Edge

http://www.edge.org

Life: What A Concept!

An *Edge* Special Event at Eastover Farm

Freeman Dyson

J. Craig Venter

George Church

Robert Shapiro

Dimitar Sasselov

Seth Lloyd

John Brockman, editor

LIFE: WHAT A CONCEPT!
An *Edge* Special Event at Eastover Farm

Freeman Dyson - J. Craig Venter - George Church
Robert Shapiro - Dimitar Sasselov - Seth Lloyd

John Brockman, editor

Published by EDGE Foundation, Inc.,
5 East 59th Street, New York, NY 10022

EDGE Foundation, Inc. is a nonprofit private operating foundation under Section 501(c)(3) of the Internal Revenue Code.

TABLE OF CONTENTS

INTRODUCTION

“Life Consists of propositions about life."
— Wallace Stevens ("Men Made out of Words")

In April, Dennis Overbye, writing in *The New York Times* "Science Times", broke the story of the discovery by Dimitar Sasselov and his colleagues of five earth-like exo-planets, one of which "might be the first habitable planet outside the solar system".

At the end of June, Craig Venter has announced the results of his lab's work on genome transplantation methods that allows for the transformation of one type of bacteria into another, dictated by the transplanted chromosome. In other words, one species becomes another. In talking to /Edge /about the research, Venter noted the following:

> Now we know we can boot up a chromosome system. It doesn't matter if the DNA is chemically made in a cell or made in a test tube. Until this development, if you made a synthetic chromosome you had the question of what do you do with it. Replacing the chromosome with existing cells, if it works, seems the most effective to way to replace one already in an existing cell systems. We didn't know if it would work or not. Now we do.

This is a major advance in the field of synthetic genomics. We now know we can create a synthetic organism. It's not a question of 'if', or 'how', but 'when', and in this regard, think weeks and months, not years.

In July, in an interesting and provocative essay in *New York Review of Books* entitled "Our Biotech Future", Freeman Dyson wrote:

> The Darwinian interlude has lasted for two or three billion years. It probably slowed down the pace of evolution considerably. The basic biochemical machinery of life had evolved rapidly during the few hundreds of millions of years of the pre-Darwinian era, and changed very little in the next two billion years of microbial evolution. Darwinian evolution is slow because individual species, once established evolve very little. With rare exceptions,

> Darwinian evolution requires established species to become extinct so that new species can replace them.
>
> Now, after three billion years, the Darwinian interlude is over. It was an interlude between two periods of horizontal gene transfer. The epoch of Darwinian evolution based on competition between species ended about ten thousand years ago, when a single species, Homo sapiens, began to dominate and reorganize the biosphere. Since that time, cultural evolution has replaced biological evolution as the main driving force of change. Cultural evolution is not Darwinian. Cultures spread by horizontal transfer of ideas more than by genetic inheritance. Cultural evolution is running a thousand times faster than Darwinian evolution, taking us into a new era of cultural interdependence which we call globalization. And now, as Homo sapiens domesticates the new biotechnology, we are reviving the ancient pre-Darwinian practice of horizontal gene transfer, moving genes easily from microbes to plants and animals, blurring the boundaries between species. We are moving rapidly into the post-Darwinian era, when species other than our own will no longer exist, and the rules of Open Source sharing will be extended from the exchange of software to the exchange of genes. Then the evolution of life will once again be communal, as it was in the good old days before separate species and intellectual property were invented.

It's clear from these developments, as well as others, that we are at the end of one empirical road and ready for adventures that will lead us into new realms.

This year's Annual *Edge* Event took place at Eastover Farm in Bethlehem, CT on Monday, August 27th. Invited to address the topic "Life: What a Concept!" were Freeman Dyson, J. Craig Venter, George Church, Robert Shapiro, Dimitar

Sasselov, and Seth Lloyd, who focused on their new, and in more than a few cases, startling research, and/or ideas in the biological sciences.

Physicist **Freeman Dyson** envisions a biotech future, which supplants physics and notes that after three billion years, the Darwinian interlude is over. He refers to an interlude between two periods of horizontal gene transfer, a subject explored in his abovementioned essay.

Craig Venter, who decoded the human genome, surprised the world in late June by announcing the results of his lab's work on genome transplantation methods that allows for the transformation of one type of bacteria into another, dictated by the transplanted chromosome. In other words, one species becomes another.

George Church, the pioneer of the Synthetic Biology revolution, thinks of the cell as operating system, and engineers taking the place of traditional biologists in retooling stripped down components of cells (bio-bricks) in much the vein as in the late 70s when electrical engineers were working their way to the first personal computer by assembling circuit boards, hard drives, monitors, etc.

Biologist **Robert Shapiro** disagrees with scientists who believe that an extreme stroke of luck was needed to get life started in a non-living environment. He favors the idea that life arose through the normal operation of the laws of physics and chemistry. If he is right, then life may be widespread in the cosmos.

Dimitar Sasselov, Planetary Astrophysicist, and Director of the Harvard Origins of Life Initiative, has made recent discoveries of exo-planets ("Super-Earths"). He

looks at new evidence to explore the question of how chemical systems become living systems.

Quantum engineer **Seth Lloyd** sees the universe as an information processing system in which simple systems such as atoms and molecules must necessarily give rise complex structures such as life, and life itself must give rise to even greater complexity, such as human beings, societies, and whatever comes next.

A small group of journalists interested in the kind of issues that are explored on *Edge* were present: Corey Powell, *Discover*, Jordan Mejias, *Frankfurter Allgemeine Zeitung*, Heidi Ledford, *Nature*, Greg Huang, *New Scientist*, Deborah Treisman, *New Yorker*, Edward Rothstein, *New York Times*, Andrian Kreye, *Süddeutsche Zeitung*, Antonio Regalado, *Wall Street Journal*. Guests included Heather Kowalski, The J. Craig Venter Institute, Ting Wu, The Wu Lab, Harvard Medical School, and the artist Stephanie Rudloe. Attending for *Edge*: Katinka Matson, Russell Weinberger, Max Brockman, and Karla Taylor.

We are witnessing a point in which the empirical has intersected with the epistemological: everything becomes new, everything is up for grabs. Big questions are being asked, questions that affect the lives of everyone on the planet. And don't even try to talk about religion: the gods are gone.

Following the theme of new technologies = new perceptions, I asked the speakers to take a third culture slant in the proceedings and explore not only the science but the potential for changes in the intellectual landscape as well.

We are pleased to present the transcripts of the talks and conversation. A photo album and streaming video clips are available on the Edge website:

http://www.edge.org/documents/life/life_index.html

John Brockman
New York
January 7 2008

FREEMAN DYSON

> *The essential idea is that you separate metabolism from replication. We know modern life has both metabolism and replication, but they're carried out by separate groups of molecules. Metabolism is carried out by proteins and all kinds of other molecules, and replication is carried out by DNA and RNA. That maybe is a clue to the fact that they started out separate rather than together. So my version of the origin of life is that it started with metabolism only.*

Freeman Dyson is professor of physics at the Institute for Advanced Study, in Princeton. His professional interests are in mathematics and astronomy. Among his many books are *Disturbing the Universe, Infinite in All Directions Origins of Life, From Eros to Gaia, Imagined Worlds, The Sun, the Genome*, *and the Internet*, and most recently *A Many Colored Glass: Reflections on the Place of Life in the Universe.*

FREEMAN DYSON: First of all I wanted to talk a bit about origin of life. To me the most interesting question in biology has always been how it all got started. That has been a hobby of mine. We're all equally ignorant, as far as I can see. That's why somebody like me can pretend to be an expert.

I was struck by the picture of early life that appeared in Carl Woese's article three years ago. He had this picture of the pre-Darwinian epoch when genetic information was open source and everything was shared between different organisms. That picture fits very nicely with my speculative version of origin of life.

The essential idea is that you separate metabolism from replication. We know modern life has both metabolism and replication, but they're carried out by separate groups of molecules. Metabolism is carried out by proteins and all kinds of small molecules, and replication is carried out by DNA and RNA. That maybe is a clue to the fact that they started out separate rather than together. So my version of the origin of life is it started with metabolism only.

You had what I call the garbage bag model. The early cells were just little bags of some kind of cell membrane, which might have been oily or it might have been a metal oxide. And inside you had a more or less random collection of organic molecules, with the characteristic that small molecules could diffuse in through the membrane, but big molecules could not diffuse out. By converting small

molecules into big molecules, you could concentrate the organic contents on the inside, so the cells would become more concentrated and the chemistry would gradually become more efficient. So these things could evolve without any kind of replication. It's a simple statistical inheritance. When a cell became so big that it got cut in half, or shaken in half, by some rainstorm or environmental disturbance, it would then produce two cells which would be its daughters, which would inherit, more or less, but only statistically, the chemical machinery inside. Evolution could work under those conditions.

LLOYD: These are naturally occurring lipid membranes?

DYSON: Yes. Which we do know exist. That's stage one of life, this garbage bag stage, where evolution is happening, but only on a statistical basis. I think it's right to call it pre-Darwinian, because Darwin himself did not use the word evolution; he was primarily interested in species, not in evolution as such.

Well then, what happened next? Stage two is when you have parasitic RNA, when RNA happens to occur in some of these cells. There's a linkage, perhaps, between metabolism and replication in the molecule ATP. We know ATP has a dual function. It is very important for metabolism, but it also is essentially a nucleotide. You only have to add two phosphates and it becomes a nucleotide. So it gives you a link between the two systems. Perhaps one of these garbage bags happened to develop ATP by a random process. ATP is very helpful to the metabolism, so these cells multiplied and became very numerous and made large quantities of ATP. Then by chance this ATP formed the adenine nucleotide, which polymerized into RNA. You had then parasitic RNA inside these cells, forming a

separate form of life, which was pure replication without metabolism. RNA could replicate itself. It couldn't metabolize, but it could grow quite nicely.

Then the RNA invented viruses. RNA found a way to package itself in a little piece of cell membrane, and travel around freely and independently. Stage two of life has the garbage bags still unorganized and chemically random, but with RNA zooming around in little packages we call viruses carrying genetic information from one cell to another. That is my version of the RNA world. It corresponds to what Manfred Eigen considered to be the beginning of life, which I regard as stage two. You have RNA living independently, replicating, traveling around, sharing genetic information between all kinds of cells. Then stage three, which I would say is the most mysterious, began when these two systems started to collaborate. It began when the invention of the ribosome, which to me is the central mystery. There's a tremendous lot to be done with investigating the archaeology of the ribosome. I hope some of you people will do it.

Once the ribosome was invented, then the two systems, the RNA world and the metabolic world, are coupled together and you get modern cells. That's stage three, but still with the genetic information being shared, mostly by viruses traveling from cell to cell, so it is open source heredity. As Carl Woese described it, evolution could be very fast.

That's roughly the situation as Carl Woese described it — you have modern cells with metabolism directed by RNA or DNA, but without any private intellectual property, so that the chemical inventions made by one cell could be shared with others. Evolution could go in parallel in many different cells, so it could go a lot

faster. The best chemical devices could be shared between different cells and combined, so evolution would go rapidly in parallel. That was probably the fastest stage of chemical evolution, when most of the basic biochemical inventions were made.

Stage four is the stage of speciation and sex, which are the next two big inventions, and that's the beginning of the Darwinian era, when species appeared. Some cells decided it was advantageous to keep their intellectual property private, to have sex only with themselves or with the members of their own species, thereby defining species. That was then the state of life for the next two billion years, the Archeozoic and Proterozoic eras. It was a rather stagnant phase of life, continued for two billion years without evolving fast.

Then you had stage five, the invention of multicellular organisms, which also involved death, another important invention.

Then after that came us — stage six. That's the end of the Darwinian era, when cultural evolution replaces biological evolution as the main driving force.

"Cultural" means that the big changes in living conditions are driven by humans spreading their technology and their ways of making a living, by learning from one another rather than by breeding. So you are spreading ideas much more rapidly than you're spreading genes.

And stage seven is what comes next.

The question is whether any of that makes sense. I think it does, but like all models, it's going to be short-lived and soon replaced by something better.

The other thing I was going to talk about was domesticated biotech, which is a completely separate subject. That comes from looking around at what's happened to physics technology in the last twenty years, with things like cell phones and iPhones and the things that I see around me at the table.

Personal computers of all kinds. Digital cameras. And the GPS navigation system. All those wonders of technology, which have suddenly descended from the sky to the earth. They have become domesticated. That has been a tremendous change, something we never predicted.

I remember when von Neumann was developing the first programmable computer at Princeton. I happened to be there, and he talked a lot about the future of computing, and he thought of computers as getting bigger and bigger and more and more expensive, so they belonged to big corporations and governments and big research labs. He never in his wildest dreams imagined computers being owned by three-year-olds, and being part of the normal upbringing of children. It's said that somebody asked him at one point, how many computers would the United States need? How large would the market be? And he answered, eighteen.

So it went in totally the opposite direction.

VENTER: Well, it went in both directions.

DYSON: To some extent, but even the biggest computers are not much bigger than they were in those days. It's remarkable — I remember the very first computer in Princeton, and it was a huge thing — a room about as big as this tent, full of machinery. This was in 1951, '52. It was actually running smoothly around '53.

VENTER: But that was less powerful than your laptop.

DYSON: Oh, much less. The total memory was four kilobytes. And he did an amazing lot with that. Especially a biologist who was there at the time, called Nils Barricelli, did simulated evolution amazingly well with a memory of four kilobytes. He developed models of evolving creatures forming an ecology, and they showed punctuated equilibrium, exactly the way real species do. It was astonishing how much he could get out of that machine.

LLOYD: The problem is that computers get faster by a factor of two every year and a half, but computer programmers conspire to make them run slightly slower every year and a half, by junking them up with all sorts of garbage.

DYSON: Because von Neumann thought that he was dealing with unreliable hardware, he made another mistake. The problem was how to write reliable software so as to deal with unreliable hardware. Now we have the opposite problem. Hardware is amazingly reliable, but software is not. It's the software that sets the limit to what you can do.

My prediction or prognostication is that the same thing is going to happen to

biotech in the next 50 years, perhaps 20 years; that it's going to be domesticated. And I take the example of the flower show in Philadelphia and the reptile show in San Diego, at both of which I saw demonstrations of the enormous market there is for people who are skilled breeders of plants and animals. And they're itching to get their hands on this new technology. As soon as it's available I believe it's going to catch fire, the way computers did when they became available to people like you.

It's essentially writing and reading DNA. Breeding new kinds of plants and trees and bushes by writing the genomes at home on your personal machine. Just a little DNA reader and a little DNA writer on your desk, and you play the game with seeds and eggs instead of with pictures on the screen. That's all.

LLOYD: One of the reasons computers became ubiquitous is the phenomenon of Moore's Law, where they became faster and more powerful by a factor of two every two years. Is there an equivalent here?

DYSON: Exactly the same thing is happening to DNA at the moment. Moore's Law is being followed as we speak, both by reading and writing machines.

LLOYD: At roughly the same rate?

DYSON: Yes.

VENTER: It's happening faster. I had this discussion with Gordon Moore and I said that sequence reading and writing was changing faster than Moore's Law, and

he said, but it won't matter, as you're ultimately dependent on Moore's Law.

DYSON: I agree with that. At the moment it's going fast.

CHURCH: Unless we build bio-computers — right now the best computers are bio-computers.

BROCKMAN: It took two weeks for a 17-year-old to hack the iPhone — and here we're talking about DNA writers and readers. That same kid is going to start making people.

