen seem to be more open to new ideas than academics, perhaps because a continuous flow of creative change is more crucial in their occupations.) Another element is the dead hand of expertise. As Kuznets tells, “Experts are usually specialists skilled in, and hence bound to, traditional views; and they are, because of their knowledge of one field, likely to be cautious and unduly conservative” (in Rosenberg, 1972).

It is a puzzle why so many people — with biologists and physicists notable among them — are so sure that there *must* be some constraint to prevent humanity from growing both ever richer and ever more populous, and why theirs is the vision of unexpandable limits. One possible explanation is that each of us tends to bring our professional modes of thought to bear on other situations even if those modes are not appropriate to the situation at hand. For example, biologists liken the human population to an animal population and then apply the animal-ecology notion of “carrying capacity,” though that notion is quite inapplicable to natural resources in a human context.

Another attraction of the closed-system vision is that the closure of the system enables one to use interesting mathematics, especially the calculus and other optimization devices. From a purely physical point of view, a proposition about finiteness (or entropy) requires a bounded system. But where is the relevant boundary for our material world? Around the earth excluding the sun? Around the earth plus sun plus solar system? Around other suns? Around a “universe” which may or may not be finite or expanding in the astronomer's eye? No boundary, no finiteness.

Still another root of the closed-system vision is the bewitching medieval notion of “first cause” or “ultimate cause,” the idea that nothing happens which is not the result of other forces. And pushing back the causal sequence in an infinite regress, it seems as if there must have been an original causal force. This sug-

gests a complete, and therefore closed, system.

For some, the closed-system vision arises because of a natural abhorrence of the loose-endedness of an open system.

An interesting example of how this vision permeates our thinking: If you say that copper might be made of other metals, hearers say "alchemy." When you point out that nuclear bombardment transmutes metals, the hearers say "not practical," implying that it never could be practical. They may be correct. But there is no *logically* binding impossibility theorem applicable here. One can only be sure that something is impossible or impractical if one can be sure that the state of knowledge will not change in the future, that is, that capacities are limited because knowledge is limited. But isn't this just what people said in the past about the possibility of finding smaller constituent parts within the "fundamental" electron, say? And about the possibility of obtaining the vast amounts of energy that we get from a small pile of stuff called uranium? Or for that matter, getting vast amounts of heat out of the black rocks that we call coal? The example of copper and "alchemy" is interesting for the infra-thinking that it brings into the open.

The psychologies of open-system and closed-system thinking must be complex and deep-rooted. Perhaps the latter is related to focusing on the social equality of distribution of a fixed pie, rather than on expanding the amount of pie to the possible neglect of equity considerations; this focus often stems from the emotion of envy (Schoeck, 1969). But whatever the roots, most puzzling is why people who are themselves creative and imaginative should lack faith in others' capacities to respond to problems and shortages with limit-expanding ideas. How can people as powerfully creative as John Bardeen (the only two-time Nobel-winning physi-

clist) and George Mitchell (the Texas oilman and developer) share this vision of limits with the non-doers?

Misunderstanding and Misapplication of the Slogan "There is no free lunch."

The slogan "There is no free lunch" seems to imply that we have to pay for everything we get. This is another case of a good thought going wrong by being applied to situations it was not designed for. This slogan was originally intended to suggest that the *government* cannot supply free lunches to *all of us*, that there is no magic trick by which we can increase our total national resources by passing laws and setting up bureaucracies; rather, we as taxpayers have to pay indirectly, some-

time. In other contexts, however, there *are* free (or below full cost) lunches all the time. None of us always pays the full cost of production for what we get. In the modern world each generation gets its lunch at a lower cost of labor than did earlier generations, because earlier generations responded to their economic problems with ingenuity and energy. Our ancestors bequeathed us the intellectual wherewithal to get our lunch, if not entirely free, at least much cheaper than if we had to start from scratch. Compare what we "pay" to what Europeans had to "pay" for lunch and the other meals a few hundred years ago. They paid most of every day's work, whereas we can buy the same amount of raw food with a small fraction of the work time it cost them. And there is no economic or physical force, and no concept in standard economic theory, that suggests that this progressive reduction in the cost of lunch cannot continue indefinitely. We eat our

cheap lunch courtesy of the sweat of our ancestors' brows in mental as well as physical labor.

Lack of Historical Perspective, and Propensity to Compare the Present to an Idealized State Rather Than to the Actual Past

It is not surprising that most people are not aware that real prices of resources were higher in past years than now; this requires adjusting for inflation, and necessitates knowledge of data back to (say) 1900 or 1800. Hence it is not surprising that views about impending resource scarcity are not informed by the contrary long-run trend of increasing availability.

It *should* surprise us, though, when mature experienced journalists in high positions, people such as James Reston and John Oakes of the *New York Times*, write about how bad things are now without reference to how things were in the past. In 1980 columnist Reston could write about "the civilized world that is now in such deep trouble," saying that "you can hardly pick up a paper these days without wondering what's wrong" and decrying our lack of leadership. Can this man have lived through the depression of the 1930's, Hitler, World War II, the Cold War, the Korean War, and the Vietnam War? And ex-senior editorial writer Oakes reproduces the pessimistic findings of the *Global 2000 Report* almost word for word, like a press-conference handout. How can he, too, have lived through such disastrous times in the past, when the environment was much more degraded and the materials more scarce, and continue to write as if the world is headed straight toward doom?

Advertisement:

Reading Breakthrough!

Increase your reading speed, and comprehension, beyond 1000 words per minute, and then to your limit! And you will be amazed at the height of your limit. You will not be skimming, you will actually be reading every word. Yet, you will be reading faster and understanding more.

(713) 879-5670

**Get a personal telephone demonstration
Feel your speed increase while on the phone**

Send Check or Money Order for \$39.95 (Tx & Hdl Incl)
To:
RCD & Associates
7503 Winkleman Rd., Houston, TX 77083

Name _____
Address _____ Apt. _____
City _____ St _____ Zip _____

Money Back Guarantee

Lack of Distinction Between the Long Run and the Short Run

The distinction between the long run and the short run is crucial to the economics of population. In the developed world, additional people — babies or immigrants — are a burden in the short run. And focusing only on the short-run burden leads to a negative judgment about population growth. But in the long run, more people mean a higher standard of living for others. So the judgment about whether more people are good or bad economically depends on how one trades off the present versus the future. By most of my calculations, the discount rate would have to be quite high for additional people not to have a positive present value.

Furthermore, short-run costs are inevitable and obvious, whereas long-run benefits are hard to foresee. If your neighbor has another child, surely your school taxes will go up and there will be more noise in your neighborhood. And when the additional child first goes to work, per-worker income will be lower than otherwise, at least for awhile. It is, however, more difficult to foresee and to understand the possible long-run benefits. Because the increase in knowledge created by more people is non-material, it is easy to overlook. Writers about population growth mention a greater number of mouths coming into the world, and more pairs of hands, but never mention more brains arriving. This emphasis on physical consumption and production may be responsible for much unsound thinking and fear about population growth.