DYSON: That's true, the driving force is the parents, not the scientists. Fertility clinics are a tremendously large and profitable branch of medicine, and that's where the action is. There's no doubt this is going into fertility clinics as well. For good or evil, that's happening.

BROCKMAN: But isn't this a watershed event because of our ideas about life? What's possible will happen. What will the societal impact be?

DYSON: It's not true that what's possible will happen. We have strict laws about experimenting with human subjects.

BROCKMAN: You can't hack an iPhone either; certain activities along these lines are illegal.

DYSON: But it's different with medicine. You do get put in jail if you break the

rules.

BROCKMAN: Not in Romania.

DYSON: There are clear similarities but also great differences. Certainly it is true that people are going to be monkeying around with humans; I totally agree with that. But I think that society will put limits on it, and that the limits are likely to be broken from time to time, but they will be there.

SHAPIRO: I just want to bring in one distinction here, because two things are getting confused. To go to computers, I remember that perhaps 30 years ago there was something called Heathkit and the idea was, why buy a computer when you can build your own computer in your basement? Well I don't see anyone constructing their own computers in their basements any more.

If you purchase from computers from Dell or from IBM they will assemble them for you. But the actual construction, the difficult part, takes place in specialized institutions and then they make their products available. Everyone has a cell phone, but I doubt that most of the people, if they dropped it, could repair their cell phones. And that the new biotechnology — while humans would get the benefit, even now I think, one can contract to put the green fluorescent proteins into all sorts of animals and one artist has done just that and arranged to have specialized laboratories put it in, and then he had an exhibit where he made it seem as if he himself had done it. That isn't the case. DNA sequencing will be done massively, and engineering will be done massively, and new organisms will be constructed. But they will be done in specialized facilities. Only the products

will become available to the general public. No child will go into his basement and set up the necessary DNA synthesizers, or DNA sequencers and proceed to make his own new organisms.

DYSON: You're thinking like von Neumann, and I disagree.

VENTER: That used to be a compliment.

DYSON: It's true: what you will sell to the kids is kits — you won't sell the whole apparatus for doing things but you will sell a kit that will do the things that are fun, just as you do with computers that are sold for children to play games. The computers only play games, they don't actually calculate numbers.

LLOYD: In fact there's a good analogy in the history of computation — 30 years ago MIT freshmen arrived having built a computer, and then shortly after that they stopped building computers. Twenty years ago, or fifteen years ago, they arrived knowing how to program computers. But nowadays when freshmen arrive, far fewer of them have actually programmed a computer before, in the sense of writing a program in a language such as Java. But they use computers far more, and they're great users of software. They know vast amounts about how computers work and what you can do with the software. Why? Because it's a lot easier to do — why program a computer if somebody can enable you to just use the software and program it — of course when you're playing Grand Theft Auto, you're effectively programming the computer at the same time. So I suspect that what Freeman says is right, people will be using this new genetic technology, but maybe there's an analog of programming in the constructing new organisms

which will enable people to do it — an analog of software so people will become the users of the software.

SHAPIRO: I see children being able to purchase lizards, say, that glow in the dark — with green fluorescence, but I don't see them creating them in their basement.

DYSON: I think both are going to happen.

SASSELOV: Maybe the question is, what is the time scale for the second thing happening? That is, by then the technology will be so developed that we may be different as a species, and not care as much as we do today whether some kid is capable of tinkering with a human. Because we will have tinkered enough, in the regulated way, by then, so that it wouldn't matter as much.

DYSON: Yes, nobody can ever know in advance; all these things always turn out differently than you expected.

LLOYD: In fact this is a real specter — because as you say, we're not allowed to tinker with humans, but we are allowed to tinker with rats, that we very rapidly will develop rats who surpass us in all abilities. Whereas we're just stuck in the dark ages.

BROCKMAN: Freeman, last night I asked Richard Dawkins if he cared to comment on your chapter suggesting "the end of the Darwinian interlude". He sent the following comment with the caveat that it is a hastily written response solely for the purpose of this meeting. He writes:

"*By Darwinian evolution he [Woese] means evolution as Darwin understood it, based on the competition for survival of noninterbreeding species."*

"With rare exceptions, Darwinian evolution requires established species to become extinct so that new species can replace them."

These two quotations from Dyson constitute a classic schoolboy howler, a catastrophic misunderstanding of Darwinian evolution. Darwinian evolution, both as Darwin understood it, and as we understand it today in rather different language, is not based on the competition for survival of species. It is based on competition for survival within species. Darwin would have said competition between individuals within every species. I would say competition between genes within gene pools. The difference between those two ways of putting it is small compared with Dyson's howler (shared by most laymen: it is the howler that I wrote *The Selfish Gene* partly to dispel, and I thought I had pretty much succeeded, but Dyson obviously hasn't read it!) that natural selection is about the differential survival or extinction of species. Of course the extinction of species is extremely important in the history of life, and there may very well be non-random aspects of it (some species are more likely to go extinct than others) but, although this may in some superficial sense resemble Darwinian selection, it is not the selection process that has driven evolution. Moreover, arms races between species constitute an important part of the competitive climate that drives Darwinian evolution. But in, for example, the arms race between predators and prey, or parasites and hosts, the competition that drives evolution is all going on within species. Individual foxes don't compete with rabbits, they compete with other individual foxes within their own species to be the ones that catch the rabbits (I would prefer to rephrase it as competition

between genes within the fox gene pool).

The rest of Dyson's piece is interesting, as you'd expect, and there really is an interesting sense in which there is an interlude between two periods of horizontal transfer (and we mustn't forget that bacteria still practice horizontal transfer and have done throughout the time when eucaryotes have been in the 'Interlude'). But the interlude in the middle is not the Darwinian Interlude, it is the Meiosis / Sex / Gene-Pool / Species Interlude. Darwinian selection between genes still goes on during eras of horizontal transfer, just as it does during the Interlude. What happened during the 3-billion-year Interlude is that genes were confined to gene pools and limited to competing with other genes within the same species. Previously (and still in bacteria) they were free to compete with other genes more widely (there was no such thing as a species outside the 'Interlude'). If a new period of horizontal transfer is indeed now dawning through technology, genes may become free to compete with other genes more widely yet again.

As I said, there are fascinating ideas in Freeman Dyson's piece. But it is a huge pity it is marred by such an elementary mistake at the heart of it.

Richard

DYSON: Good. Yes, I have two responses.

First, what I wrote is not a howler and Dawkins is wrong. And I have read his book.

Species once established evolve very little, and the big steps in evolution mostly

occur at speciation events when new species appear with new adaptations. The reason for this is that the rate of evolution of a population is roughly proportional to the inverse square root of the population size. So big steps are most likely when populations are small, giving rise to the "punctuated equilibrium" that is seen in the fossil record. The competition is between the new species with a small population adapting fast to new conditions and the old species with a big population adapting slowly.

Second, it is absurd to think that group selection is less important than individual selection. Consider for example Dodo A and Dodo B, competing for mates and progeny in the dodo population on Mauritius. Dodo A competes much better and has greater fitness, as measured by individual selection. Dodo A mates more often and has many more grandchildren than Dodo B. A hundred years later, the species is extinct and the fitness of A and B are both reduced to zero. Selection operating at the species level trumps selection at the individual level. Selection at the species level wiped out both A and B because the species neglected to maintain the ability to fly, which was essential to survival when human predators appeared on the island. This situation is not peculiar to dodos. It arises throughout the course of evolution, whenever environmental changes cause species to become extinct.

In my opinion, both these responses are valid, but the second one goes more directly to the issue that divides Dawkins and myself.

VENTER: I have trouble with some of the fundamental terms. What's your definition of "species"? That's something I have great difficulty with lately out of

our research.

DYSON: Yes, it is a problem — it's supposed to be just a population that breeds within the population but not outside, but of course there are all sorts of exceptions.

VENTER: That ignores most of biology.

DYSON: Yes, so I don't know what the real definition is. But that's the conventional definition.

VENTER: It's a human definition.

DYSON: It is fuzzy. Like most things.

LLOYD: So for sexually reproducing species, then, it's less fuzzy than for bacteria.

DYSON: Right.

VENTER: But it really comes down to one or two recognition molecules that determine the species — if it's based on interbreeding, it's the sperm recognition sites, right?

DYSON: Yes.

VENTER: So that determines the species, then.

DYSON: Well, amongst other things.

CHURCH: Chromosome dynamics, morphology, behavior — many things. Depending on how complex the organism is.

VENTER: It's easy to tell a human from a giraffe, and we can call that a different species.

DYSON: One of the books that I've learned most from, is *The Beak of the Finch*, which describes evolution as it's observed in the Galapagos by Peter and Rosemary Grant. It's remarkable that they can actually see from year to year species starting to hybridize when conditions are good and then separating again when conditions are bad. So even on a year-to-year time scale you can actually see this happening, that species are not well-defined.

LLOYD: Sorry, I'm not familiar with this work. So they hybridize when times are good, and when times are bad they separate into smaller populations. Is this so that they can evolve more rapidly?

DYSON: Yes. So they can specialize. Because in bad times you have to specialize on chewing particular seeds.

VENTER: During droughts, all that was left were these really hard seeds. Finches that survive have Arnold Schwarzenegger beaks.

DYSON: Not only those — you can also have a separate population which specializes on the small seeds, which have small beaks. It happens because of the geography that you have violent swings in climate. During El Niño conditions are wet, and between El Niños, conditions are dry. So selection is brutal — almost every year about half of them get selected out.

VENTER: One of the highlights of my round-the-world expedition was meeting up with the Grants in the Galapagos, and their little tent on the site of Daphne Major. They spent three months on this island in this little tent, there's no fresh water, there's nothing there. And they live off of bottled water and cans of tuna fish. And I took them a bottle of chilled champagne. It became a happier eco-system. Remarkable what they've done.

DYSON: The enormous advantage that they had was that the birds are completely tame. You can just walk up to a bird and put a ring around its leg and it doesn't fly away. That's what made it all possible. They know every bird personally.

VENTER: Better than tame — if you walk on their path, the boobies and stuff will peck at your leg. It's their island. The humans become non-tame after a while. But so that's an important part of the definition. Are the finches with the larger beaks a different species, in your view?

DYSON: Yes, according to Darwin they are. In fact they do interbreed quite extensively.

VENTER: So two base pair change in a genome could be sufficient to create a new species out of 1.5 billion.

DYSON: Yes.

VENTER: I'm not sure everybody will buy that definition... So that makes you a very different species than George.

DYSON: The real problem is the lawyers. You have the endangered species act; that means you have to make a legal definition of the species.

CHURCH: That's true. We're all endangered.

LLOYD: I gather human beings are a genetically very non-diverse species. We take two squirrels on this tree right here — they're much farther apart genetically than we are with any other human being on the face of the earth. So we're inclined to see things in our own light.

VENTER: What's your evidence for that?

CHURCH: It's true for chimpanzees; I don't know about squirrels.

LLOYD: But homo sapiens is a quite recent species — and also the mitochondrial DNA evidence suggests that we're descended from common ancestors in the not very distant past — within the last hundred thousand years or so. So there seems to have been a genetic bottleneck in the human species, compared with hominids

as a whole, within the last hundred thousand years. Which makes us much less diverse than, for instance, squirrels.

SHAPIRO: The thrust of what Freeman was saying if we accept most of what he said, which I certainly do, is that concepts like species and interbreeding are about to become in a sense extinct. Because entering the new era, laboratories will exist which will recreate species or combine qualities of one species with qualities of another and it will be up to the designer the extent to which they interbreed or interbreed with existing organisms and so on. So that perceivably, if civilization continues we will then be in charge of what species may come into being and what species do not.

LLOYD: I have a query: is that actually important, actually? Freeman, you said we reached the end of Darwinian evolution, where human beings are the dominant species on earth, and species that can't co-evolve with humans are probably doomed. But this means that in this end of Darwinian evolution, then genes are no longer so important, and instead ideas, which can be generated more rapidly, and — dare I even say — things like computations and software are more important. Are you envisaging an era where genetic information returns to the predominant position that it had for billions of years on earth?

DYSON: No, I don't look very far. I'm quite conservative as far as human society is concerned. We would be wise to keep ourselves as much as possible the way we are, and I hope we'll be successful in it. I don't see any great likelihood if you monkey around with humans that you'll produce anything much better.

BROCKMAN: This sounds like an engineer's approach, rather than a thinker's approach. As a scientist, aren't you talking about a huge watershed concerning our ideas of what it means to be human or even what it means to be alive? Can you imagine what ideological factions or religious groups would do with some of the statements that have been made this afternoon?

LLOYD: Ironically many religions are sets of ideas, and one of the things that many religions tend to do is to try to sequester themselves genetically. Keep the gene pool within this religion from people within this religion — prevent intermarriage with people of other faiths. You could say religion is almost an attempt by ideas to get back to the good old days of rapid evolution via genetic engineering in small populations.

DYSON: I'm not familiar with this feeling that culture is collapsing. All these millions of people who are now publishing blogs on the Web are to my mind producing something you might call culture. Of very uneven quality, but it's easier to publish now than it used to be. And that to me is not necessarily a disaster. It may be a step forward.

LLOYD: In fact it's easier to preserve information as well. In the past one of the main problems with culture is it would disappear because there was only one copy. When there's only one copy, things get easily destroyed. And yes, maybe because in the United States we don't have as much culture so we're not so worried about losing it.

Perhaps worrying about the wholesale copying and monkeying with genetic

information might open people's eyes to the danger of copying and monkeying with ordinary cultural information — for instance, violating copyrights. While I am usually for any kind of information manipulation I can think of, it does seem a little strange to try to manipulate human genomes. Of course, the primary way of manipulating genomes in the past, which people have been doing for ages, is by breeding. People are rather squeamish about attempts to manipulate human genomes to create perfect human beings just by breeding. This is an old fear among people and an old temptation as well. We may not be so culturally bereft with the mechanisms that we need to cope with these kinds of issues as we might think. It is scary. But anything fun is scary.

PRESS: Is open source sort of an inexorable direction that we're moving in — as people blog openly, and copyrighted music seems to be losing out to open and tradable music — is that the way you expect it's going to be with genomics as well, that ultimately this information is all going to be openly and freely available, and that's the way this whole system is going to progress?

DYSON: Not necessarily. Bill Gates is still around. But that remains to be seen. Clearly this is the alternative.

CHURCH: Genomics for the most part has been quite open historically — even in profit-making sectors they will publish papers and so forth, and the genome project went so far as to try to publish things within one week of collecting the data. So it's really quite aggressive so far. Almost every genome that you could possibly want, including some that some people would prefer not to be in open source, like small pox, which Craig helped to do, and the 1918 flu virus — all

those things are available. So I think that is a trend.

DYSON: It's unfortunate that small pox is out there — the world would be a lot safer if that hadn't been published.

VENTER: I can disagree very violently with you on that.

DYSON: Good. That's a minor exception, but as a general rule, openness is by far preferable.

VENTER: Even with that, I think I could convince you openness is far more important. There were two states that were funding an incredible amount of secret research — the U.S. and the former Soviet Union — on trying to modify small poxes, make them more dangerous, et cetera. So if it was not open source, those states would be the only ones with access to this information. There would be nothing out there for either tracking it, understanding it, making better vaccines, et cetera, if it was even a real threat. And on the synthetic biology side, it's a very, very low threat because the DNA is not infective. It's a hypothetical threat that people like to use to scare people, but in reality it's really not one.

CHURCH: DNA is not infective but you can make infective viruses with the DNA in the lab?

VENTER: Hypothetically. But nobody's done it yet.