Disbelief in the Relationship Between Population and Knowledge Creation

To many, it is implausible that additional people cause more technical knowledge and advance in productivity, *ceteris paribus*. One source of misunderstanding is the common belief that new technical knowledge usually arises spontaneously, and without connection to social needs. But there is now ample evidence that increased output and investment in a given industry induce more inventions to be made and applied. This “demand-side effect,” as economists call it, can be seen in systematic studies of learning-by-doing, where the time required to complete an airplane or ship decreases as more units are made. The effect can also be seen in systematic studies of comparative productivity in the industries in the U.K. and

in Canada that are relatively large and relatively small compared to the same industries in the U.S. (see Simon, 1981, 2nd edition forthcoming, Chapter 14) And Bernal in *Science and Industry in the Nineteenth Century* (1953/1970) provides additional evidence — case studies of steel; electricity, light and power; chemistry, bacteriology and biochemistry; and the theory of heat and energy in the 19th century — showing that innovations respond to economic demand. In the case of electricity, for example, “The barrier, or rather the absence of stimulus to advance, was economic. Electricity developed quickly when it paid, not a moment before (p. 131).” And a large population size and density imply higher total demand, *ceteris paribus*, which is why Edison’s first street lighting was in New York City rather than in Montana. It is also clear that countries with more people produce more knowledge, assuming income is the same, e.g., Sweden vs. the U.S. And Bernal shows how the power of final demand works indirectly, too. “Once electric distribution on a large scale was proved feasible and immensely profitable, then came a demand for large efficient power sources (p. 129),” leading to the development of turbines. And the development of light bulbs led to advances in creating vacua, after the subject “had stagnated for about two hundred years.... Here was another clear case of the law of supply and demand in the development of science and technology (p. 125).”

On the “supply side” there is also much misunderstanding, especially the belief that the number of potential inventors does not matter. One source of this misunderstanding for some is the idea that, to paraphrase, “One need only contrast innovation and creativity in tiny Athens in the Golden Age with monstrous Calcutta” now, or Calcutta with Budapest of the 1930’s, to see that more people do not imply more technical knowledge being produced. This argument leaves out the all-things-equal clause; Calcutta is poor. And, underlying this argument is the implied (but unwarranted) assumption that Calcutta is poor *because* it has so many people.

If we make more appropriate comparisons — comparing Greece to itself and Rome to itself during periods with different population sizes and growth rates, and industries of various sizes in different countries now — we find that a larger population is associated with *more* knowledge and productivity, because there are more potential inventors and adopters of new technology. Graphs that plot the numbers of great discoveries, and the population sizes in various centuries

in Greece and Rome, bring out this conclusion very nicely (Simon, 1980, or 1981, pp. 200-201).

On the related question of whether material well-being can be improved through there being more ordinary persons — not geniuses — who contribute to our knowledge in their everyday work, the story of electricity and power production is again illuminating. Bernal describes the “stumbling progress of the first fifty years from 1831 to 1881... the effort put into the development (1831-1881)... was small.” The people who made the necessary technical developments “were not geniuses... and others no more gifted could have hit upon these ideas earlier if the field had attracted enough workers (pp. 130-131). As said by Soichino Honda, the inventor and founder of the Japanese motorcycle and auto firm bearing his name, “Where 100 people think, there are 100 powers; if 1,000 people think, there are 1,000 powers” (*The Wall Street Journal*, Feb. 1, 1982, p. 15.)

Confusion Between Trends and Levels, Between Whether Things are Getting Better or Getting Worse and Whether They are Good Now

A frequent and crucial error in the thinking of the doomsayers is neglect of the lessons that experience teaches. And often the doomsayers criticize their opponents on the grounds that we are extrapolating from the past on the assumption that the past usually bears some resemblance to the future. These critics prefer that we form our conclusions purely by analysis of the structural elements, physical and otherwise, that they decide are the most important variables. This is ironic, because to the extent that we have knowledge of each of these elements, that knowledge is based upon experience — systematic and otherwise — of the operation of those elements in history and in scientific experience. As Macaulay, I believe, put it about 150 years ago, If we cannot learn from history, what *can* we learn from? This is not to say that the future is simply an extension of the past; the number of horses in the U.S. did not continue growing throughout the 20th century (although I have read that there are more horses alive in the U.S. now than at the turn of the century). Nor am I downplaying the key role of theory, which is a generalized and formalized structure that embodies our accumulated experience in a particular

field. What I *am* saying is that to assume that the future will bear no resemblance to the past in a particular context, without extraordinarily weighty reasons to believe that there has been a turning point, is to court serious error.

Many make an unwarranted logical leap from the fact that things are not good in some places to the fact that things are getting worse everywhere. This leap is coupled with a lack of historical perspective — for example, a sense of how much worse-off Mexico City and its inhabitants were twenty or fifty years ago compared to today. And when the doomsayers cannot avoid admitting that at least some of the trends in the past have not been toward things getting worse, but rather toward things getting better, they often reply: But history is not a good guide in this connection, because we are now at a turning point in history.

For example, when I say that the history of humankind is the history of people responding to existing and impending problems with solutions that leave us better off than if the solution had never arisen, others sometimes poke fun at the notion that experience of the past is a sound basis for forecasting the future. The issue of whether we are now at a turning point needs some close attention to dispatch it satisfactorily (see my 1984, also in Simon, 1990, Selection 47). But all throughout history people have felt that they are at a turning point, and it hasn't turned out to be so. More generally, if we cannot base our judgments about the future largely upon past experience, in conjunction with reasonable theoretical explanations of that experience, then all of our experience is without value. I doubt that many people really do wish to reject experience as a teacher in this manner.

Belief That What is in Print and What is Said "Officially," Must Be Right

Consider this statement from a recent letter:

[Y]ou said that the transformation of farmland to urban use is far less than society is led to believe. I find this very outraging because I think you made a very blind statement. You have given many people the idea that we're not really losing that much farmland than what the Government or Department of Agriculture or Farmers claim. I have enclosed a pamphlet [from the "official" National Agricultural Lands Study] which is proof to my claim.

Well, you might say, the writer doesn't sound very sophisticated. But the "offi-

cial" label which gives statements the full authority of the federal government was prominent in most newspaper and television stories about the "Vanishing Farmland Crisis" of the early 1980s. If that document (and the report of the National Agricultural Lands Study) had not displayed a government label, they probably would never even have appeared in print, let alone been discussed in national publications, because their level of technical competence and scientific proof was so low. (On second thought, maybe my confidence is misplaced. *The Limits to Growth* was widely publicized even without an official label. But that book had the backing of the wealthy Club of Rome and a hired public relations agency; the whole story was told in *Science*.)

Differences in Conceptions of Human Nature

The main interest of David Hume and Adam Smith was human nature, and they came to study economics as an outgrowth of that interest. Differences in conceptions of human nature are at the root of much disagreement about economic issues, and evidence about the validity of these different views is relevant to decisions about the economic issues themselves. (Unfortunately for the discipline of economics, that explicit focus of attention has been lost in the mathematics that constitutes so much of modern "sophisticated" and "rigorous" and "elegant" economics.)

For example, the doomsayers who desire more government intervention in the production and consumption of natural resources, and the optimists who argue for non-intervention of the government in resource markets, differ in their views of how individuals and private enterprises behave in the face of economic opportunity; they also differ in their views of the performance of government personnel and agencies when entrusted with economic tasks. This thought first struck me as I was out jogging one morning near Asheville, N. Carolina, and I found myself on "Old Toll Road" going up a secluded mountain. Long ago, private enterprisers must have built that road in hopes of making a profit from traffic in that difficult country. And the end result of their private desires was a benefit for the public that continues until today. Interventionists are likely to believe that if government does not provide such services, they will not be provided at all.

I am not suggesting that government should play no role in our economy. Space certainly would have been explored later without government action (or without

competition between countries, perhaps), which might (or might not) have been an economic loss. But given opportunity, private enterprises will supply more ventures than doomsayers expect, more quickly, and at less cost to public — especially in the field of natural resources — partly because individuals bear costs of the failing ventures, rather than taxpayers.

Concerning the difference in views of public and private performance: I imagined a conversation with a (say) potato-chip distributor about possible competition by a government agency. I guessed that his/her first reaction would be to laugh at the possibility that a government bureau could even come close to her/his prices and quality without massive subsidies. But then she/he would reflect that if a government agency got into trouble because it could not compete, it could lower its prices to the competitive level, lose money, and then reach into the public pocket to make up the losses. That would be less funny, and not unrealistic; indeed, it is an accurate description of socialist enterprise East and West.

Another difference in views of human nature concerns its changeability. Reformers, starting perhaps most vividly with William Godwin (to whose writing Malthus' *Essay on Population* was a response) believe that human nature is quite malleable — for example, that self-interested behavior can be reduced by the proper social environment. This belief is very important in Marxism; it implies that one can design a social system that has particular desired properties, and then expect people to be molded to fit that system. In contrast, the Scottish moralists — David Hume, Adam Smith, and their teachers and friends — tended to see human nature as relatively immutable, which implies choosing a social and economic system that produces the best results given that fixed human nature.

MISUNDERSTANDING OF THE NECESSARY NATURE OF AN EFFICIENT MODERN SOCIAL SYSTEM WITH MANY PARTICIPANTS, GOODS, AND PRODUCTION TECHNIQUES²

People Yearn for an Organization of Society Which Reflects the Best Aspects of the Mode of Organization of the Family

In a family, members share goods out of love and altruism, and their decisions about individual and family activities are (at least sometimes) affected by caring thoughts for one another. But this mode of social organization cannot work nearly as effectively when a) individuals cannot know the preferences of all others in the society, b) their capacity for empathizing with another is diminished by lack of intimate relationship, c) there is no accepted hierarchy as there is between parents and children, and d) the number of goods and possible transactions is very large. But many persons find it abhorrent to turn over the function of distribution to the impersonal market. And market distribution seems especially abhorrent when the goods seem to have (though they may well not have) a particularly inelastic supply and are especially important to physical survival — for example, food, land, and clean air and water.

Belief in the Need for Centralized Control of Important Activities

Hayek (1952) thinks that the belief in centralized control of economic activity in society is a misplaced analogy to the way engineers plan a dam or bridge, and he traces socialist theory back to the creation of the great engineering schools in France at the turn of the 19th century. Whether or not his account of intellectual history is correct, Hayek's analysis of the contemporary sources of the belief in the need for control is sound, I think. Many people believe that without planning and controls, the system just cannot work well. For example, in a debate over whether Champaign County, Illinois, should permit rezoning of farmland for industry, people were heard to say, "I'm for growth, but for *controlled* growth, of course." When you ask them why growth must be controlled by a planner or an agency, they look at you blankly, as if you are lacking in elementary intelligence.

Many seem to fear that chaos is the inevitable result of lack of centralized control. Hayek argues that this belief in the need for control is related to lack of understanding of how a large group of people, acting without any pre-arrangement, can develop an orderly structure of production and exchange based on individual desires and perceptions of other's desires and intentions. He also mentions the common failure to understand the difficulty of organizing an economy nearly as well by central planning, even with the aid of unlimited computing capacity and the most detailed imaginable information

gathering, as with a market. These are subtle ideas, not easy to grasp, so it is not surprising that even well-educated laypersons often have not thought them through and do not understand them.

Two other aspects of a market-directed economy that often are not understood, and whose lack of understanding leads to the call for a directed society and to worry about resources in a market society: a) the capacity of markets to deal with the future; and b) the capacity of correctly structured markets to deal with externalities. There is no space to say more about these matters here, however.

Another possible reason people believe in the need for a centrally-directed society is the belief that *others* who are not so well-educated and intelligent cannot figure out how to conduct self-supporting lives that will also thereby contribute to economy and society. The belief that welfare support will be necessary for immigrants — who are often thought (wrongly) to come to the U.S. with little education and knowledge of English — stems from the arrogance of educated people.

Beckmann and others have suggested that this view fits with intellectuals' desire to be needed by the society, and with their belief that their trained intellects should therefore achieve for them places of special importance and reward in the economy and society. As Beckmann says about a capitalistic society, "The highly skilled jetliner pilot and the lowly cleaner of sewage systems get a reward beyond dollars — the heady knowledge that they are voluntarily supported because they are genuinely needed. Such a reward is unknown to the professor of Turkish medieval poetry" (1978).³ In Western civi-

lization this is an old story with Plato. As Popper put it, Plato "charmed all intellectuals with his brilliance, flattering and thrilling them by his demand that the learned should rule" (1966, p. 199).

Along with this lack of belief in poor people's capacities to run their own lives well is likely to come disbelief that others — and especially the uneducated and poor — can really create resources by way of creating new ideas. Perhaps this disbelief is due in large part to common lack of understanding of how such human intervention lies behind the resources that we take for granted, e.g., the fertile Midwestern prairie that was a malarial swamp before settlers drained it.

Belief That Externalities of Self-interested Actions are Usually Bad

Environmentalists worry that the unintended by-products — the "externalities" — of humankind's economic activities (especially those that affect the environment) are malign even if the direct effects of production and trade can be benign. But I believe a case can be made that even activities that are not intentionally con-

structive usually leave a positive legacy to subsequent generations. That is, even the unintended aspects of humans' use of land (and of other raw materials) tend to be profitable for those who come afterward.

Take as an example the "borrow pits" by the sides of turnpikes, from which earth is taken for road-building. At first the pits seem a despoliation of nature, a scar upon the land. But borrow pits turn out to be useful for fishing lakes and reservoirs, and the land they are on is likely to be more valuable than if the pits had never been dug.

Another example is a garbage dump. Later generations may find dumps profitable sources of recyclable materials. Even a pumped-out oil well — that is, the empty hole — probably has more value to subsequent generations than does a similar spot without a hole. The hole may be used as a storage place for oil or other fluids, or for some as-yet-unknown purposes. And the casing that is left in the dry well might be reclaimed profitably by future generations.

The explanation of this general phenomenon is that humans' activities tend to increase the order and decrease the randomness of nature. We tend to bring like elements together, to concentrate them. This properly can be exploited by subsequent generations. Furthermore, humans perceive order, and create it. One can see this if one looks from an airplane for the signs of human habitation. Where there are people (ants, too, of course) there will be straight lines and smooth curves; otherwise, the face of nature is not neat or ordered.

Many acts that we tend to think of as despoiling the land actually bestow increased wealth upon subsequent generations. Of course this proposition is hard to test. But perhaps a mental comparison will help. Ask yourself which areas in central Illinois will seem more valuable to subsequent generations — the places where cities now are, or the places where farmlands are?

One sees evidence of this delayed benefit in the Middle East. For hundreds of years until recently, Turks and Arabs occupied structures originally built by the Romans 2,000 years ago. The ancient buildings saved the late-comers the trouble of doing their own construction. Another example is the use of dressed stones in locations far away from where they were dressed. One finds the lintels of doorways from ancient Palestinian synagogues in contemporary homes in Syria.