CHURCH: With other pox viruses you can do it — so it's not that hypothetical.

VENTER: There are probably a few thousand pox viruses out and very closely related species that could easily become small pox. I'll argue for open source of information — my genome is on the Internet, but I'm much more selective who I share my biological materials with. There's open source and there's open source.

SHAPIRO: You did raise an interesting point there, though, because genetic privacy is something which is often debated — the rights of individuals to genetic privacy, not to have their genomes known.

VENTER: But that's driven by fear, not by knowledge.

SHAPIRO: But what I'm saying is, that genetic privacy actually maybe impossible. Let us say that I wish that he hadn't put his genome on the Internet and wanted it secretive, say he was running for public office and had some gene for some mental instability, and therefore wanted no one to have his genome; yet someone wanted his genome. All I'd need to do is swipe your glass, and shake your hand.

VENTER: This is issue that we could talk about that George and I have been facing that's counteractive to what our government is doing. Francis Collins is setting up data bases, where you have to have retinal scans and finger prints to have access, and we're publishing our data on the Internet. So, open source is not a guarantee of any means at all. We hope by making human genetic data available, people will find in fact that it's almost impossible for your scenario, wherein you can look at one gene and say this person's going to have mental

illness. Even the entire genetic code doesn't provide that answer. You have to know the environment; you have to know a lot of other things.

Perhaps 50 years from now we can get much closer to those answers of predicting things, but we are not just genetic animals. My dangerous idea is that we're probably far more genetic animals than society is willing to accept. But we're not purely genetic animals, so I don't think it's going to be as predictive as some people think.

SHAPIRO: Well, certain specific things will be predictive — for example, Huntington's disease is due to a repeat of certain letters in DNA.

VENTER: There are some very rare exceptions, yes.

SHAPIRO: You can even tell what onset is likely at what age by counting the number of repeats that are present.

VENTER: But that's the exception that doesn't make the rule. That's what every geneticist has used as the few early examples of success in genetics of single gene disorders.

SHAPIRO: But there are cases where individuals themselves didn't want to know whether or not they had inherited the gene for Huntington's disease, or if they did, whether they were going to have a severe form. Yet if some external person wanted to inform himself as to whether that individual did carry the gene, it would almost be impossible to prevent that individual from getting the

information. You would practically have to live in seclusion, with all of your clothing, all of your artifacts destroyed on contact.

LLOYD: It's interesting because in fact the digital nature of genetic information, the fact that it's seven billion bits that can easily be written into a computer hard drive, makes genetic information much more like the information in computers and it can be manipulated in that way. Whereas strangely enough, our mental information, the information that's in our brains, is much less digital in a fashion, and much harder to get hold of.

And in fact it does suggest that, since this information has been digitized, and will continue to be digitized and manipulated, and be more available, the question of how secrecy and privacy for genes is rather similar to the privacy of your iPhone? How privately are you allowed to keep the information in your iPhone? How privately are you allowed to keep the information in your genes? Because it will be available, and it will be possible to get it and to digitize it so then the question is, do you need codes for protecting your genetic code? Maybe everybody should be issued their own public key cryptic system so they and only they can have access to their own genetic code.

CHURCH: We're kind of in a state of change where we're deciding what's the right thing. For example, consider our faces. Some people keep their faces completely masked; in most situations it's considered anti-social to keep your face completely masked. Like walking into a bank, for example. But it's extraordinarily revealing — it not only reveals something about your physiology, your current health, your relationship with the person you're talking to, whether

you're angry or very happy — it's very revealing. And so we've made a conscious decision in society, for the most part, to not keep that private. We might do the same thing for genomes, it could be, who are we protecting? But it's an open question.

SHAPIRO: Well we shed cells so easily unlike faces that it's almost impossible to keep your genome private, if there is someone out there determined to have it.

CHURCH: I agree with you. We'll all become bubble people, living in our little hermetically sealed bubbles so nobody can get in.

LLOYD: Who steals my genome steals trash, right?

J. CRAIG VENTER

I have come to think of life in much more a gene-centric view than even a genome-centric view, although it kind of oscillates. And when we talk about the transplant work, genome-centric becomes more important than gene-centric. From the first third of the Sorcerer II expedition we discovered roughly 6 million new genes that has doubled the number in the public databases when we put them in a few months ago, and in 2008 we are likely to double that entire number again. We're just at the tip of the iceberg of what the divergence is on this planet. We are in a linear phase of gene discovery maybe in a linear phase of unique biological entities if you call those species, discovery, and I think eventually we can have databases that represent the gene repertoire of our planet.

One question is, can we extrapolate back from this data set to describe the most recent common ancestor. I don't necessarily buy that there is a single ancestor. It's counterintuitive to me. I think we

may have thousands of recent common ancestors and they are not necessarily so common.

J. Craig Venter is one of leading scientists of the 21st century for his visionary contributions in genomic research. He is founder and president of the J. Craig Venter Institute. The Venter Institute conducts basic research that advances the science of genomics; specializes inhuman genome based medicine, infectious disease, environmental genomics and synthetic genomics and synthetic life, and explores the ethical and policy implications of genomic discoveries and advances. The Venter Institute employes more than 400 scientist and staff in Rockville, Md and in La Jolla, Ca. He is the author of *A Life Decoded: My Genome: My Life.*

J. CRAIG VENTER: Seth's statement about digitization is basically what I've spent the last fifteen years of my career doing, digitizing biology. That's what DNA sequencing has been about. I view biology as an analog world that DNA sequencing has taking into the digital world . I'll talk about some of the observations that we have made for a few minutes, and then I will talk about once we can read the genetic code, we've now started the phase where we can write it. And how that is going to be the end of Darwinism.

On the reading side, some of you have heard of our Sorcerer II expedition for the last few years where we've been just shotgun sequencing the ocean. We've just applied the same tools we developed for sequencing the human genome to the

environment, and we could apply it to any environment; we could dig up some soil here, or take water from the pond, and discover biology at a scale that people really have not even imagined.

The world of microbiology as we've come to know it is based on over a hundred year old technology of seeing what will grow in culture. Only about a tenth of a percent of microbiological organisms, will grow in the lab using traditional techniques. We decided to go straight to the DNA world to shotgun sequence what's there; using very simple techniques of filtering seawater into different size fractions, and sequencing everything at once that's in the fractions.

For example, we discovered that almost every microorganism in the upper parts of the ocean has a photoreceptor similar to the ones in our own eyes. We knew there were one or two bacterial rhodopsins, but people thought they were rare molecules; it turns out it's probably one of the largest gene families on the planet. It's the same gene family that we have in our own eyes — our own rhodopsins, our visual pigments. Only instead of just capturing light information, these organisms capture a light, and convert it into cellular energy — non-photosynthetic, totally separate mechanism. When we set out to go to the Sargasso Sea surrounding Bermuda, all the marine microbiologists told us nothing was there, it was a desert, and we'd find only a few organisms. Instead we found tens of thousands of organisms in just a barrel of seawater. And the reason they said we wouldn't find anything is that there are no nutrients there; they said there are no nutrients, therefore there'll be no life. It turns out they don't need the nutrients because their energy is derived directly from sunlight.

LLOYD: Do you think perhaps that first use of rhodopsin was to harvest energy?

VENTER: Quite possibly. And then it was adapted for visual pigments because it was a light-recognition molecule. And the other aspect of it is there's a wide range of new ones — and we have thousands of these now — and so lining up the proteins of all these, there's a single amino acid residue that determines the wavelength of light that the receptors see.

There's a single base change in the genetic code that determines the amino acid responsible for the wave length of light seen by the receptor. Changing one base in the genetic code can switch the light seen from for example blue to green. We found when we went back looking at the distribution of where these different rhodopsin molecules are, they totally segregate based on the color of water. The photoreceptors in the organisms in the deep indigo blue of the Sargasso Sea see blue light. You get into coastal water where there's a lot of chlorophyll, they see primarily green light. To me this is classical Darwinian selection. A single base pair determines the switch between blue and green. And whatever wavelength of light, it clearly provides a survival advantage for that environment.

It turns out, just looking at the populations, that this switch between blue and green has happened at least four times, back and forth. And so in one hand it seems like a classical reformation of Darwinian thinking. On the other hand, under each type of 16 S RNA, we have in fact hundreds to thousands of different cells, different genomes. Are they different species?

These are the types of species and definitions of life that I've been lately been devoting much of my professional career to. And it turns out these diverse organism share most of the gene content with each other and most of the gene order is conserved. But the sequence variations are as high as 60% between cell types. And according to classical Darwinism, this should not be the case, and it's not what anybody expected — there should be that one or a few species that survived and all the others died out.

It turns out these are ancient lineages, all existing in parallel, all virtually identical to each other but not quite identical.

Maybe they just have that equivalent base change to where one can see the equivalent of blue light and the other can see green light and we can't see that because we don't know what all the other molecular switch changes mean. But we find this over and over again with every microbial species. I'm forced to use the term species, because there's no better common word, but species' is a very vague term. It's sort of a group of closely related organisms, much like the human population is a group of closely related organisms. If it's a new species because it has that one base change and it can see blue light instead of green light, it's a phenotypic difference, it's a survival difference; it's got nothing to do with sexual reproduction, but it's roughly the same order of change of genetic information.

I have come to think of life in much more a gene-centric view than even a genome-centric view, although it kind of oscillates. And when we talk about the transplant work, genome-centric becomes more important than gene-centric. From the first third of the Sorcerer II expedition we discovered roughly 6 million new

genes that has doubled the number in the public databases when we put them in a few months ago, and in 2008 we are likely to double that entire number again. We're just at the tip of the iceberg of what the divergence is on this planet. We are in a linear phase of gene discovery maybe in a linear phase of unique biological entities if you call those species, discovery, and I think eventually we can have databases that represent the gene repertoire of our planet.

One question is, can we extrapolate back from this data set to describe the most recent common ancestor. I don't necessarily buy that there is a single ancestor. It's counterintuitive to me. I think we may have thousands of recent common ancestors and they are not necessarily so common.

Other things you can throw into the mix: we have organisms that could absolutely survive long-term space flight, they can take millions of Rads of ionizing radiation, they can be totally desiccated; when they reach an aqueous source they can repair their genome and start replicating again. Thus, you could potentially view evolution as a six- to seven-to eight billion year event, not a three- to four-billion year event, if life can travel around the universe. That adds a lot of dimensionality to things when we think of life in other planets and galexcies We exchange roughly a hundred kilograms of material annually with Mars. So we're exchanging biological material and biological information. To me it's just a matter of time until life is found on Mars. It's inevitable. It won't tell us whether it originated on Mars, or originated on Earth, but there'll be common overlap. We won't know if we don't know our own planet's genetic repertoire, which we're in the earliest stages of discovering. There are the evolutionary aspects, the origin of life aspects to this, which make it very intriguing.

But thinking about the next stage makes it more fun if you start to think of our present set of 10 million genes in our data bases soon to be 20 million unique genes.

I spent early parts of my career trying to make site-directed mutants of neurotransmitter receptors, where you change a single base pair and study the change in function of the new altered protein. Now you can just go to a database of our gene sets and find 35,000 variants of a single gene that you know survived and work in nature. It is a very different view of the world, and it will change the experimental approaches. But this new data set gives us a wonderful new repertoire, if you view these genes as the design components of the future.

For me, it all started in 1995 when we sequenced the first two genomes of living entities. The first was *Haemophilus influenzae*, the second was Mycoplasma genitalium. Mycoplasma genitalium had only 550 some odd genes; *Haemophilus influenzae* had 1,800. And so this was the first comparative genomics that could be done for living organisms. We just started asking very simple questions for example — if one species needed 1,800 genes and the other needed 550, are there species that can get by with less? Can you define a minimal genetic operating system for life? Could we define life at a genetic level? Obviously extremely naive questions but the view of biochemistry and genomics by the scientific community was very limited as well. For example when we published the *Haemophilus influenzae genome* a well known biochemist at Stanford University said we obviously assembled it wrong because it didn't have a complete TCA cycle. And everybody knew that every organism had a complete glycolytic

pathway and a complete TCA cycle. And Haemophilus only has half of one.

SHAPIRO: Therefore it's not an organism.

VENTER: No, therefore they assumed we made a mistake in the sequencing and the assembly. Now we see every repertoire under the sun, for example the third organism that we sequenced was the first Achaea that we did with Carl Woese, it was methanococcus jannaschii, which has neither a TCA cycle nor glycolysis. It makes all its cellular energy by methanogenesis, going from CO2 to methane, using hydrogen as its energy source. CO2 is its carbon source for all the carbon in the cell.

What Freeman was talking about when he said you can separate metabolism from replication is certainly true, and we have at least 20 different modules we could plug in for metabolism. There's not a universal genetic code that will be an operating system — it can be a choice of methane production, or of glucose metabolism, or anaerobic metabolism, oxidative metabolism or some other approach. So the naïve assumption that we could even define an operating system clearly went out the window.

We first we tried to knock out genes in *Mycoplasma genitalium* to see if we could get to a minimal genome. We could only knock out one gene at a time, and we could knock out about 100, one at a time, but that didn't tell us whether if you knocked out all 100 that you can get a living cell from it. But we also learned other things for example — that essentiality for life is very much a relative term. We can have a gene that is absolutely essential for life in one circumstance that is

not another — the simplest example I give is that there are two genes in *Mycoplasma genitalium* for sugar transport, one for fructose and one for glucose; if you have both sugars in the media and you knock out either of the transport genes, you would say they are non-essential genes as the cell lives. If you only have glucose as your sugar in the media, and you knock out the glucose transporter, the cell dies and you say, ah, that's an essential gene. All these things are definitional, based on two things: the genetic code and the environment.

Right now we're all focused on the genetic code because it's something we can define and the environment is so many orders of magnitude more complex to define, but we're having this trouble with a single cell with a few hundred genes; we as humans have a hundred trillion cells with 23 thousand or so genes, and an infinite number of combinations, so defining our environment is going to be a lot more complicated than that for a single cell. We decided the only way to answer these questions was to make a synthetic chromosome to understand minimal cellular life.

We actually set out with some simplistic experiments in 1995 to make the phi-X174 viral genome. My colleagues in this — and I have an incredible set of colleagues — including Ham Smith who discovered restriction endonucleases that led to his Nobel Prize in 1978, and Clyde Hutchison, who actually was in Sanger's lab when he sequenced phi-X174 and was one of the inventors of site-directed mutagenesis. We thought we'd just synthesize a set of overlapping small olgionucleotides , anneal them together with a replicating enzyme and we'd get the complete genome. We decided to start with phi-X174 because of its historic value and because there are very few base pairs in it that can be varied and still

get a functional virus. We figured it would be a real test, because you had to have it accurately synthesized to get a functional virus from it. And even though we had selection by infectivity, which gives us a million to ten million-fold selection, even though we got full-length genome molecules, none of them were viable.

And I use the term 'viable' loosely with viruses — it's not clear what a viable virus is, but it's one that can infect a cell and start the replication machinery going. At the same time — John, you raise the question of whether this all happens in a moral vacuum — we decided that, if the only approach forward is to synthesize a genome, and then to start modifying it, that we were going to be creating new biological entities, or species, that hadn't existed before.

It's true if you knock out a gene, if one base-paired difference is sufficient for something new, this is nothing new in molecular biology, except if we try to do this deliberately and then try and design things from it. So we asked Art Kaplan at the University of Pennsylvania to spend time with his group looking at this, and his team brought in every major religion to see if it is okay to synthesize life in the laboratory. After a year and a half study basically none of the religious groups objected to it, because they couldn't find anything in their scriptures that said it shouldn't be done.