A related trait of mind is appropriate for safety engineers but paralyzes the social will and causes rejection of new tech-

nical possibilities when mis-applied to thinking about natural resources and the environment. This way of thinking focuses only upon the dangers of a projected line of activity, and urges us to "play it safe." When discussing a social scheme, a talented game theorist and I kept disagreeing about whether particular systems would work or not. Finally we discovered that I focused upon the *aggregate* effects on *average*, whereas he focused upon "worst-case analysis," which he said is characteristic of his trade. And worst-case analysis causes one to reject as not attractive many possibilities that on average are desirable. Much of the thinking of the environmental movement seems to be worst-case analysis.

Most puzzling is why people who are themselves creative and imaginative should lack faith in others' capacities to respond to problems and shortages with limit-expanding ideas.

Another analogy in another context: Nathan Leopold, of the Loeb-Leopold murder case, wrote in his fascinating autobiography that it is extraordinarily difficult to persuade prison administrations to accept new ideas for running the prison because they know that a thousand pairs of eyes are looking for the slightest loophole in the new setup which can be exploited for escape or other troublemaking. But as Einstein said about nature, God may be tricky but he (sic) is not malignly trying to do us in. And our situation with respect to resources and the environment is not like that of a prison, and we need not think as do prison administrators or safety engineers.

Nuclear power debates provide many instances of what we might call the Leopold-safety engineers syndrome. Those who are against nuclear power point to scenarios conceivably leading to, say, 50,000 deaths. Proponents of nuclear power point out that the risk of such a scenario occurring is minuscule, and the "expected number of deaths" — using "expected" in the statistical sense — is very small. The anti-nukes are not impressed by such a probabilistic argument, saying that the worst case has a meaning to us that cannot be treated as part of any set of averages. Nor are anti-nukes impressed

by other examples of similarly large worst-case risks that we routinely accept, such as those of power-providing dams that might break and kill hundreds of thousands of people, or airplanes falling from the sky into stadia seating 70,000 people where all might be killed risks that are probabilistically greater than those from nuclear energy. There seems to be a value judgment at the bottom of the argument, a value which cannot be rebutted logically any more than other values can be rebutted logically. But it is possible to point out costs of such policies that are being neglected in the discussion. It is appropriate for a safety engineer not to be concerned with the costs of avoiding a dangerous activity, because the cost/benefit calculation will be made at higher levels of management. But in discussion of such activities as nuclear power, it would seem that all discussants have an obligation to have a balanced view and not just focus on one side of the matter, because there is no arbiter in a court of public opinion who will take into account all sides of the matter, as higher levels of management are responsible for doing in an industrial setting. Also, it seems appropriate to point out in such discussions that if we routinely follow such a line of thought, lives will be shorter and poorer, and fewer people will get a chance to enjoy life, because of the life-shortening effects of air pollution from coal and the industrial accidents that kill so many people in coal-mining and petroleum operations.

The case of hydroponic vegetable growing may sharpen the argument. Hydroponics is now a profitable operation around Washington, D. C. for a good many farmers during the months when vegetables are not grown outdoors nearby (Shelley Davis, "Roots Under Water", *The Washington Post*, April 15, 1984, pp. D1, D4). Hydroponic farming takes up only about one twelfth as much land as does ordinary agriculture, the article points out. Shortage of cropland for growing food is one of the common arguments why population growth should slacken now and must eventually cease. But the mention of hydroponic farming usually evokes a long series of what-if objections. What if there will be a shortage of water? Of chemicals? Of sunlight? Of glass to build greenhouses? And on and on. It is impossible to rule out every imaginative scenario without detailed analysis. And of course there is always the seemingly-unrebuttable objection: This cannot go on *forever*. We would even run out of room on earth for hydroponic farming. (Of course there is plenty of room in space for spaceships carrying hydroponic farms, a

possibility for which the technology is already available without even waiting for further developments. And hydroponic farms can be operated as multi-story plants with artificial light.) Each of these questions is offered as argument against change and growth; the questioner would have us proceed as if hydroponic farming is not a real option.

REFERENCES

Beckmann, Petr, *What Attracts Intellectuals to Socialism?* (Boulder, Colo: Golem Press, Box 1342, Boulder, Colo 80306, 1978).
Bernal, J. D. *Science and Industry in the Nineteenth Century* (London: Routledge and Kegan Paul, 1953/1970).
Ehrlich, Paul R., and Anne H. Ehrlich, *The Population Explosion* (New York: Simon and Schuster, 1990).
Gore, Albert, Jr., *Earth in the Balance*.
Hayek, Friedrich, *The Counter-Revolution of Science* (Glencoe: The Free Press, 1952).
Kristol, Irving, *Two Cheers for Capitalism* (New York: Basic Books, 1978).
Meadows, Donella H.; Dennis L. Meadows; Jorgen Randers, and William W. Bahrens III, *The Limits to Growth*. (New York, Potomac Assoc., 1972).
Meadows, Donella H., Dennis L. Meadows, and Jorgen Randers, *Beyond the Limits* (Post Mills, Vermont: Chelsea Green, 1992).

National Academy of Sciences, *Population Growth and Economic Development* (Washington: National Academy Press, 1986).
Popper, Karl R., *The Open Society and Its Enemies*, 5th ed., (Princeton: PUP, 1966).
Rosenberg, Nathan, *Technology and American Economic Growth* (New York: Harper, 1972).
Schoeck, Helmut, *Envy* (Indianapolis: Liberty Press, 1969).
Simon, Julian L. *The Ultimate Resource* (Princeton: PUP, 1981, 2nd edition forthcoming).
— — — “Does Doom Loom?” *Reason*, April, 1984, pp.31-35;
— — — “Resources, Population, Environment: An Over-supply of False Bad News,” *Science*, 208, June 27, 1980, pp.1431-1437.
— — —, *Population Matters* (New Brunswick, N.J.: Transaction, 1990).
United Nations Fund for Population Activities, *Population, Resources, and the Environment* (New York: UNFPA, 1991).
von Mises, Ludwig, *The Anti-Capitalist Mentality* (South Holland, Ill., Libertarian Press, 1972).

NOTES

1. Much of the unsound thinking about the nature of natural resources and their supply, and about the effect of additional people upon environment, resources, and the living standard, is discussed in my 1981 *The Ultimate Resource*, and therefore will not be repeated here. More generally, I have drawn upon my various writings for words and thoughts contained here

without troubling either reader or writer with quotation notes or citations.

In an important sense, the heart of the economics of population and resources is the kind of thinking that is brought to bear upon the subject — what we might call the “metaeconomics”. The needed kind of thinking — focusing on the indirect, long-run, diffuse influences rather than on the immediate and direct effects — does not excite the mind as do those two old bewitchers, exponential growth and diminishing returns. That perhaps explains why so many persons become and remain bankrupt about population and resources.

2. This section is heavily influenced by Hayek’s works. There also is a fair amount of common ground here with the literature on why people are attracted to socialism (e.g., Kristol, 1978; Beckmann, 1978; Mises, 1972), because resources and environment are part of the “economic problem” that socialism purports to “solve.”

3. Consider this remark about summer work cleaning up garbage from Lake Michigan beaches:

I remember the mornings when the beach was particularly filthy — the Fifth of July was always the worst — and halfway through the job, looking back and seeing only the bare golden sand where before there had been a half-ton of garbage.

I learned that summer the palpable satisfaction of doing a job well, even if that job is picking up garbage. (R. Simon, 1982, p.7)

The Transhuman Taste Reviews of Extropian Interest

Theories of Everything: The Quest for Ultimate Explanation

by John D. Barrow

Ballantine Books, New York, 1991

Paperback \$12.00, ISBN 0-449-90738-4

Reviewed by Peter McCluskey

Theories of Everything is a good exploration of the most fundamental questions of physics, mathematics, and why we exist. It contains a number of the fascinating ideas that were discussed in *The Anthropic Cosmological Principle* (ACP), but is simple enough that it can be understood by people with no mathematical background (it has about ten simple equations, versus enough tensor equations in ACP to scare off most non-physicists). It has enough new ideas so that people who have read ACP will still find significant parts of it interesting.

This is the book that Hawking's *Brief History of Time* should have been. It describes attempts to unify forces, particles, space, time, and the initial conditions of the universe into a single coherent explanation, possibly even a single equation. The term "theory of everything" may be a somewhat exaggerated way to describe a unified theory encompassing the essentials of physics, since it would no more explain the complete configuration of the universe than a perfect theory of thermodynamics would enable us to reliably forecast next year's weather. All it would take is a little bit of chaos to make prediction intractable.

Barrow summarizes the history of comprehensive theories, such as ancient creation myths, and Roger Boscovitch's *Theory of Natural Philosophy* (1758), which proposed a unification of the gravitational, electrical and magnetic forces.

Barrow recounts the historical separation between laws describing the evolution of the universe and the initial conditions that are treated as arbitrary inputs to the equations in those laws, and describes hints that a theory of quantum gravity would abolish that distinction, along with the distinction between time and space at the beginning. Since time is not an explicit component of quantum theories of cosmology, time can be regarded as a phenomenon which applies only to some parts of the universe. Closer to what we think of as the origin (English is not well

suited for handling directions when they are not readily divided into space and time), a model of 4 spatial dimensions could be a more useful way of describing things. In this case, the problem of explaining the "beginning" of the universe fades away, as there is no starting point to such a universe.

In ACP, Barrow and Tipler had appeared to prove that the universe had a finite past by showing that if intelligent life had ever existed in our infinite causal past, it would have taken over all that we can see. Barrow now describes a theory of an "eternal" inflationary universe in which regions of a universe which is infinite in both time and space can undergo an inflationary expansion of the kind that is postulated to have occurred shortly after the "big bang" in order to explain the large-scale uniformity of the visible part of the universe. Each such expanded region sees a dense big-bang-like past. As long as life cannot evolve in the superhot unexpanded

regions of the universe, it can exist in infinitely many inflated regions without being able to overrun newly created ones because the rapid expansion in the early phases of each region insures that the lightspeed barrier isolates each region.

If even the infinite spatial part of this idea is correct, then our search for universal theories is seriously hampered by the likelihood that the visible portion of the universe has been selected for its ability to evolve life, and is therefore probably very unrepresentative of the conditions that a truly general theory would say are typical. So what appear to be universal laws may be special cases of something much more general.

There are some numbers such as the ratio of the proton mass to the electron mass, which appear at first glance to be totally arbitrary, and therefore impossible to reduce to a principle derived from some more basic law. One of the ways they could have been determined, Barrow suggests, is by the effects of numerous small wormholes (little larger than the Planck length) connecting our region of the universe with distant parts of itself or with other regions which would otherwise be

independent universes. The effects of quasi-random connections of this nature would reportedly be to introduce random shifts in the values of the "constants" observed within each "normal" region. One example is the cosmological constant (a term Einstein added to General Relativity to allow the theory to predict a stationary rather than expanding or contracting universe), which is expected to be extremely close to zero for most possible wormhole interconnections.

Barrow expresses concern about the possibility that a Theory of Everything could be shown to be the only logically possible theory, and thus be an analytic truth, rather than a scientifically falsifiable truth. I find this to be disappointing, as deductively proven truths have withstood the test of time better than theories which are subject to falsifiability. While I doubt that such an analytic truth will be found, I consider falsifiability to be a crutch to use when nothing better is available, and want analytic truths to replace unprovable theories wherever possible.

In Our Own Image: Building an Artificial Person

by Maureen Caudill

Oxford University Press, 1992

242 pp hardcover; \$22.00

ISBN 0-19-507338-X

Reviewed by Derek Zahn

Books on Artificial Intelligence (AI) written for general audiences usually proceed either historically — tracing the development of the major AI ideas by focusing on the scientists responsible for them — or by presenting a case for the centrality of particular techniques or problems. In *In Our Own Image: Building an Artificial Person*, Maureen Caudill takes a different path; using the popular media image of The Android as a goal to be reached, she explores the problems that must be solved to achieve that goal. The exploration visits most of the research areas of AI, resulting in a satisfying and readable summary of the state of the art.

The book has two sections; the first concentrates on techniques and research; the second looks at more philosophical issues. Of these, the first is far superior. She focuses in turn on the different capabilities that an android must have — sight, movement, memory, reasoning, speech, etc, and evaluates the methods that AI has developed to give machines those capabilities.

Despite the broad scope of such a project, in less than 200 pages Caudill covers the territory skillfully, from the early insights into reasoning and logic, through the focus on knowledge representation, to the more recent emphasis on learning and parallel-processing alternatives to hard-wired computer programs. The explanations are necessarily only overviews, but they are very clear and easy to follow. For example, the neural network architectures developed by Carpenter and Grossberg, which embody their Adaptive Resonance Theory, are notoriously difficult to understand — even for AI researchers. Caudill does a superb job of making their operating principles clear. In fact, her explanations of various neural network techniques were for me the strongest parts of the book.

As the chapters unfold, one of the central tensions in AI is revealed: the contrast in approach between traditional "symbolic" AI — in which aspects of the world are represented as discrete abstract symbols which are then operated on by the computer (Expert Systems exemplify this approach) — and "connectionist" AI, which focuses on adaptable, "subsymbolic" processing mod-

elled loosely on the way the brain works. Caudill believes that these approaches can and will be merged in the quest for a functioning android.

I had a few negative reactions as well, though. First, her discussions of some of the newest directions that AI has taken — the use of evolutionary techniques, for example, are divorced from the main text and included in her philosophical speculations, without discussing how these emerging techniques will help achieve the project's goals. Other hot topics are not even mentioned, for instance: modelling computation as a free market, the Rodney Brooks-led turn away from the functional decomposition of intelligence, Minsky's *Society of Mind* model, and "distributed" AI, in which problem-solving takes place in the interaction between semi-autonomous agents who negotiate with one another. Also, the interesting engineering issues concerning the construction of an android's musculature, skin, and other "hardware", are not reported — which is appropriate for an

introduction to AI but is a sad oversight for a book putatively about androids.

My most severe disappointment was Caudill's failure to address the integration issue. AI researchers have divided cognition up into parts along much the same lines as Caudill, and (as she points out) the progress has been promising. But a central puzzle that AI is trying (and failing, currently) to solve is how to put the pieces together. It seems that many of the hardest issues have somehow fallen between the cracks of the problem division. Caudill would have served her readers better by reporting this in a straightforward way, even though it might have tempered her optimistic prognosis for rapid progress along well-defined paths.