BROCKMAN: Four thousand years ago.

VENTER: In fact, they mostly came up with the opposite of this notion of "playing God" that the lay press talks about all the time with regard to this and other related work. Every major religion basically said that part of their dictate to

humanity is that you're supposed to try and use knowledge to benefit humanity. The report they published in *Science* in 1999 noted that what we were doing and the way we were approaching it was very reasonable and we should proceed. The only caution was about biological terrorists using our techniques to try and make biological weapons.

We began with that as a basis. Then the entire project was postponed because I had the opportunity to sequence the human genome. We gave up on synthesizing new life forms for a while.

LLOYD: That's interesting, so you actually came around to sequencing the human genome from the perspective of trying to construct an artificial organism.

VENTER: They were parallel tracks. One didn't lead to another. I wouldn't want you to make that intellectual leap.

LLOYD: But you were trying to synthesize life then stopped doing it in order to sequence.

VENTER: Right. And so we've spent basically since 2002 starting in full form to find ways to synthesize genomes. We started back with phi-X174 and Ham and Clyde came up with some new ways of error-correction and improving synthesis. DNA synthesis is a degenerative process where the longer you make the molecules, the more errors are in them — a, N minus 1 system. Trying to just synthesize off of chemical synthesizers and get really accurate molecules is currently not possible.

But it was an exciting leap for us when we actually made this chromosome, injected it into E. coli, and the next thing we knew E. coli started using that synthetic genome to start producing the phi x 174 viral particles that then started killing the E.coli. It was clearly that this human-made piece of chemical DNA software was now building its own hardware. One of the exciting parts of synthetic biology and synthetic genomics is that that is possible.

There are two real components or questions to this, and there's even more in terms of when you think of the implications. One is, can you make these large molecules, and the answer is now absolutely 'yes'; we can make whole chromosome-size macromolecules of DNA. At the same time as we were thinking about the synthesis, we were thinking, how do we boot-up the chromosome in a cell? In the process we initially thought, well, you'd like a ghost cell that just has the ribosomes and other cytoplasmic components in it but is devoid of its chromosome, and we tried numerous ways to get rid of a chromosome in a bacterial cell. Then Ham Smith came up with the notion of, well maybe we don't really need to get rid of it, we'll just put the new chromosome into a cell and when they segregate maybe the chemically-made chromosome will go in one daughter cell cell and the other one will go another way. — actually it is much more complex than that.

A short while ago we published a paper in *Science* on genome transportation, where we took a purified chromosome from one species, made sure it was totally devoid of any protein, and put that chromosome into another bacterial species, and it's the ultimate identity theft because the new chromosome we put in

completely took over the cell, and the cell converted completely into the cell dictated by the new chromosome. We put an antibiotic selectable marker gene into the transplanted chromosome so we could select for those cells with the new chromosome — the story on how the selection happens and why one chromosome survives is actually much more interesting and deals with an important part of evolution, but needless to say the new chromosome dictated everything. All the proteins changed over to that. The phenotype of the cell, everything changed, converted from the old species into the new species.

DYSON: How many generations did that take?

VENTER: It's not clear. It could have happened in the first couple of generations — until you get enough cells that you can see and do biochemistry on, they've gone through, you know, dozens of generations. It would be nice to do stop-flow experiments and see what happened in those initial phases. But we tried to take an EM-micrographs to see if there were hybrid cells that had both sets of chromosomes in them, but we did not find any evidence for hybrid cells. You can see why restriction enzymes were so important for cellular evolution, because speciation could have problems with DNA uptake, if whoever had the dominant genome could just immediately take over your species and transform you into them. It's true identity theft at the molecular level.

LLOYD: And you wonder why people are worried when you describe it like that — I feel that it's happening to me right now!

SHAPIRO: How genetically different were the two cells?

VENTER: They were roughly of the equivalent of man versus mouse. They were closely related mycoplasmas, and it turns out the restriction barriers are really major barriers. It became clear to me for the first time — how important restriction enzymes were for evolution, because if foreign DNA can just go in and take over the cell, obviously that's what you want protection against. It's their equivalent of an immune system. To actually do genome transplants on a reasonable scale, we have to overcome the restriction bariers. In the case that worked, the chromosome that we transplanted had a restriction enzyme that we think chewed up the chromosome that was actually in the cell. The one in the cell did not have a restriction enzyme, so there was no restriction barrier.

SASSELOV: Is your gut feeling that there is no need for a hybrid generation to develop, that you jump entirely from one to the other.

VENTER: I don't know.

SASSELOV: But do you have a gut feeling for it?

VENTER: Gene expression and new protein synthesis can happen pretty quickly. If we're transplanting in a new chromosome and it immediately starts to get transcribed events could happen quite rapidly. Turning over the membrane and all its content could perhaps take a couple of cell divisions and generations to get them completely changed over. After several generations we ran 2D Gels and some protein sequencing — every protein there was the ones dictated by the new chromosome.

The genome transplant experiment was the key one for synthetic life. But I think definitions are really important here; none of us are talking about creating life from scratch, because that's not what's happening. We're taking two approaches; we're taking the genome transplant approach and we have a team working on trying to isolate every protein from the cell to see if we can put the proteins together, with the chromosome, with some lipids, and get spontaneous cell and life formation.

LLOYD: You're saying it's much less like open source software and more like Microsoft, whose software actively resists operating with exterior software.

VENTER: Absolutely. You could not have speciation without it.

These things get down to basic definitions of life. The lay press likes to talk about creating life from scratch. But while we can create and develop new species, we're not creating life from scratch. We talked about the ribosome; we tried to make synthetic ribosomes, starting with the genetic code and building them — the ribosome is such an incredibly beautiful complex entity, you can make synthetic ribosomes, but they don't function totally yet. Nobody knows how to get ones that can actually do protein synthesis. But starting with an intact ribosome is cheating anyway right?
That is not building life from scratch but relying on billions of years of evolution.

When starting with an existing protein synthesis machinery we can create new life forms, we can create a synthetic chromosome that we can now do transplants of,

and develop new species with very unique properties — so we can create human-made species — but we're not really creating life from scratch. You can boot up a system but right now all life derives from other living entities. What we're doing is really no different, because we're just putting a new operating system into a living cell.

If George or anybody else doing this can take even the raw protein components and lipids and boot up a cell and get it activated and go from the chemical molecules to a living cell, that's a big conceptual barrier that remains to be passed. But even that's cheating. If you look at the Urey-Miller experiments of the chemicals that can be made in certain environments and start with basic amino acids and nucleic acids and all of a sudden you get life out of that — that would be creating life from scratch. Starting with a ribosome is sort of like building a Tesla electric car now, but that wasn't the first car ever made. They borrowed a lot of components from other cars.

SHAPIRO: At the basic level I'm sure that even the DNA synthesizers you use and the nucleotides are derived by degrading bulk DNA taken from other organisms. They're not synthesized in the laboratory.

VENTER: I think they're all chemically made, aren't they?

SHAPIRO: No, because they are all one mirror form, not the other one. It would be a chemical nuisance to have to resolve all of the nucleotides and then put them into the machines and synthesize them.

CHURCH: They're certainly more synthetic than they were in the past. They were harvested from salmon DNA and yeast and now they're actually made from very simple precursors. But you're right, the chiral forms, the riboses?

SHAPIRO: Ultimately come from life.

CHURCH: Although you can, and I'll talk about it in my time, make chiral forms.

SHAPIRO: You could in principle, but in practice what's happening is people are getting the basic parts out of life and putting them back into life.

CHURCH: You have to ask why? Why does it matter?

SHAPIRO: Saves time.

VENTER: But these are important notions that get lost in the over-simplified interpretation of what we are doing. I don't think we are creating life. We're coming up with new modified life forms, and we should be able to go from the digital world right to the analog world in the computer, and we have a team working on a program to do that, designing a species in the computer.

It's only a short time away from doing that just to have systems crank out synthetic chromosomes. In fact I've talked to various funders about trying to design a robot that could build a million chromosomes a day. Because then we can have a new field that I call combinatorial genomics. We don't know answers to seemingly simple questions such as; is gene-order important? If you just

scramble the order of the same genes does life work the same way? We know with operating systems and operons gene order is important but in the genome as a whole it is perhaps not; — that was one of the biggest surprises with the human genome, multi-subunit proteins such as the nicotinic acetylcholine receptor that has five sub-units which were all in different chromosomes. The assumption was that they would all be together in order on one chromosome. Maybe as long as all the parts are there, that's all that matters; we just need the gene set and the gene repertoire. There are so many different directions we can go with this conversation, maybe I should just stop here.

BROCKMAN: Where is it all going?

VENTER: I disagree with this recreational use — who cares whether you have green fluorescent protein in fish or something. Hopefully these are fads that will go away very quickly — maybe they're important for inducing biological concepts — but there are far more pressing issues.

To me the biggest issue, that's why we decide to put most effort there, is what we're doing to our own environment and atmosphere by taking billions of tons of oil and coal, burning it, and having the CO2 go into the atmosphere. Right now it's over four billion tons a year — and the estimation in 50 years is that it's going to be on the order of 12 billion tons a year of CO2. It's a big experiment that we are doing with our planet that hasn't happened during the existence of human life. It's happened to the
planet before. It's a dangerous experiment to do. We can only estimate outcomes from it. But we have to have some potential solutions.

While everybody's looking just to physics for the solutions, I have been arguing that biology could play a major role, if not the ultimate role, in the solution. And that's why we started Synthetic Genomics, the company, to try and design genomes to make new fuels.

DYSON: We were doing that at Oak Ridge 30 years ago.

VENTER: People have been looking at biology, but they were looking for naturally occurring organisms to do this. Right? That worked well during World War II .

DYSON: It was clear that biology would be the right way to do it. Even 30 years ago.

VENTER: There has been an effort within some DOE labs to find and use natural organisms to produce hydrogen or other potential fuels but the efforts have essentially gone nowhere due to the scale of the problem and the efforts have been very limited. It has been a battle in the Department of Energy as to the importance of biology. Many in congress have the niave idea that it should only be done at the NIH.

LLOYD: If we destroy our idea-based culture by raising the ocean levels by a hundred meters, we can just return to the good old days of Darwinian natural selection. Go back to Stage Five, right?

VENTER: The choices are doing something to change the environment or being able to engineer the human genetic code to be able to survive in different environments than we can currently.When we can actually design and build millions of new organisms a day, single cell organisms — the first single-cell organisms are months away, the first synthetic eukaryotic cells are less than five years away, and multi-cellular systems are not orders of magnitude more complex to do. In fact in some respects they're easier to do because the ocean is a multi-cellular genome system, just different cells provide different components — the cells can just be associated loosely in the environment, but there's only a small number of cells that actually fix nitrogen that provide that for the whole pool. It's a cooperative environment, and having the cells get together to do it I think is not much more complicated than engineering some good stem cells, if we want to do that.

SASSELOV: Craig, I wanted to connect what you are saying with what Freeman was saying about the Darwinian era. To me his idea is not so much about the end of a process but the beginning of a new phenomenology, from the big picture point of view. Don't you feel that creating those species or whatever you may want to call them, synthetic life, even if it's not creating life from scratch, basically starts a new phenomenology in the universe that hasn't been observed before, because you have a complex chemistry that reached the stage at which it actually changes and produces viable complex chemistry that can continue even without its own existence. In other words, if we do not continue as a species and our technological civilization comes to an end, those species will actually continue to exist on this planet potentially could go to other places.

VENTER: Yes it is an important conceptual change

LLOYD: Human beings, you could even say, could identify the beginning of modern human culture with the ability to modify grains genetically so that they could be grown in large quantities and people could go from being simply hunter-gatherers to actually being farmers. And certainly people have been genetically modifying their world around them for tens of thousands of years, in a variety of different ways, starting merely by first picking the variety of grains that have things that are easier to separate out from their husks. The recent discovery of the precursor of corn, or maize, was a big surprise, because the actual precursor looks nothing whatsoever like the corn that we actually eat, and over the course of thousands of years the people merely by selecting out the varieties they wanted made corn as we know.

VENTER: We've been doing blind genetic experimentation by mixing whole genomes together for a long time. It's amazing how little concern there's been over the agricultural practices for the last millennium, right? Just blindly mixing any species together — if they will mix together it gets done. And if you do it with intelligent forethought, it's more dangerous, somehow.

LLOYD: I guess one of the worst consequences of hybridization was that the species that were introduced were so successful that we got rid of a huge amount of biodiversity by replacing what used to be a much larger set of different kinds of corn with just a couple of varieties which are then much more vulnerable to various blights.

VENTER: People don't think about these issues too much — when you start to look at the ocean as a sea of life — a million bacteria per milliliter of bacteria and Achaea — ten million viruses per milliliter — when you think of a supertanker taking on ballast water, going to a port where it fills up with oil, dumping the ballast water, it's put hundreds of thousands of gallons, basically billions to trillions of new organisms that have never seen that environment before, into that environment. And you wonder why we end up with environmental problems such as the zebra mussels in the great lakes.

BROCKMAN: What drives the decision as to what projects you pursue at the Institute? Available funding? Areas that you and your colleagues picked out? What you can get away with? Are you looking for another definition of life?

And finally, what's your fantasy if you had all the money you wanted and needed and you weren't going to get hassled by the government, the press, what would you do?

VENTER: I'm one of the few scientists actually in the situation to live most of his fantasies every day. My institute budget varies between $80 million and $100 million dollars a year in funding, the majority of which comes from federal sources, but an increasing percentage comes from the wealth that's been generated in this country.

The Gordon and Betty Moore Foundation — Gordon was the founder of Intel — is an example of people who have made tremendous private wealth, and put it back into science and to the benefit of society. It's a unique thing in this country

— I don't see it anywhere else in the world. More and more of our research funds come from those sources.

In addition, out of founders stock from Celera and Human Genome Sciences, I created my own endowment that can fund new ideas when they occur not a year or two later. I've found that's the key that's happened in creative science, at least in my case, both having resources and ability to do experiments when I thought of the ideas.

LLOYD: I want you to stop describing this, because I just might have to leave the table if you continue. It's too depressing for me.

VENTER: We could have an entire session on just how dismal new science funding really is in this country. We celebrate the breakthroughs, but to me they happen at one one-thousandth of the rate that they should be happening.

But I've had the privileged situation by creating the environment that I want so myself and my colleagues do things driven totally by our intellectual ideas. And then we spend our seed money on it and then we try and find other sources to fund them to the next stage.

BROCKMAN: What are the ideas that are too dangerous to pursue, that you want to pursue?

VENTER: There isn't anything now within technical capability that's worth doing that is too dangerous to pursue. Our knowledge is so primitive of the human

genome that to start engineering it is just stupid. Hopefully in 50 or 100 years our knowledge will be sufficient that we could do that intelligently. In the long run genetic manipulation of humans is not only inevitable and it's probably a very good idea.

GEORGE CHURCH

> *Many of the people here worry about what life is, but maybe in a slightly more general way, not just ribosomes, but inorganic life. Would we know it if we saw it? It's important as we go and discover other worlds, as we start creating more complicated robots, and so forth, to know, where do we draw the line?*

George Church is Professor of Genetics at Harvard Medical School and Director of the Center for Computational Genetics. He invented the broadly applied concepts of molecular multiplexing and tags, homologous recombination methods, and array DNA synthesizers. Technology transfer of automated sequencing & software to Genome Therapeutics Corp. resulted in the first commercial genome sequence (the human pathogen, H. pylori,1994). He has served in advisory roles for 12 journals, 5 granting agencies and 22 biotech

companies. Current research focuses on integrating biosystems-modeling with personal genomics & synthetic biology.