The second (much shorter) segment of the book takes a rather cursory look at the philosophical and social consequences of androids. She hits the main obvious points: What does it mean for something to be alive? How will society react to intelligent machines? Should they have rights? Would you let your daughter marry one? How will we feel when they surpass us? But the analysis is rather mechanical, and is completely divorced from the concurrent changes that will be sweeping society as other technologies are developed in the coming decades. Some readers who have been living in caves may not have seen these issues before, but most have — especially those who will actually buy her book. I'd suggest that the movie *Blade Runner* and a couple episodes of *Star Trek* do at least as good a job of making you think about the ramifications of this fundamentally important component of our future. Much written science fiction does even better.

But that is a minor part of her book. For readers who are largely mystified by the methods and issues of AI, I firmly recommend *In Our Own Image*, on the strength of Caudill's considerable talent for clear explanation. In fact, were I teaching an introductory AI course, I would consider assigning this book as a first week's introductory reading. However, if you have taken such a course or have independently followed AI, much of this book will probably be rehashes of things you already know (though the neural net material still may make it worthwhile).

These particular strengths and weaknesses aside, Caudill has done a service to Extropians with this book, by bringing the challenge more into public view. The creation of Androids will increase the intelligence that can be applied to progress in all endeavors, but additionally — as the book's title suggests — it will help us understand the human mind and so empower our own self-transformations in ways as yet unforeseen.

Mirror Worlds — Or the day software puts the Universe in a shoebox... How it will happen and what it will mean

by David Gelernter.

Oxford University Press, 1992. New York.

ISBN: 0-19-506812-2, 0-19-507906-X

Reviewed by Harry Shapiro Hawk

Extropians will find Gelernter's views appealing in many respects. He writes about how groups of uncoupled entities can communicate and work together by choice. That is certainly a metaphor for an idealized Extropian community.

Mirror Worlds is Gelernter's prediction of a software revolution that will allow us to enter into a real-time, dynamic model of our world. We can enter it for pleasure and profit. We can send our software agents to do our bidding. It is like the box Gibson's character 'Bobby' is plugged into; a simulation of our Universe running forward and in real-time. Gelernter contends "[it] will allow us to explore the world in unprecedented depth and detail without leaving the comfort of home."

Mirror Worlds is also a technical primer for creating highly modular software for parallel and massively parallel computers. It primes us with Gelernter's view of computer architecture and how that can represent a "fine art form" in its own right. Gelernter's views have historically been unorthodox, especially his strong believe in parallelism. You can expect to find his views refreshing, even if you don't completely agree with them. Like the clearly more Extropian-minded, Hans Moravec, Gelernter presents an interesting thought experiment, based on his actual research, of where he feels the future of computer science and our society will go.

The prologue starts the book appealingly, Gelernter poetically writing, "technology is the ocean on a bright cool Spring day. Sparkling in the far distance; breath-takingly cold; exhilarating once you've plunged in... The cold ocean is coming to meet you," and extols, "Why not give it a try? Hold your breath. Let's plunge."

The first chapter is devoted to explaining, in a non-technical way, what a simulation like his proposed *Mirror Worlds* would be like. He compares the view of software in 1991 to the view of technology in 1791:

"People were pretty sure, in 1791, that the industrial revolution had 'happened'. It was history... In retrospect, little had changed... In 1791, the industrial revolution was merely building up a head of steam... Glancing backwards from a vantage-point two centuries hence, 1991 will look a lot like 1791. The real software revolution won't have much to do with fancy robots, computers in education, ... or the other hot topics that dominate this month's hit parade. It will have to do with... *Mirror Worlds*... Today, software as a building material resembles mosaic tile. In the future, software will metamorphose into a something more like stone or steel or concrete. The metamorphosis has in fact (just) begun."

He then spends the remainder of the book writing about creating and using software tools. These tools are: Tuples which are passive data, Infomachines which are tuples actively running "their" programs, Ensembles which are groups of Infomachines, Trellises, which are layers of specialized Ensembles, and finally FGP machines (F=Fetch, G=Generalize, P=Project) which are collections of multiple Trellises being manipulated to extract inference and conclusions from a vast sea of tuples which he calls Tuplespace. When he says vast sea of tuples he really means *vast*. He imagines every sensor like I.C.U medical equipment and every datum like a bid or an ask on some trading floor being available.

I first heard about David Gelernter when reading an article he had written for the August, 1989 issue of *Scientific American* [Vol. 261; No. 2; Pg. 66]. He was explaining a new method for writing programs that relied upon "anonymous uncoupled communication," where each "component that produces data need not know who will use it or when" and "components that require data need not know who produces it." Gelernter had developed such a system while a graduate

student; he called it Linda. Linda allows multiple processes to run and share data across multiple CPUs.

In the *SciAm* article Gelernter wrote:

Linda programs inhabit it what we call 'tuple space.' (A tuple is a chunk of data; the term is a generalization of terms such as quadruple and quintuple.) Passive tuples are just data available for reading or processing. Active tuples are subprograms, all executing simultaneously, that consume and produce other tuples. Active tuples turn into passive ones, available for reading or processing, once they have finished executing.

He depicted a tuple space structured into a "trellis: a row of modules at the bottom connected to sensors in the real world, a second, higher row to refine the data and make connections between different items, a third row for further refining and so on. Two-way communication between rows permits the lower-level modules to alter their actions in response to queries or comments passed down from upper levels..."

With many trellises connected into different parts of the real world: hospitals, traffic systems, your home's heating system, the FAA flight control database, and many more, at the highest level of abstraction you would have a software model of the "universe." That is a Mirror World.

Gelernter has clearly refined his thinking since the 1989 *SciAM* publication. He also presents work by various graduate students of his, who have been actually building the types of systems he describes. He presents the work of Researcher Scott Fertig who created a FGP environment for radiology diagnosis. This system with a very limited database of 70 cases was able to correctly determine that 1) a new case was unlike anything it had seen before and 2) the breast tumor in question was malignant.

In this radiology example, there was a human feeding in data to the computer. Gelernter sees a different type of human interaction in future systems. He sees a human plugged into the top level of a trellis. He even shows a illustration with a bunch of floating heads above and seemingly connected to a trellis. For any proto-posthuman, this is an example of how a human consciousness could retain its identity, do profitable work and be connected to millions of other minds through a vast data-scape (Gelernter's Tuplespace). Gelernter writes:

A Trellis, it turns out is like a crystal... When you turn it on it vibrates at a certain frequency... In concept each

Trellis element is an Infomachine. All these Infomachines run separately and simultaneously... In practice, we do things somewhat differently... Workers collaborate to make the whole thing work predictably. Predictability is crucial. We run the Trellis in a series of sweeps... we instruct each worker to run through its list... sticking with each one just long enough... When it is done, it waits until all the rest have finished. Then all the workers proceed into the next sweep. Hence the 'frequency' of a Trellis... In a fast Trellis, sweeps are short; the frequency is high... In a slow Trellis, the opposite... It's easy to imagine a Trellis that includes human elements alongside the software ones. In a... Trellis, lower rungs act 'instinctively.' Higher levels look for the big picture... In these areas, we could use people to realize some of the higher-rung elements... At some level of the hierarchy, human elements start to intermingle with software ones.

Gelernter calls such combinations, Turingware.

What is largely missing from the book is an in-depth focus on the agoric aspects of the computer models Gelernter presents. He does briefly mention Bernardo Huberman's work at Xerox. I would like to see more about how bits of tuples are bought, sold and traded. How FGP machines will evolve a 'correct' ratio between buying expensive data versus the cost of building complex inference engines. However, despite the lack of such agoric concepts I don't see Gelernter's overall architecture precluding them.

In one other general sense is the overall work non-Extropian. For example, early in the book Gelernter spends time writing how a Mirror World connected into the heart of a democratic government will allow citizens of that democracy actually seek out and find how it is actually working and thus make it work better. I see such systems working in the opposite direction; namely showing why a strong central government can never work. Of course in an idealized Extropian PPL¹ society such a system would allow each community member to monitor, to what ever degree desired, how well each citizen was meeting their contractual obligations.

While not mentioning memes, Gelernter does write about people's fear of the unknown. He, most likely, unknowingly tells how Mirror Worlds will fill the receptor site for that fear, "I've claimed that Mirror Worlds are a development of a large potential importance. This is why... They will make the world

run better and smoother... My guess is that, by offering topsight to the millions (not merely to the visionaries who have monopolized it in the past), they speak directly to the large, perpetually unsatisfied human craving to understand 'what's going on,' to see things whole. For 'reasons' that transcend the rational, they will be hard to resist." I think that memetics adds enough reason so that we can rationally understand why humans want to understand the 'whole.' I fear Gelernter means that such a better understanding of the whole will lead to better governments rather than simply better societies.

Nevertheless, Gelernter sees that Mirror Worlds will force people to model and interact with accurate rational models of the world. This is something that too many people don't do. If Mirror Worlds can do this, then I think we will be working our way towards a better society.

If these ideas intrigue you or you want more technical details, I strongly recommend *Mirror Worlds* as a good starting point. Since Linda is commercially (and competitively) available from several sources, for those inclined to tinker, these ideas can soon become very real.

¹Privately Produced Law. See the article of this title by Tom Morrow in *Extropy* #7.

Call For Papers:

Extro 1:

The First Extropy Institute Conference on Transhumanist Thought

San Francisco, California, April 30 - May 1, 1994

Extropianism is a transhumanist philosophy: Like humanism it values reason and sees no ground for believing in supernatural external forces controlling our destiny. But transhumanism goes further in calling us to push beyond the simply human stage of evolution. Where others see difficulties, we see challenges. Where others give up, we move forward. Where others say enough is enough, we say: Forward! Upward! Outward! We espouse personal, social, and technological evolution into ever higher forms. Extropy Institute, and its publication *Extropy: The Journal of Transhumanist Thought*, seeks to promote and develop these ideas, which are summarized in the five Extropian Principles:

- Boundless Expansion
- Self-Transformation
- Dynamic Optimism
- Intelligent Technology
- Spontaneous Order

Extro 1 will be a rich, intellectually invigorating gathering designed to help push outward the boundaries of progress and possibility. It will be both a serious study and a joyful celebration of humanity's limitless potential and how it will be achieved. Besides presentations of accepted papers, the conference will feature lectures by leading thinkers, panel discussions, the first Extropy Awards banquet, and other events.

Submitted papers should as much as possible exploit interdisciplinary connections, rather than presenting results in a particular narrow sub-field. They should be aimed at an intelligent, educated and interested audience that is not necessarily familiar with the detailed background

of any field. Necessary mathematical formalism is encouraged, but detailed proofs of marginally significant results are not. Accuracy, rigor, and rationality are of course expected, but breadth of vision is also important.

The following is a non-exhaustive list of appropriate topics for papers, to give an idea of what the conference is all about:

- ◆ Polycentric law and voluntaristic, market-based societies
- ◆ Neuroscience and the transfer of consciousness into non-biological media (uploading)
- ◆ Artificial Intelligence, genetic algorithms, and neural networks
- ◆ Futurism
- ◆ Nanotechnology
- ◆ Cryptography and technological aids to privacy
- ◆ Logic and guides to effective thinking, such as General Semantics, fuzzy logic
- ◆ Biostasis
- ◆ Philosophy of mind, self, and identity
- ◆ Space travel, exploitation, and habitation
- ◆ Virtual reality
- ◆ Life-extension
- ◆ Biological and neurological augmentation
- ◆ Idea Futures, Hypertext, and other information technologies
- ◆ Transhumanist eupraxophy
- ◆ Self-organization and complexity
- ◆ Artificial life

Papers must be written in English, must be 5,000-12,000 words in length, and must begin with an abstract of not more than 400 words. Papers must include a separate cover page (not part of the paper itself) containing the title, author, postal address, and email address if available. Submissions should not have been previously published or submitted to any journals or refereed conferences or workshops. Accepted papers must be presented at the conference.

Submission deadline is December 15, 1994. Authors will be notified of review decisions by January 15, 1994. Camera ready copies of accepted papers are due back by February 15, 1994 for inclusion in the Conference proceedings.

Please mail four (4) copies of papers to:
Extropy Institute
Extro 1 Conference
11860 Magnolia Avenue, Suite R
Riverside, CA 92503

Questions can be directed to derek@cs.wisc.edu.

Proceedings will be available to conference attendees at the conference, and will be available afterwards from Extropy Institute. This call for papers and other information about Extro 1 can be retrieved via anonymous ftp from [lynx.cs.wisc.edu](ftp://lynx.cs.wisc.edu/pub/Extro-1) in the directory `pub/Extro-1`. Also included there are the full text of the Extropian Principles, and information about Extropy Institute and the journal *Extropy*.

Extro 1 Conference Team

David Krieger, American Information Exchange
Max More, Extropy Institute
Derek Zahn, University of Wisconsin

Forward! Upward! Outward!

classified ads:

Cross verbal swords in *The (Libertarian) Connection*, open-forum magazine since 1968. Subscribers may insert one page/issue free, unedited. Lots of stimulating conversation. Eight issues (one year) \$20. Strauss, Box 3343X, Fairfax, VA 22038.

For sale: **Laserjet II Plus** 300dpi laser printer, with Postscript Pacific Page cartridge. Well under a year old. \$500 or best offer. Contact the Extropy Institute office: 909-688-2323

Results of ExI Pledge Drive, held on the Extropians e-mail list from October 21-30. (Originally reported in *Exponent* #3) Our thanks again: Without you, this issue of *Extropy* probably wouldn't exist.