GEORGE CHURCH: We've heard a little bit about the ancient past of biology, and possible futures, and I'd like to frame what I'm talking about in terms of four subjects that elaborate on that. In terms of past and future, what have we learned from the past, how does that help us design the future, what would we like it to do in the future, how do we know what we should be doing? This sounds like a moral or ethical issue, but it's actually a very practical one too.

One of the things we've learned from the past is that diversity and dispersion are good. How do we inject that into a technological context? That brings the second topic, which is, if we're going to do something, if we have some idea what direction we want to go in, what sort of useful constructions we would like to make, say with biology, what would those useful constructs be? By useful we might mean that the benefits outweigh the costs — and the risks. Not simply costs, you have to have risks, and humans as a species have trouble estimating the long tails of some of the risks, which have big consequences and unintended consequences. So that's utility. 1) What we learn from the future and the past 2) the utility 3) kind of a generalization of life.

Many of the people here worry about what life is, but maybe in a slightly more general way, not just ribosomes, but inorganic life. Would we know it if we saw it? It's important as we go and discover other worlds, as we start creating more

complicated robots, and so forth, to know, where do we draw the line? I think that's interesting. And then finally — that's kind of generalizational life, at a basic level — but 4) the kind of life that we are particularly enamored of — partly because of egocentricity, but also for very philosophical reasons — is intelligent life. But how do we talk about that?

As a scientific discipline, many people have casually dismissed Intelligent Design without carefully defining what they mean by intelligence or what they mean by design. Science and math have long histories of proving things, and not just accepting intuition — Fermat's last theorem was not proven until it was proven. And I think we're in a similar space with intelligent design. What Freeman suggests is that we are moving into a phase which is different, not only in that it's like
Web 2.0 where we're all sharing all of our parts like we used to, but maybe more fundamentally, we're moving into intelligent design big-time, and we need to understand what that means, and what we should be designing.

In terms of utility: people might have huge disagreements about even within a religion as to the right thing to do. You might say, thou shalt not kill, or the same person a few days later might say you must kill — you must kill a lot of people. What might we all agree on? Well, we might agree on that it's not a good idea to wipe out the entire intelligent species in the universe. Even if you believe in the afterlife. Then you might say, well, we shouldn't kill off the afterlife. There's some basic thing that we like, and it has to do with complexity, and what we mean by intelligence. We'd like to preserve that somehow, and I apologize for that being fairly philosophical, but when we start to construct life — maybe not create

it, as Craig has carefully pointed out — what is it we're trying to do? We're trying to make things that are more complex — but it's not just complex, you take a rock: essentially quite complex; if you take a leaf, it's complex in a different way. Take that leaf and then rearrange all its atoms so it's just a bunch of salts in a rock — ammonium carbonate, and silicon dioxide, and potassium phosphate, and sodium sulfate — it's the same atoms but in a form that a mineralogist would recognize. The mineral is still quite complex — if you wanted to transmit over the Internet the structure of that mineral, it would take a lot of bits.

Both Shannon's theory and chemical entropy would say that's a very complicated thing. But what we mean by a living complexity is more like, you've taken something very rare like that mineral, made almost an exact copy of it, and that's replicated complexity. This very unlikely object isn't really any more unlikely than another rock because they're kind of random — their compositional nature is known. But if you made an exact copy of that rock, or nearly an exact one, that would be interesting. That would be indicative that there's some sort of living process, some living thing was involved. It could have been some 3-D photocopier, but that 3-D photocopier was probably made by an intelligent being. It could be that the rock had the ability to replicate.

That's what we mean when we're talking about basic life. And that's sort of what we're trying to get at when we're doing synthetic biology; we're trying to increase diversity, increase replicated complexity, and maintain our ability to continue to do that for many many years, and we don't want to endanger that by doing something that's too risky.

LLOYD: Are you imputing that there's a virtue in increased complexity somehow?

CHURCH: I'm trying to make that argument. There might be a virtue in carefully contemplating not just short-term diversity, but longish-term, to the extent that we can calculate that, which we can't right now. But it's desirable to be able to calculate that as well as we can, and I'll just take a leap at defining intelligence, too, while we're in dangerous territories here for me.

There's analytic intelligence and synthetic. And I would argue that life is sort of this replicated complexity, or mutual information, where given the molecules in this leaf we can predict the arrangement of structures in the other leaf. In other words, we know a lot about this thing — even within the leaf there's replicated complexity that is somewhat predictable — and so that mapping, that mutual information, is predictive life in general. But the mutual information is something where one structure will reflect the structure of something at a distance — especially if you can reflect something distant in time without actually causing it. Intelligence is anticipating things in the future, without causing them. That would be analytic intelligence. Again, it's replicated complexity — or mutual information, even better — where there's a relationship between the two locations, but you anticipate. That's analytic.

Synthetic's harder, because if you synthesize something you've used your analytic intelligence to make a plan and then make a replicated complexity of some sort at a distance, but you've done it. There's a cause and effect. And how do you distinguish between that and, say, the sun having patterned light onto the earth,

and that has a cause and effect which isn't necessarily what we would recognize as intelligence. I'm still struggling with this, but I think synthetic intelligence would be something that in some way or another would enhance the analytic intelligence — the ability of ours to predict what's going to happen in the future. So we synthesize something that will increase our ability to, say, survive as a species, to get off the planet because we know an asteroid is coming ? Various things that we would recognize as long-term intelligent behavior. Something we need to embrace is our ability to do that. And it might be inorganic life that does that, in the form of computers, or maybe some hybrid.

I'll come back down to earth a little more and talk about synthetic biology, but part of getting at intelligence is assaying what kind of intelligence is already on the planet. And what's really remarkable is, we're still far ahead of our computers in that we have 6.5 billion geniuses on the planet. Some of them are undernourished, some of them are undereducated, but there really aren't any computers that are comparable to people — to brains — so far. That may not last forever, but it's certainly true now. We need to assess that diversity and not cure it — as we get better at personalized medicine, the goal is not to cure our diversity, but to enable it — to make it so we can all enjoy our lives and contribute to all the other things we've been talking about today.

Personal genomics is getting into the analytic phase of figuring out what we all are capable of; synthetic biology is still very primitive. We're interested in useful things, like making fuels. Craig already set this up: if we're going to burn carbon, of which there are vast amounts in the earth, we might want to have some way of recovering that carbon and maybe burning it again. If we run out of petroleum, it's

not only useful as a fuel for cars, trucks, and planes, but it's also the source of a lot our constructive plastics. We need a replacement for petroleum, at least in the short term. By short term we may be talking about 30 years, and for most of us that's pretty long-term. One of the companies I have founded is called LS9 in California, and they are making synthetic petroleum as an example of what you can do with synthetic biology. Hydrocarbons, which are compatible with current engines — cars, diesels and jets — rather than requiring a new infrastructural change, and that's basically metabolic engineering and exchanging DNA from many organisms just like Freeman said — not just one gene at a time, but whole systems, whole metabolic systems, and using what we know to expedite that.

Another company I've been involved in is called Codon Devices and it makes the DNA for companies like LS9. They have a capacity of about two million base pairs of very highly polished DNA per month, which goes out to many biotech pharmaceutical companies. It's not making whole genomes, or whole chromosomes, like Craig is talking about, but it has the capacity to make several such genomes per month. And it is important to have this, and to have it regulated. As we started my part of this conversation, I really think that's important, and Codon Devices is joined with about a dozen other companies to increase their ability monitor DNA synthesis worldwide. I started this with a white paper many years ago, tried to bring government attention to this, and it became clear that the government isn't going to act unless it's sure that industry isn't going to be hurt. It isn't always a formula that will work — but if industry can get its act together and do it voluntarily — it's not really self-monitoring, it's just trying to get something going — then governments can point to it and say, oh, that's already working, they're sharing resources, so the cost of software is distributed, the cost of

monitoring local regulations is distributed, and then the government can say, okay, we'll make the law. At which point internationally they can say, many governments are acting; we can make international law. I hope that's the way that these things work out.

That's synthetic biology, making biofuels, helping improve stem cell biology; you don't necessarily have to start with a germ line if you can change people's soma. That may be less risky, it may allow you to do more rapid prototyping and alleviate people's concerns. The germ line is a very special set of concerns, and there's quite a bit you can do — there's actually more you can do in a way — because predicting what somebody's going to be like from a fertilized egg is very hard, but predicting what they're going to be like once they're 30 years old is actually quite a bit easier. It's not really prediction, it's just observation. It may be harder to fix, but at least you'll develop tools and you'll be more cautious.

There's a great deal of progress there, with now being able to establish more pluripotent cells from, say, mammalian skin cells, which will enable synthetic biology to move in that direction. And there's even some projects that we're working on as part of the NSF-funded project with Berkeley, USCF, and MIT, to engineer bacterial cells that are compatible with mammalian immune systems so they can motor around inside your blood stream and do sensing and actuating. For example, they will hone to tumors, they will sense their presence (there will be a thousand times higher concentration near a tumor) — all the parts of this are working, but the whole is not working — and they'll sense that they're there, they'll invade the cell by expressing an invasive protein, and then inside the cell they will make a drug that will destroy the tumor cell. That kind of capability of

working well within a mammalian immune system can range from using your own cells that are perfectly immune-compatible to using these engineered bacterial cells.

There are many other things I could say, but they could be more easily said in the context of questioning. Hopefully I've brought up enough provocative points for you to ask interesting questions.

BROCKMAN: How is you work different that Craig's?

CHURCH: Well, he's much more productive.

VENTER: I use George's techniques.

CHURCH: There, isn't that sweet of him? We develop technology for the most part. Usually we try to enable other groups to do production, so I'll start a company or I'll work with some genome centers to get our sequencing or synthesis technologies to work. From a synthetic standpoint, the major difference is that Craig is a little more interested in making a synthetic genome from scratch; we're mainly interested in making variations on genomes, although I'm sure both of us would be comfortable doing the other one as well.

He mentioned making combinatorial a million chromosomes, well, for instance, we do that, we do lab evolution, and we make millions of chromosomes that go into competition with one another. You either do it by recombining every base pair, spontaneously, in which case you can only get one mutation at a time —

we've gotten up to three or four mutations serially in that way — or we can do site-directed mutagenesis, and we have a new automated method of doing that. Or we can get a series of 23 mutations in nine days — one at a time, for up to three hundreds of days.

One thing that is interesting that I didn't mention — and Craig was kind of getting at, whether we're building things up from atoms, and Bob Shapiro mentioned chirality — is we are trying to make a mirror image biological world starting with DNA polymerase, and that does in a certain sense require starting at a much more fundamental level than synthetic biologists usually do. Synthetic biologists in our classes will PCR genes up by the cells — that's very dependent on life. Craig and I will synthesize genes from nucleotides — that's basically just dependent on knowledge, and a few chiral centers, but if you actually flip the chirality, now you're really closer to dealing with atoms.

LLOYD: Do you do this for safety reasons?

CHURCH: That's right, Seth: for every one of these things, we should ask why we're doing this, okay. Why we're doing molecules is obvious. Why we're doing stem cells and pharmaceuticals is obvious. We're changing the genetic code in the normal chirality for safety reasons, and to extend the number of amino acids. A chirality, you could do the same thing, in each of these cases you have to disable it in some way, because changing the chirality makes it incompatible to the rest of the world, but that can make it more or a threat or less of a threat depending on what other things you do. It could be more of a threat because now not only is it resistant to phages it's resistant to enzymes, like proteases and ribonucleases, and

at least existing antibodies. Now if you put in a mirror image cell, you would get new antibodies and that's not a problem.

The other reason we're interested in it, aside from safety, which is always something we're interested in and a major theme in synthetic biology and Codon Devices, is, let's say you can evolve DNA and RNA molecules that will bind to your favorite thing. In a certain sense this is morphogenesis from scratch: you make a completely random selection of polynucleotides and you can get some that will bind your favorite surface or molecule. In a way you didn't define that surface, you found it, in a random collection. That's getting close to what we'd like to get.

But when you use those in a practical setting, they get degraded by biological fluids. But if you made the mirror-image form of it — and since you're starting from scratch, you don't have any preconceived rules, you're just evolving by selection of binding — you can start with a mirror-image nucleic acid set, a library containing trillions of molecules, and you'll find something that binds your favorite molecule and is resistant to the enzymes. That's one motivation for the first thing we're making, which is DNA polymerase. We want to be able to mirror image PCR, polymerase chain reaction, where you can amplify up DNA with a mirror-image polymerase. A post-doc has gone through a prototype polymerase which is for a medium size — 353 amino acids long — and he's made all but four of the peptide bonds now. So we're very close to getting that first polymerase.

The mirror image nucleotide part is actually fairly simple because you can use the same machines to make mirror image DNA. The peptide synthesis machines are

much more primitive than the DNA synthesis machines, so a lot of this is done currently by hand. But the goal then so we can make mirror image DNA is to make mirror image proteins, and there we have to make all of the ribosome from scratch, and so we wouldn't make all the ribosome just for fun, but this is something that we think is useful. And that's about 25 times more bonds to be made than just making the DNA polymerase. But, as Craig will greatly attest, scaling up by a factor of 25 is not that big a deal. In the genome project we went basically a hundred thousand fold scale-up from where we were at the time we started. And now we're talking about doing many many genomes. I think we will be able to make a mirror image DNA polymerase and ribosome, in which case you can start programming it straight from the computer. Once you have it all you can start making mirror-image proteins.

VENTER: There's the big assumption that the mirror image of all these things will have the same activity.

CHURCH: Will have the mirror activity.

VENTER: Will have the mirror activity on the other chiral molecules. Is there any evidence for that?

CHURCH: There is, a little bit. I wish there weren't. Then we could get all the glory — but I think the HIV protease has been made in mirror form and has been shown to be inhibited by things that are in the mirror form. Very small things have been made — and crystallography has shown that the mirror image polymers — they make up the mirror-image monomer; the mirror-image polymer is flipped.

Almost every time I mention this there's a subset of people who feel that it couldn't be otherwise, and there's a subset that say, Prove it. I'm happy either way.

SASSELOV: That brings me to a question here. It seems like when you talk about synthetic biology, you feel that in the next few years things will evolve to what Craig called "life from scratch". Do you think there will be a clear watershed along the way, or will it really get there incrementally?

CHURCH: Well almost certainly it's going to be incremental. There'll be many milestones. Certainly Craig's article in *Science* "Transplantation", was a milestone. When he does it with synthetics, that will be one. Another will occur if we manage to get synthetic ribosomes that are in a chiral form. There'll be many milestones, but every one of them you can trace to something incrementally similar in the recent past.

SASSELOV: But you feel there is not one big gap, which needs to be crossed and then you are there.

VENTER: The gap of taking inanimate objects and not starting with life and getting life from it is a hurdle that when crossed — it's inevitable that it will be crossed, and relatively soon, somewhere, but intellectually it's not a gap, but I think until it's done it's a big conceptual gap to society.

CHURCH: That's almost philosophical, whether you've got the atom from carbon dioxide or whether you got it from Ribose, but I see most of the gaps as practical

ones. This field will go faster the more useful it is, and people will resist it less the higher the benefit-to-cost and risk ratio. Most people accept evolution, even creationists accept microevolution. If we start getting macroevolution in the lab, then they'll accept the macroevolution to whatever extent it is useful and obvious. If it's not demonstrated in the lab then you might reasonably say 'I don't care', or, 'Prove it'. The scientists should be saying, ‘Prove it.' 'Do it in the lab.' Now some things you argue can't be, but I actually think macro-evolution is something that might be possible — were certainly doing micro-evolution big-time — a lot of companies now depend on pretty amazing changes in the structure — you could argue though that they all have intelligent design somewhere in the process, but I think the less intelligent the design and the more macro the evolution, the more people will accept it as acceptable.

BROCKMAN: Can you talk about Bio-Fab Labs and their self-replicating nature — in terms of the discussion about fooling around with the human genome and playing God?

CHURCH: Craig pointed out that some religions think that humans should do technology, whether you want to call it playing God or working at it, or — you're certainly not creating a universe. You're constructing things.

VENTER: You're only so big.

CHURCH: Pretty small. Pretty small. So you know, Bio-Fab Labs are very much a continuation of all the other engineering disciplines — civil engineering, electrical engineering, mechanical engineering, chemical engineering. Ironically,

genetic engineering was really not what most engineers would recognize as an engineering discipline when the term was coined. They do recognize, or they are part of, the revolution that's now finally making it an engineering discipline, with interchangeable parts, hierarchical design, interoperable systems, specification sheets, that kind of stuff. Stuff that only an engineer could love.

BROCKMAN: What is the difference between Neil Gershenfeld's Fab Lab and your Bio-Fab Labs.

CHURCH: I was just at the annual meeting of the Fab Labs of the world, that Neil Gershenfeld organized in Chicago, and I did make that comparison. On the plus side, the current generation of Fab Labs interoperate well with computers, while biology basically doesn't — with the exception of what Craig and I have been talking about.

Take native biology. It's very hard to stick a WiFi onto a corn plant. While in Fab Lab it's all about that, interacting. On the positive side, despite some efforts, there is no inorganic or non-life technology, despite very sophisticated Fab Labs, capable of making itself. A Fab Lab can't make itself without a huge amount of human intervention, which is something that the most elementary bacterium can do. And even with human intervention there's not some compact Fab Lab; it's something that's spread out over continents — there'll be one place that makes the integrated circuits, another place that makes the nice steel bars that you use, another place that turns petroleum into plastic, and so forth. It's not an obsession with them, but something that Fab Labs toy with is making a compact desktop device that could make copies of itself. They, to some extent like us, want to have

an open source environment. They already do this in the Fab Labs — they'll send over the Internet plans for making a chair, or a house, and they'll make it, in another country, without actually physically transferring a person or a device. That's very exciting and something that we share in common.

PRESS: It sounds like the kinds of tools you're talking about can also be very useful for actually doing some sort of hard experimentation, or at least testing different theories of the origin of life. Not starting from the bottom up, you can at least approach it from the top down and sort of pick apart the different models — is that a direction that you're involved in, or are the people using your tools doing that?

CHURCH: I'm a little more interested in the future than the past — but I don't dismiss it either. For example, on the top of Freeman's wish list was ribosome archaeology. And Dimitar asked, Is there a milestone that we think is significant.

The ribosome, both looking at the past and at the future, is a very significant structure — it's the most complicated thing that is present in all organisms. Craig does comparative genomics, and you find that almost the only thing that's in common across all organisms is the ribosome. And it's recognizable; it's highly conserved. So the question is, how did that thing come to be? And if I were to be an intelligent design defender, that's what I would focus on; how did the ribosome come to be?

The only way we're going to become good scientists and prove that it could come into being spontaneously is to develop a much better in vitro system where you

can make smaller versions of the ribosome that still work, and make all kinds of variations on it to do really useful things but that are really wildly different, and so forth, and get real familiarity with this really complicated machine. Because it does a really great thing: it does this mutual information trick, but not from changing something kind of trivial, from DNA to RNA; that's really easy. It can change from DNA three nucleotides into one amino acid. That's really marvelous. We need to understand that better.

VENTER: And you can't have life without it.

CHURCH: Definitely. It's common to all life. We need to understand that, and the way we're going to fund it — there's not that much funding for prebiotic science, but if there's a lot of funding for understanding the ribosome in the future and in the present, inevitably it will much enable studies of it in the archaeological and ancient biology sense.

VENTER: But using these tools, it's my hope we can do something similar to what you suggest. We can extrapolate back once we have the database of Planet Earth genes to what might have been a precursor species, and then we should be able to build that in the lab and see if it was really viable, and then start to do component mixtures to see if you can spontaneously generate such things.

CHURCH: But isn't it the case that, if we take all the life forms we have so far, isn't the minimum for the ribosome about 53 proteins and 3 polynucleotides? And hasn't that kind of already reached a plateau where adding more genomes doesn't reduce that number of proteins?

VENTER: Below ribosomes, yes: you certainly can't get below that. But you have to have self-replication.

CHURCH: But that's what we need to do — otherwise they'll call it irreducible complexity. If you say you can't get below a ribosome, we're in trouble, right? We have to find a ribosome that can do its trick with less than 53 proteins.

VENTER: In the RNA world, you didn't need ribosomes.

CHURCH: But we need to construct that. Nobody has constructed a ribosome that works well without proteins.

VENTER: Yes.

SHAPIRO: I can only suggest that a ribosome forming spontaneously has about the same probability as an eye forming spontaneously.

CHURCH: It won't form spontaneously; we'll do it bit by bit.

SHAPIRO: Both are obviously products of long evolution of preexisting life through the process of trial and error.

CHURCH: But none of us has recreated that any.

SHAPIRO: There must have been much more primitive ways of putting together

catalysts.

CHURCH: But prove it.

VENTER: You need to improve DNA synthesis a little bit more so that it's 3 or 4 orders of magnitude faster. Then you can make a seemingly infinite pool of nucleotides and start to get — to me the key thing about Darwinian evolution is selection. Biology is a hundred percent dependent on selection. No matter what we do in synthetic biology, synthetic genomes, we're doing selection. It's just not natural selection any more. It's intelligently designed selection, so it's a unique subset. But selection is always part of it. We're not that far away from being able to do these experiments. It's very hard to do now, because nobody would spend the money to make all these different related molecules to see if we can get spontaneous ribosome formation, but within a decade it will be doable.

LLOYD: I would be a little bit worried. If I look at Freeman's two steps that precede formation of ribosomes, ribosome is step three, with collaboration and intervention of the ribosome and you have these two steps prior to that. Before is the parasitic stage and use of ATP, and then prior to that just the kind of garbage bags on their own. There could have been a lot of events of natural selection going on to get to the stage — it could be a very very very long process, with Avogadro's number of events. There are not enough graduate student lifetimes in the world, even with lots of private money invested in it to actually try to explore all those. Even if life just happened here on Earth, it is something that happened globally, and it went on for quite a long time. I'm saying this is a positive sense — the fact that you can't do this in the lab, even though people who do intelligent

design will say, 'ah ha, see, irreducible complexity' — it might in fact be that it was very complex in the sense of requiring a long and complicated process or computation to arrive at.

CHURCH: What we can do in the lab, though, is to reconstruct intermediates and characterize them and say, okay, here's something we found valuable in the lab that has fewer proteins, a slightly different reaction — and make a plausible time line to say, okay, given that there were 10^44 water molecules on this planet and we can't reconstruct that in the lab, and we can't use the water molecules on the planet to do a *BR lab*, maybe there was a fairly small number of environments that actually did the trick and if we construct intermediates that are kind of convincing, then we could do small pieces of that pathway in lab time frames.

LLOYD: Well that would certainly be my hope. I'm just saying that that's a hope.

VENTER: But the power of selection can give you at least 7 to 8 orders of magnitude of selectivity.

LLOYD: Sure, absolutely — it's definitely worth doing, absolutely, yes.

VENTER: That's what came out of the phi X work: you can make ten to the 6, 10^7 different molecules out of the assembly, and if they're viable you'll select them.

CHURCH: But Seth was saying that if we try to do the whole process, from primordial soup to ribosomes, we haven't got 10^8 leaders times 10^9 years to do

it.

LLOYD: You're trying to reproduce this metabolic phase that Freeman was talking about, and in some sense once you're already at the digital phase, as you yourself said, life is not just genes, it's the machinery required to take those genes and then reproduce them, which means viruses and cells, and somehow monkeying with the program might be easier than creating viable new programs from analogs of things that are out there. This seems to me potentially easier than trying to construct this process that led up to ribosome where you don't even know what it was at all in the first place.

DYSON: You have to look for something that components of ribosomes might have done in order to evolve.

LLOYD: Right — there's the example of rhodopsin that you came up with, which provides an inspiring example. It could easily be that the rhodopsin showed up as an earlier version of photosynthesis — less efficient, but hey, still good for harvesting energy — and then only later, oh, look, it also could be used as a sensor — discovered by natural selection. Somehow natural selection is full of all these tricky little switches, which makes it very hard to trace back what happened.

VENTER: Well, we can trace back to where it switched from light to chemicals, and became the key driver of nervous systems.

LLOYD: Actually, Bob, I'd be interested in what do you think if you had to bet on the success of this particular venture of trying to recreate a ribosome from scratch

— trying to come up with a pathway.

SHAPIRO: You can synthesize in the laboratory a ribosome from scratch, undoubtedly.

CHURCH: You mean evolve a ribosome?

LLOYD: Evolve a ribosome.

VENTER: We have synthetic ribosomes in our lab, they're just not totally efficient right now. We didn't design them; we're copying the design.

SHAPIRO: What I would say, and Freeman is probably in my camp, but I hear as I listen around the table example after example of what I call DNA-centric thinking. Of equating life with DNA.

My problem is I know too much about DNA, I spent my life in DNA chemistry, and to me it looks like a highly evolved organism. Life started without DNA, without RNA, and undoubtedly without proteins, and was yet alive. And in fact it's the same error that was made when microbiologists only characterized what they could clone up may be made by only identifying what's alive by looking at its DNA content, because you may then be missing things that are alive but lack DNA.

Life undoubtedly had to start with what nature gave us and there's a different approach, which is called the bottom-up approach, where you try to use physical

principles and ask what would what we would regard as inanimate matter do when subjected to an appropriate environment and some liberal supply of free energy, and what combinations of those might actually work to kick off the living process.

Now we're an example of one very successful conclusion of the living process but not necessarily the only example, nor need life necessarily have our exact set of components. There's famous set of experiments from about ten years ago when Albert Eschenmoser, a brilliant Swiss synthetic chemist, set out to prove why nature had a select DNA. With enormous Swiss skill and manpower he set students out to make DNA-like molecules using different sugars, one after the other, expecting that in every instance he would fail. But in fact he succeeded and he found that different sugars in many cases was superior to DNA. They had greater stability; they had fewer complications in replication.

I thought that he would arrange to have the Swiss government declare that from now on every Swiss life form would adapt his symbiosis and dispense with DNA as quickly as possible. There's PAN, and someone else came up with TNA — there's endless ones — and so to me DNA is probably what evolution stumbled upon through accident, and it's the easiest thing that could be come upon by slow trial and error that would make a molecule that could be replicated by proteins and that's how it came into being.

Now to me, first as an imagination experiment but ultimately in laboratory experiments would be to try and see where else starting from simple chemicals and energy, you might go in the direction of evolution.

ROBERT SHAPIRO

I looked at the papers published on the origin of life and decided that it was absurd that the thought of nature of its own volition putting together a DNA or an RNA molecule was unbelievable.

I'm always running out of metaphors to try and explain what the difficulty is. But suppose you took Scrabble sets, or any word game sets, blocks with letters, containing every language on Earth, and you heap them together and you then took a scoop and you scooped into that heap, and you flung it out on the lawn there, and the letters fell into a line which contained the words "To be or not to be, that is the question," that is roughly the odds of an RNA molecule, given no feedback — and there would be no feedback, because it wouldn't be functional until it attained a certain length and could copy itself — appearing on the Earth.

Robert Shapiro is professor emeritus of chemistry and senior research scientist at New York University. He has written four books for the general public: *Life*

Beyond Earth (with Gerald Feinberg); *Origins, a Skeptic's Guide to the Creation of Life on Earth*; *The Human Blueprint* (on the effort to read the human genome); and *Planetary Dreams* (on the search for life in our Solar System).

ROBERT SHAPIRO: I was originally an organic chemist — perhaps the only one of the six of us — and worked in the field of organic synthesis, and then I got my PhD, which was in 1959, believe it or not. I had realized that there was a lot of action in Cambridge, England, which was basically organic chemistry, and I went to work with a gentleman named Alexander Todd, promoted eventually to Lord Todd, and I published one paper with him, which was the closest I ever got to the Lord. I then spent decades running a laboratory in DNA chemistry, and so many people were working on DNA synthesis — which has been put to good use as you can see — that I decided to do the opposite, and studied the chemistry of how DNA could be kicked to Hell by environmental agents. Among the most lethal environmental agents I discovered for DNA — pardon me, I'm about to imbibe it — was water. Because water does nasty things to DNA. For example, there's a process I heard you mention called DNA animation, where it kicks off part of the coding part of DNA from the units — that was discovered in my laboratory.

Another thing water does is help the information units fall off of DNA, which is called depurination and ought to apply only one of the subunits — but works under physiological conditions for the pyrimidines as well, and I helped elaborate the mechanism by which water helped destroy that part of DNA structure. I realized what a fragile and vulnerable molecule it was, even if was the center of

Earth life. After water, or competing with water, the other thing that really does damage to DNA, that is very much the center of hot research now — again I can't tell you to stop using it — is oxygen. If you don't drink the water and don't breathe the air, as Tom Lehrer used to say, and you should be perfectly safe.

However, around the year 1980 a physicist friend, Gerald Feinberg, now deceased, needed an organic or biochemist to collaborate with on the book he wanted to write called "Life Beyond Earth," which was based on his conception of physical principles. And he loved to think in exotic terms. One of his favorite inventions was a particle he named the tachyon.

Now it's known that there's a barrier at the speed of light, and something that's moving slower than the speed of life can't exceed it, but he conceived of a particle that's moving faster than the speed of light and can't slow down to the speed of light. Something like the bus in the movie "Speed." Except that was about 59 miles an hour, which is a little slower than light.

He named that the tachyon, and that was part of his claim to fame. He had quite an imagination, and by contact with him I developed an imagination too and we began to think, what is there that's special about life, that binds it not only to DNA and RNA and proteins, but to even carbon chemistry, or, in his mind, and I'm not prepared to speak on his topics, to chemistry at all.

In this book he liked to fancy life and liquid helium, plasma life in the center of the sun, laser life, which is all energy and didn't depend on matter at all, and we had fun, but the idea was to shake up conventional thinking of biologists where

they only would recognize as life something that could be cultured by them and then published. Albert Sangiorgi once said that a drug is something that, injected into an animal, produces a paper. A microorganism is something which when put in one of its favorite culture media leads to a paper. Nowadays you might say it leads to a DNA sequence, which would be a different argument.

VENTER: But eventually a paper.

SHAPIRO: Eventually a paper. The term exobiology was coined by Joshua Lederberg in the early 1960s to describe any life forms found outside of Earth. In the 1990s it got resurrected at astro-biology, which defined itself as the science that encompasses all life, past, present, and future. Which swallows all of biology and then burps — and if anything exo is ever found, it will swallow that too. So it can't be accused, as exobiology once was, of being the only subject that had no subject matter to study.

In any event, I helped him write that book, and tried to use my imagination as best I could to envisage other life forms, different than Earth.