Karl Waldman	\$300**
"CAW member & Extropian"	\$300
Harry Shapiro	\$240+
T.O. Morrow	\$100
Eirikur Hallgrimmsson	\$100
Sam Shipman	\$100
David Nelson	\$100
Chris Moriondo	\$100
Lee Nussbaum	\$80++
Perry Metzger	\$60
Tim May	\$60
Bob Grahame	\$53
Peter C. McCluskey	\$50
Mark Desilets	\$50
Mark Venture	\$50
Fred Moulton	\$50
Alexander Chislenko	\$50
Kennita Watson	\$50
Richard Kennaway	\$50
Anthony Garcia	\$50
Carl Feynman	\$50
Steve Whitrow	\$50
Bill Eichman	\$30
Scott Meeks	\$30
Robert Brooks	\$30
Anonymous	\$20
Anonymous	\$20
Mike Linksvayer	\$20
Kevin Q. Brown	\$20
Andrea Gallagher	\$20
Chris Rasch	\$20

**pledge of \$25/month for one year
+matching grants totalling \$20/month
++pledge of \$50 now and \$30 in March

Extropy#11 was produced on a Gateway 486DX2 50 with 8Mb of RAM, a 230Mb hard disk, 14" Crystal Scan monitor, using Pagemaker 5.0, Word for Windows 2.0 and Aldus Freehand 3.1. Proofs were printed at 600dpi on an HP Laserjet 4 with 6Mb RAM. Layout by Max More.
This issue was printed on a web press by Canyon Printing, Inc., Anaheim, CA.
Printrun: 3,200

CONTRIBUTORS

Harry Shapiro Hawk is the Manager of Computer services at Warwick Baker & Fiore Advertising. He holds a Master degree in Interactive Telecommunication from New York University. He is the Electronic Communications Officer for Extropy Institute and manages the Internet based Extropians mailing list. habs@panix.com

David Krieger: David is Marketplace Administrator at the American Information Exchange (AMIx), Science Editor for *Extropy*, an Extropy Institute Director, and a former Technical Consultant to *Star Trek: The Next Generation*. dkrieger@netcom.com

Peter McCluskey: Peter received a B.S. in Biology from Yale in 1978, and an M.S. in Computer Science from Brown in 1993. His master's thesis compared several neural network algorithms and genetic programming for predicting the stock market and other time series. He is currently designing financial software for the MacGregor Group in Boston. pcm@macgreg.com

Ralph Merkle: Dr. Merkle received his Ph.D. from Stanford University in 1979 where he co-invented public key cryptography. He pursued research in this area at Bell Northern Research until he joined Elxsi (a computer start-up) in 1981. He left Elxsi to join Xerox PARC in 1988, where he is pursuing research in computational nanotechnology. Dr. Merkle has six patents and has published extensively. merkle.pa@xerox.com

Max More: Editor of *Extropy* and President of Extropy Institute, Max received his BA in Philosophy, Politics, & Economics from Oxford University, and now is writing his Ph.D dissertation at the University of Southern California on "The Diachronic Self: Identity, Continuity, Transformation." more@usc.edu or 76436.3157@compuserve.com

Mike Price: BSc and MSc in Theoretical Physics from Imperial College, London. Specialist interests in cosmology and quantum field theories. Freelance systems engineer, software engineer, and project manager. Next issue joins *Extropy's* Editorial Committee as Theoretical Physics Editor. price@price.demon.co.uk

Julian L. Simon: Perhaps call him Senior Junior — teaches business administration at the University of Maryland. He studies the economics of population, and is now exploring and promulgating the resampling method of doing all statistics — a method (including the bootstrap) that he developed a quarter century ago. He invented and promoted the volunteer scheme for handling airline oversales, from which you may have benefitted.

Ralph Whelan: Vice President of Alcor Life Extension Foundation, Editor of Alcor's *Cryonics* magazine, and a director of Extropy Institute. His main interest is music, which he views as a tremendously precise and complex form of thought/communication. His biggest distraction is computer art/animation 71532.2442@compuserve.com

Derek Zahn: Derek Zahn recently escaped the higher educational system, fortunately only suffering a master's degree in Computer Science at the University of Wisconsin. He is presently Principal Optimist at Luminous Software, and maintains the ExI-Essay FTP archive. derek@cs.wisc.edu

ExI Audio Tapes, Books, T-Shirt

Everything is Getting Better and Better – I'll Bet On It!

by Julian Simon.

Economist Julian Simon uses hard data to counter prevailing gloomy beliefs about the current state of the world and its direction. Practically all measures of human well-being substantiate the Extropian's dynamically optimistic views: Life does tend to improve, though only through the efforts and applied intelligence of free persons. This tape makes an effective introduction to Simon's ideas, and, lent out to your pessimistic friends, will serve as a valuable intellectual catalyst.

\$10.95 (Members \$9.95)

EC1 (1-hour audio)

Bionomics On Trial: A Discussion With Michael Rothschild

Rothschild outlines the main contentions of his book *Bionomics: Economy As Ecosystem*, and responds to audience questions. Topics discussed include electronic ecosystems; how bionomics effectively draws 'liberals' into support for free markets; the relation between Austrian/process economics and bionomics; the role of government; how far the economy as ecosystem analogy can be pushed.

\$12.95 (Members: \$11.50) EC2 (80 minutes audio)

Postage:

\$1 per tape. Overseas orders: Surface mail – \$1.50 first tape, \$1.25 each additional tape. Contact ExI for airmail rates.

BOOKS AVAILABLE (postage rates in parentheses):

Maureen Caudill, *In Our Own Image: Building An Artificial Person* \$23.00 (\$2.00)

K. Eric Drexler, *Nanosystems: Molecular Machinery, Manufacturing, and Computation* \$24.95 (\$3.00)

Alan Lakein, *How to Get Control of Your Time and Your Life* \$4.95 (\$1.00)

Marc Steigler, *The Gentle Seduction* \$3.50 (\$1.20)

Vernor Vinge, *A Fire Upon the Deep* \$5.99 (\$1.50)

Forward! Upward! Outward! Into your T-Shirt!

Strike terror into the hearts of gloomy pessimists, life-haters, and statist!

Proudly display your transhumanist colors! Yes, Extropy T-shirts are once again available in both men's Large and eXtra-Large. These shirts, in blue, feature the five-spiral Extropy logo in black and gold up front, and the motto "Forward! Upward! Outward! Into the Future!" on the back.

The cost is \$16 per shirt (\$14 for ExI members), plus \$1 postage.

Check or money order in US dollars drawn on a US bank, payable to "Extropy Institute." Mail your order to: Extropy Institute, Dept. S, 11860 Magnolia Avenue, Suite R, Riverside, CA 92503.

How to Join the ExI Virtual Community

Extropy Institute sponsors, through the endeavor of ExI Electronic Communications Officer Harry S. Hawk, a number of electronic fora for sharing libertarian, life-extensionist, pro-technology and other Extropian ideas with bright, like-minded individuals around the globe. The rapid evolution of these fora means that many of the details below have changed since last issue.

The most popular service is the Extropians e-mail list, which boils over with lively discussion and debate on numerous topics. To join, send a request to:

extropians-request@extropy.org

When sending your add request, indicate whether you want real time or digest mode. (If unsure, try the digest mode first!) The Extropians list is using the most advanced information-filtering software, allowing you to select which messages (topic, author, etc.) you receive and how you receive them.

There is also an Extropian conference on the Well, one of the longest-running professionally run BBS systems. On the Well, send mail to habs.

Another service is the ExI Essay list, for posting longer, more carefully prepared electronic manuscripts. To get on this list, send a message to:

exi-essay-request@gnu.ai.mit.edu

Those ExI-Essay papers containing explicit notices granting permission for redistribution are available by anonymous FTP at [lynx.cs.wisc.edu](ftp://lynx.cs.wisc.edu) (IP address 128.105.2.216). A list of available essays along with their file names is in the file `pub/ExI-Essay/INDEX`. Any questions should go to Derek Zahn at derek@cs.wisc.edu.

There are also two "local" lists for announcements and discussions in the San Francisco Bay Area and the Boston area. To join these lists, send messages to:

exi-bay-request@gnu.ai.mit.edu

exi-bos-request@gnu.ai.mit.edu