For example, water, thought so precious as a solvent, is hardly indispensable. In fact, for organic chemists, almost all of the famous reactions in most of the chemistry they do are run in solvents in which water is rigorously excluded in and the first step you take in doing the work up is to shake your product solvent, which might be chloroform or something like that, with water, and then you throw away the water because it has the uninteresting stuff. So that most known organic chemistry is the chemistry of gasoline, so to speak. And water chemistry is the

exception. That being the case, it was not hard to envisage life functioning in organic solvents.

Of course at that point there was no thought that there might be any organic solvent anywhere in the universe. Now we know otherwise — there are lakes of liquid methane on Titan. Pretty chilly, but part of our thinking was to free ourselves from what you might call the tyranny of the covalent bond — thinking that covalent bonds were the only thing of which life could be constructed. And at low temperatures, much weaker bonds that play a minor role in Earth biology, unless they're multiplied by the thousands, as they are in DNA, could fulfill the function of covalent bonds, whether a hydrogen bond broke, or made, might be momentous event.

Or at other higher temperatures, most carbon compounds are incinerated or fly apart and the bonds break, but silicates are perfectly stable so I can see magna life — life in molten silicon — living inside of the Earth.

Now some of these things would be difficult to test for, but the idea to get across is that there's nothing sacred about the idea of water — that a fluid would do anything that could promote reactions or interactions. And then we went to carbon and decided that yes, carbon was very convenient. As Gary used to say, all of my best friends are made of carbon. But we couldn't see why it was indispensable, in the realm where there's a hundred plus elements, some of which could bond with others, silicon-silicon bonds are weaker than carbon, but they might function at weaker temperatures. Silicon-to-oxygen-to-silicon bonds are much stronger than carbon-carbon bonds and they could function at higher temperatures.

At this point the Viking experiments hadn't been run, and having explored Mars and come out with ambiguous results, well, I went to American Association for Advancement of Science meeting and a congressman was the key speaker, and he said, you gave us a billion dollars for the Viking mission to tell us whether or not there was life on Mars, and we have demanded the answer, and 'no' is an acceptable answer, but the one thing we won't tolerate is a request for another billion. Thank you.

NASA promptly moved into the No, there is no life on Mars position. There's something called the Munch Report, which talked about all the interesting geology that could be done on Mars without paying any attention to life, and there was a whole decade in which nothing flew to Mars, because to investigate the geology of Mars, wind erosion, water without reference to life, well, there was no public interest and no congressional funding.

In the interim, looking for topics to get amused with, I got into the question of the origin of life, and knowing the DNA chemistry that I did know — and helped write — I looked at the papers published on the origin of life and decided that it was absurd that the thought of nature of its own volition putting together a DNA or an RNA molecule was unbelievable.

I'm always running out of metaphors to try and explain what the difficulty is. But suppose you took Scrabble sets, or any word game sets, blocks with letters, containing every language on Earth, and you heap them together and you then took a scoop and you scooped into that heap, and you flung it out on the lawn

there, and the letters fell into a line which contained the words "To be or not to be, that is the question," that is roughly the odds of an RNA molecule, given no feedback — and there would be no feedback, because it wouldn't be functional until it attained a certain length and could copy itself — appearing on the Earth.

Christian de Duve, the Nobel laureate, once wrote a letter to *Nature* which was headed, 'Did God Make RNA?' Because it's hard to think of any other manner in which RNA out of purely abiotic chemistry would assemble itself on the early Earth. Seeing this area called prebiotic chemistry, I decided, my major attention still being funded by National Cancer Institute, and devoted to how chemicals can rip apart DNA as a hobby essentially, I started publishing papers which disassembled — deconstructed, if our German friend wants — so-called prebiotic chemistry, and showed that in every case the result was due to the flagrant interference of the investigator in biasing the results to attain the results that he wanted.

At one point I went and spoke to the now, unfortunately, late Stanley Miller, and asked him about the circumstances of his famous Miller-Urey experiment — the one with the electric lightning and amino acids were formed — and he handed me a biographical piece he himself had written to something called the *Transactions of the Copernican Society* or something like that, and he described how in building his apparatus he was concerned with questions of safety, because if you take a flask and you mix it with methane and hydrogen and ammonia, the most likely result is BOOM, with flying glass in all directions, which is definitely not publishable.

With regard to safety, he built a certain apparatus, let it run for a number of days, and at the end of the days he looked at what he'd found and he found the class of chemicals called hydrocarbon — the stuff that makes up the lakes on Titan but no amino acids whatsoever. And he looked at this and he said, this isn't interesting. And he threw it out. He redesigned the equipment: he said, I was over-cautious. This is not likely to explode. He interchanged the spark and the condenser and he re-ran the experiment, and this time he got amino acids and not hydrocarbons, and he said, Ah ha! And he published.

Thus we have the famous Miller-Urey experiment showing the inevitability of amino acids on the primitive Earth. And of course the apparatus itself has no resemblance whatsoever to the primitive Earth. One of the popular magazines said that if his apparatus had been left on for a million years, something like the first living creature might have crawled out of it. And I say, if he'd left his apparatus on for a million years, he would have run up one hell of an electric bill. But nothing further would have happened because the spark was in the atmosphere and he'd used up all of the chemicals with carbon in the atmosphere, and the amino acids, which aren't volatile — they don't fly, so to speak — were safely ensconced in the water solution, and the water solution was a collection of non-volatile compounds, well, and the volatile compounds ended up in — so when an experiment goes wrong in organic chemistry you get a black gook and you reach for the potassium bichromate and sulfuric acid — mixed together it's a called cleaning solution — that cleans out about 90 percent of the failed organic experiments that are ever run.

You use that and you can get rid of the tars in about 80 to 90 percent of his

carbon, this stuff that had unfortunately flown again and again until it got zapped and ended up as tars on the wall of his flask. Well, this was the best prebiotic experiment ever run, because at least he started with components that hypothetically could have been on the early Earth.

Since then, so-called prebiotic chemistry, which is of course falsely named, because we have no reason to believe that what they're doing would ever lead to life — I just call it 'investigator influenced abiotic organic chemistry' — has fallen into the same trap. In the proceedings of the National Academy of Sciences about two months ago there was a paper — I think it was theoretical — they showed that in certain hydro-thermal events, convection forces and other attractive forces, about which I am unable to comment, would serve to concentrate organic molecules, so that organic molecules would get much more concentrated in the bottom of this than they would in the ordinary ocean.

Very nice, perhaps it's a good place for the origin of life, and interesting finding, but then there was another commentary paper in the Proceedings by another invited commentator, who said, Great advance for RNA world because if you put nucleotides in, they'll be concentrated enough to form RNA; and if you put RNA in, the RNA will come together and form aggregates, giving you much more chance of forming a ribosome or whatever. I looked at the paper and thought, How did nucleotides come in? How did RNA come in? How did anything come in?

The point is, you would take whatever mess prebiotic chemistry gives you and you would concentrate that mess so it's relevant to RNA or the origin of life — it's

all in the eye of the beholder. And almost all of prebiotic chemistry is like this; they take chemicals of their own selection.

People were talking about Steve Benner and his borate paper where he selected, of his own free will, the chemical formaldehyde, the chemical acid-aldehyde, and the mineral borate, and he decided to mix them together and got a product that he himself said was significant in leading to the origin of RNA world, and I, looking at the same thing, see only the hands of Steve Benner reaching to the shelf of organic chemicals, picking formaldehyde, and from another shelf, picking acid-aldehyde, etc. Excluding them carefully. Picking a mineral which occurs only in selective places on the Earth and putting it in in heavy doses. And at the end getting a complex of ribose and borate, which by itself would be of no use for making RNA, because the borate loves to hold onto the ribose, and as long as it holds onto the ribose it can't be used to make RNA. If it lets go of the ribose, then the ribose becomes vulnerable to destruction by all the other environmental agents.

The half-life of pure ribose in solution, a different experiment and a very good one, by Stanley Miller is of the order of one or two hours, and all of the other sugars prominent in Earth biology have similar instability.

I was publishing papers like this and I got the reputation, or the nickname in the laboratory of the prebiotic chemist, of 'Dr. No'. If someone wanted a paper murdered, send it to me as a referee. And so on. At some point, someone said, Shapiro, you've got to be positive somewhere. So how did life start? And do we have any examples of authentic abiotic chemistry, not subject to investigator

interference?

The only true samples we have are those meteorites, which are scooped up quickly and often fallen in an unspoiled place — there was a famous meteorite that fell in France in a sheep field in the 1840s and led to dreadful chemistry of people seeing all sorts of bio molecules in it, not surprisingly. But if you took pristine meteorites and look inside, what you see are a predominance of simple organic compounds. The smaller the organic compound, the more likely it is to be present. The larger it is, the less likely it is to be present. Amino acids, yes, but the simplest ones. Over a hundred of them. All the simplest ones, some of which, coincidentally, overlap the unique set of 20 that coincide with Earth life, but not containing the larger amino acids that overlap with Earth life. And no sample of a nucleotide, the building block of RNA or DNA, has ever been discovered in a natural source apart from Earth life. Or even take off the phosphate, one of the three parts, and no nucleoside has ever been put together. Nature has no inclination whatsoever to build nucleosides or nucleotides that we can detect, and the pharmaceutical industry has discovered this.

Life had to start with the mess — a miscellaneous mixture of organic chemistry to begin with. How do you organize this? You have to have a preponderance of some chemicals or lacking others would be against the second law of thermo-dynamics — it violates a concept that as a non-physicist that I barely grasp called 'entropy'.

How does one get around this? Entropy is like a business. It doesn't matter if one subsidiary of the business loses money as long as the others show enough profit to offset it. What you need is a larger system, the environment, and part of it absorbs

energy and gets organized, and in payment for that, the rest of the environment gets disorganized, usually by going up a little bit in temperature, which is the common denominator of entropy. If you convert other kinds of energy to heat, you can pay for a lot of organization.

With just these concepts and with a lot of papers scattered in the literature that hadn't been pulled together, including some excellent ones by Harold Morowitz, a physicist who can be very eloquent on the subject. You get the idea that life could start in mixtures of simple molecules, provided that the organization of these molecules was intimately connected to the release of some energy as heat if there was coupling.

In the simplest case, and there may be many more elaborate cases, they found that the energy wouldn't be released unless some chemical transformations took place. If the chemical transformations took place then the energy was released, a lot of it is heat. If this just went on continuously, all you do is use up the energy. Release all of it and you've converted one chemical to another. Big deal. To get things interesting, you have to close the cycle where the chemicals can be recycled by processes of their own, and then go through it again, releasing more energy. And once you have that, you can then develop nodes — because organic chemistry is very robust, there are reaction pathways leading everywhere, which is why it's such a mess.

Well the fact that this cycle, this energy-driven cycle, was working would suck the material out of all these other side-reactions into the main cycle. Occasionally some step, due to change in environment, and the main cycle would be blocked:

the acidity would change, the temperature would change, and then the machine would cease to turn. But in wandering around the many pathways, a bypass might be found and the cycle would reestablish itself again. And in searching for that bypass some catalyst might be found that would unblock the main cycle.

The main idea is that you get a network of reactions, which would feed into each other, and the net result would be that the latent energy would be released. One of the most familiar kinds of chemical energy for the chemist is redox energy; you get an electron-poor reagent (an oxiding agent technically), an electron-rich reagent (a reducing agent technically), and the electron-rich reagent gives electrons to the electron-poor reagent.

In planets from the Earth are subject to a sort of tension because the inside, which is pure iron, is very electron-rich, while the outside, due to continual escape of hydrogen into space because water gets broken up by radiation, is electron-poor, so that at various places on the Earth there will be interfaces where electron-rich molecules are interfaced with electron-poor molecules. These are then prominent sites for the origin of life. Everyone can have his own favorite site. Some argue for the interiors of volcanoes, some argue for vents, some argue for the monolayer of the space of the ocean.

The idea is that this is inherent in the laws of chemistry and physics. One doesn't need a freak set of perhaps a hundred consecutive reactions that will be needed to make an RNA, and life becomes a probable thing that can be generated through the action of the laws of chemistry and physics, provided certain conditions are met. You must have the energy. It's good to have some container or compartment,

because if your products just diffuse away from each other and get lost and cease to react with one another you'll eventually extinguish the cycle. You need a compartment, you need a source of energy, you need to couple the energy to the chemistry involved, and you need a sufficiently rich chemistry to allow for this network of pathways to establish itself. Having been given this, you can then start to get evolution.

This segment is part of the environment, which you can call an organism, a cell, a life-container, starts to change to adapt, become more efficient at harnessing the energy. If, as Freeman Dyson said, you want to diversify, well, if it were a spherical object, something might shatter it into two pieces. Each piece would contain half of the necessary set of molecules; it would have two daughters. This is all analog. Digital information storage is nowhere in it.

One way of finding out, to make this clear, who is at this table today would be to hand around a piece of paper and have everyone write on a list their names, at the end of which you'd have a list which has the attendants at this table. That's DNA. However much more simply, one knows who's at this table if you're familiar enough with the people by looking around the table. And I can see that Freeman Dyson, for example, is present, but another origin-of-life chemist name Jack Sostack is definitely not present. Craig Venter is present, but many DNA sequences are not. Fred Sanger, who is the person who started that ball rolling, is definitely not present; he retired when he was 65. But this is information, too. Information just stored by the presence of the people at the table, and this was called by Doron Lancet, an Israeli theoretician, the compositional genome.

And the other part of life that involves giant molecules are the enzymes, which are wonderful catalysts, but at the start of life you may have been able to get by without wonderful catalysts — you don't need to speed things up a million-fold, which is what enzymes can do. Small molecules can in instances speed things up a thousand-fold, and if you have the right system a thousand-fold may be enough. The picture that emerges is that in a planet like Earth there may have been dozens or hundreds of separate starts of Earth — we are all the children of the most successful one, which as far as we know, dominates on the Earth. We need not be the only life form present on the Earth.

Paul Davies and others have started research, much cheaper than the space program, to try to detect life that has no DNA, and once I talked to a person who also knew something about life and bacteria, namely Joshua Lederberg, and he said an experiment he'd always wanted to run was to have a culture medium that only contained radio-active phosphate. No ordinary phosphate. Now this wouldn't bother small molecules, but if you ever tried to build an RNA or a DNA, the radio-activity emitted would inevitably shatter it into pieces. Anything that could grow in that medium would automatically not be using DNA or RNA as its central function.

It may be that novel life forms are floating by us on this table. One way of doing it would be to set up a truly extraordinary culture medium, which excluded phosphate or some other ingredient thought essential for all of Earth life, which is really better described as life as we know it, and see if life as we don't know it just happens — be sending a spore that was floating back and fell in by accident and flourished.

The same concepts apply to the search for extra-terrestrial life. If you're only looking for Earth-like planets, if you're only looking for liquid water, if you're only limiting yourself to carbon compounds, you may be missing the organisms that exist. And when people want to send DNA sequences to Mars, or anti-bodies to Mars, I say, tell me another one. The two things eventually come together. If you believe that you need to generate an RNA molecule to get to start life, then the odds are so staggering that the odds are really that Earth is the only place is the universe, allowing perhaps that it spread by panspermia that has life.

On the other hand, if you believe that life could start with good molecules, given enough energy, then the universe may be rich with start-ups, and then there may be some series of levels that you have to go through, higher and higher, in order to get life more and more advanced. And we may be one of the relatively small number of places where you have intelligent life. I don't know, we haven't detected other intelligent life yet.

VENTER: Other places may not view us as intelligent life.

SHAPIRO: They have decided that one of the eight dimensions of string theory — one of the alternative universes that are now postulated by the anthropic people, are much more habitable than this. Life has a difficult job getting started. I admit this is one extreme view of life, but it's one that makes life, as Stuart Kauffman put it, something that the universe has in a sense expected, and what one does with that fact, I leave to you.

I'm not a theologian, I'm an agnostic, which says that I really do not know what's going on. But that at least in the origin of life we have a problem that can be solved not too difficultly in a laboratory, by getting the right set of molecules, by getting an appropriate source of energy — okay, we cheat a little bit, we use a beaker as the container rather than some membrane, which is perhaps more difficult to achieve than is commonly understood, and we just try to see what happens.

Does it all turn into tar, or are molecules always cut off from the energy as in the Miller experiment, or do you get ever-increasing and intricate cycles of reactions with some of the original compounds vanishing and others increasing in great numbers. If you are, then you may have caught a picture of the start of this universal phenomenon and you have to then find bigger and bigger vats or divide it up. Don't interfere, don't look for the results you expect, just let nature teach you, see what nature wants to do, given the bait of releasing this energy, which is what nature does seem to want to do.

I've spoken at length and enough so that I have enough voice to answer any question or accept any insults from the DNA-philes that are present at the table.

LLOYD: Richard Dawkins wrote a special email for you.

SHAPIRO: Richard Dawkins wrote a wonderful book, but the place where he absolutely blew it was in a section on the origin of life. He took all of the supposed irreducible complexity and said that natural selection can explain all this, you need no other rule. But then when it gets to how does one get to natural

selection, he has no other recourse — he's not a chemist — than to invoke some improbable event; he says that we need a vast improbable event and he goes anthropic and says, well maybe there are many universes and we happen to be the lucky one.

But the same type of reasoning that Richard Dawkins uses to explain evolution could apply equally well to what could be called thermo-dynamic evolution. In fact natural selection may just be one special case of thermo-dynamic evolution; there may be other forms of evolution undetected by us. So his schoolboy howler is the section on the origin of life. He writes brilliantly elsewhere. If you want to write that to him, Freeman — He hasn't written an email to me; I keep my hands off evolution, I don't claim to know very much about it.

VENTER: Maybe I come at this as a basic experimentalist — the theory behind theory is that you come up with truly testable ideas. Otherwise it's no different than faith. It might as well be a religion if there's no evidence for it. So how do you get it past your religion phase?

SHAPIRO: Instead of looking for elaborate new ways to make ribose or to connect nucleotides, if an equal amount of money were invested into telling people to just look at coupled reactions where energy is discharged in matter.

VENTER: If we're looking for life here other than DNA-based life, what should we be looking for that we're not?

SHAPIRO: Well, now we're talking about two different questions. One is how

you start with inanimate life, chaotic mixtures of chemicals and under the influence you organize them, which is an event that took place here three and a half billion years ago. How you look for a shadow biosphere, life that is very different than ours, well there are very many schemes and all of them are worth trying. Graham Cairns-Smith has numerous schemes for detecting life made entirely of minerals.

I said How would one detect advanced mineral life and he said, well there are always fresh starts. And he would invest money in looking for unusual minerals — following the activity of minerals that are out of place, or growing unusually in different areas, or having interesting interactions with organic compounds, unusual catalytic effects. I would have told him my idea would just be to set up a culture medium consisting of rich minerals — someone was talking jokingly about the ultimate diet where you're not a meat-eater, you're not a vegetarian, you're not even a vegan because even vegetables are alive, you're a mineral. You eat only minerals and breathe only carbon dioxides, and that's the ultimate in dietary purity.

LLOYD: Now you've destroyed this. If now there are mineral societies, you can't even eat just minerals any more.

SHAPIRO: I may leave us all to starve, but I would just take the roof of one of those buildings in Glasgow, which looks as industrial and polluted as any place on the planet, though maybe China would be better these days, and you put in a rich mineral bath and you have a filter that excluded anything somehow, any virus or living bacterium, and you just see if anything grew. Nothing more expensive than

that, you know?

VENTER: What's missing is your definition of life. It seems to encompass a broader range than maybe my definition would.

SHAPIRO: There was a wonderful paper written by Chris Chyba and Carol Cleland about three years ago about definitions of life, and how even defining what definition is can get you into philosophical doo-doo. And it's best to look for phenomena that by their properties we would be happy to classify as alive, and to not worry too much about definition.

Of course we'd want something that didn't extinguish immediately. That wouldn't be a good kind of life. One could consider a Zhabotinsky reaction as alive, or a thunderstorm, or a hail storm — but they don't evolve, they dissipate, so that isn't interesting life. What we're really interested in is interesting life — something which becomes more and more complex and adapts so it resists being extinguished.

VENTER: Does it need to be self-replicating?

SHAPIRO: It needs to be reproduced. The idea of a replicator, of DNA copying itself. I have a tie like that: it shows nucleotides swimming up to DNA, and miraculously one strand forms a double helix, but anyone who teaches biochemistry knows that doesn't happen — no way. There are dozens of proteins that come in and get involved in the action, and untwist the twists of DNA, and prime it and close the gaps in DNA.

VENTER: I wasn't describing a mechanism, just, the term 'self-replicating'.

SHAPIRO: DNA isn't self-replicating.

VENTER: No, I'm not talking about DNA.

SHAPIRO: And RNA as far as I know isn't — virus needs an entire cell filled with ribosomes and god knows what — mitochondria.

VENTER: Methanococcus is self-replicating.

SHAPIRO: Methanococcus is self-replicating, and if it lives and grows and changes eventually into different strains, that's alive.

LLOYD: So is a virus alive?

SHAPIRO: That's a question of how you want to define it.

VENTER: Is it not self-replicating.

LLOYD: I'm not self-replicating either. I have children and neither of them look anything like me.

SHAPIRO: The difficulties in these definitions are notorious. Is a nun alive? She's certainly not replicating. Is a mule alive? It has most of other properties, but it's

sterile and has no offspring.

CHURCH: Its cells are alive.

SHAPIRO: Its cells are alive.

VENTER: If we're looking for life, it helps to know what we're looking for in some form.

SHAPIRO: Yes, I would design missions to Mars to follow the carbon, not the water. They've detected methane now in the atmosphere, and I would have orbiters that sniff that methane and looked for the place where it was coming out of the ground and then analyze whatever organic chemicals might be emitted there. Out of the nature and identity of those organic chemicals, I would come to a conclusion about whether something of interest is present there or not, and decide if missions should be flown to investigate that site in greater depth.

VENTER: My other question is, I don't understand your dismissal of Stanley Miller's experiments.

SHAPIRO: I'm not dismissing — he was really addressing a very separate and valid scientific question, which is how did Earth come into its carbon reservoir. I mean we know there is carbon in various forms on the Earth.

VENTER: But isn't the spontaneous formation of the amino acids that he showed still valid? And if not, why not?

SHAPIRO: Well, if Earth had an atmosphere of methane, predominantly methane, which is a highly debatable subject, which the geologists are arguing about...

VENTER: There's a question how applicable it is to the Earth, but the question is, with that mixture did he not really get spontaneous formation of amino acids.

SHAPIRO: He got spontaneous formation of the two simplest amino acids in reasonable amount, plus very trace amounts of other amino acids, and some amino acids that are not present and made no use of in biology, and the organic compound he got in greatest amount in something called formic acid, which is secreted by some ants as part of their venom but plays no other role as such in Earth biology.

VENTER: So you're saying his experiments are valid, they're just not applicable.

SHAPIRO: They don't tell us anything about the origin of life, they tell us one way in which more complex carbon compounds, more complex than methane, were formed. Other alternatives are that they were brought in by heavy in fall of meteorites or comets, or else that they got trapped in the center of the Earth and were out gassed by volcanoes, and volcanoes to this day are out gassing carbon monoxide and methane phyol, that by itself could not start life alone. There's another possible source of the carbon. His was an authentic contribution to that question, it was not a contribution to the origin of life per se. It didn't answer what happened next, which is the crucial question.

LLOYD: It sounds consistent with the finding of simple amino acids in meteorites. He did a simple experiment in which he produced simple amino acids, and then we see these simple amino acids in meteorites, so it almost sounds like it's okay to assume that simple amino acids would be available as part of the chemical soup from which you could concoct life.

SHAPIRO: I would say it's okay to assume that simple amino acids were present in the early oceans, but at great dilution such that ten to the minus fifth molar is the best estimate that I've seen — or lower. And at that dilution very little of interest was likely to have happened. Places where for some reason carbon compounds were concentrated, due to whatever effect and by whatever mechanism, are perhaps more likely environments as the place where life began. And there may have been more than one solution, there may have been many starts of life, and we're simply sitting here as the children of the most successful one.

PRESS: Are you suggesting we should be doing a million Miller-Urey experiments, do combinatorial exobiology, look at many many different types of conditions and see where you get the kinds of simple molecules that you're talking about?

SHAPIRO: No, we can draw clues from existing life as to the type of molecules one should be looking at. Catalysis by minerals, which is catching on, was an important part, and certainly oxidation and reduction by minerals was another important part, so I would look at interfaces between minerals and aqueous media as a very promising place to look. One might look at places where primitive

organic compounds or even carbon monoxide itself was emitted from the Earth as one possibility, and deep-sea vents offer this reason very popular places.

Günter Wächtershäuser, who's only occasionally leaking what he's doing, is working on minerals and for a time he very heavily favored iron pyrites as a very likely mineral because it's oxidization and reducible, and because certain of our key enzymes contain clusters of iron and sulphur, which he regards as authentic relics of the oldest life.

If you think that phosphate was important from the very beginning, then those places on Earth which are very rich in phosphate, or better yet in pyrophosphate, minerals of pyrophosphate, if any existed, might be very likely places — those exposed to water. Again, if those minerals had any tendency to collect organic compounds in their interstices, that would be another site worth investigating. All of these sites are worth investigating.

In some places you might start a form of life that doesn't lead in the direction of the life we know. So much the better. The events that started out type of life may have been in an environment, which no longer exists on the face of this planet. And environments that we might investigate today, or ones we might concoct in a test tube, might lead to different kinds of life again.

But the study reactions where energy, say, present in minerals are discharged and there's a steady flow of simple organics in some sort of flow reactor, and one just sampled what the content of organic chemicals formed in such reactions was, without any prejudice, would be very interesting. One could take hints, if

Wächtershäuser likes, pyrites: very good, that could be the mineral of choice. If Steve Benner thinks that life started in borates, well, wonderful. Let him pour streams of simple formaldehyde through borates from now until doomsday, and if cycles begin I will bow my head in his direction.

VENTER: But maybe we're looking for an answer that didn't even exist, maybe it didn't originate of Earth at all. Given the statistical probability of some of the events you're talking about, panspermia has a wonderfully high probability if life exists almost everywhere else in the universe.

LLOYD: But it had to start somewhere.

VENTER: Yes, but we could be looking at totally the wrong environments.

SHAPIRO: But there are so many different environments, so much energy of all sorts being shot out of volcanoes, given off by lightning, pouring down in all sort of radiation — this is a very likely place, much better than, say, the moon.

SASSELOV: That's certainly a possibility, but we'll have an opportunity to actually experimentally look for that because there is not that much time in this young universe to spread the seeds so to say. So these environments must be very common.

VENTER: Yes, and they should be common.

SHAPIRO: I'm very in favor of environments that look for life on Mars. I

wouldn't start by looking for silicon life, because we don't know anything about silicon life, and we do know something about carbon life — so if there was carbon life but it was different than our own I think we have a fair chance of identifying it somehow.

VENTER: If there is silicon life, do they live in glass houses?

SHAPIRO: They may excrete glass houses — that could be their waste product.

LLOYD: Which brings me to a question about this hypothetical non-carbon-based life that might be here on Earth. When you have one of these systems, it has a metabolism — so it's taking in energy, and different chemicals, and excreting energy in other forms of a higher entropy content but presumably not at maximum entropy, the stuff that they're excreting. When that happens, you have a source of free energy and one thing that life is very good at is finding sources of free energy and then taking advantage of it. It seems strange to talk about another thing out there metabolizing that's different from ordinary life, and yet that's hard for us to detect. Surely something like that would actually be being taken advantage of by regular life — wouldn't we be able to see its signature there?

SHAPIRO: Perhaps, if regular life existed in that environment. That environment might be inimical to regular life. For example because it was utterly depleted in phosphate, to name one thing, or because the temperature was much too cold for regular life to exist, but life that worked on much weaker chemical bonds might flourish there.

LLOYD: What I'm suggesting, if the stuff is there and it's in an area where there are other living things around, then other living things are presumably eating this, or eating the byproducts of this.

VENTER: Something like methanococcus does live on just minerals and inorganics — it's a true autotroph. Autotrophs exist quite abundantly in nature. CO2, hydrogen — it doesn't take any organics presumably, or doesn't need to, for life. So you could say that that's evidence for mineral forms of life.

PRESS: I just want to clarify the definition that you're working with here — you're talking about looking for life in terms of a metabolism and energetics — you're not really dealing with the whole question of self-replication, is that right?

SHAPIRO: 'Replication' is a word that's gotten too tied to the idea of DNA, digital information storage. I'm talking about the topic that Freeman referred to earlier, which has the broader name of 'reproduction'. Somehow the entity, call it what you will, becomes two entities, or three, each of which has the functionality and the capabilities of the original entity.

DYSON: I'd like to raise another question, which is the question of a vacuum life — most of the universe if a good vacuum, most of the habitat, probably 99.9 percent of the real estate, is small objects — asteroids and comets, dust, grains, all kinds of things. We are the exceptions; we happen to live in an atmosphere. We have been biased to think that you should look on other planets.

But in point of fact I consider it vastly more likely if there is any external life that

it's not on planets at all. And of course once life is in a vacuum it has the enormous advantage — it's much easier then for it to spread from one place to another. Life that's adapted to an atmosphere first of all has a very hard time to get away from a planet, and secondly it's very unlikely to hit another planet with an atmosphere. So the probability that it will spread is very small.

SASSELOV: Freeman, do you suggest that this is a possible pathway to life? Or only life that's adapted to live in vacuum?

DYSON: I say both. It could be. I'm assuming it originated in a vacuum, because that's where most of the habitats are.

LLOYD: Yes: you need a source of free energy, and you need flows of materials. Well the surface of a comet has sources of free energy, and just in the same way that you were saying a good place to look for prebiotic life, as it were, on Earth, would be at a mineral-aqueous interface, the interface between the surface of a comet of a vacuum has the same features that you might require from this mineral-aqueous interface.

DYSON: Right — not that we shouldn't look for planets, but that it's stupid to concentrate one's attention on planets to the extent that we are.

DIMITAR SASSELOV

Is Earth the ideal planet for life? What is the future of life in our universe? We often imagine our place in the universe in the same way we experience our lives and the places we inhabit. We imagine a practically static eternal universe where we, and life in general, are born, grow up, and mature; we are merely one of numerous generations.

This is so untrue! We now know that the universe is 14 and Earth life is 4 billion years old: life and the universe are almost peers. If the universe were a 55-year old, life would be a 16-year old teenager. The universe is nowhere close to being static and unchanging either.

Together with this realization of our changing universe, we are

now facing a second, seemingly unrelated realization: there is a new kind of planet out there which have been named super-Earths, that can provide to life all that our little Earth does. And more.

Dimitar Sasselov is Professor of Astronomy at Harvard University and Director, Harvard Origins of Life Initiative. Most recently his research has led him to explore the nature of planets orbiting other stars. Using novel techniques, he has discovered a few such planets, and his hope is to use these techniques to find planets like Earth. He is the founder and director of the new Harvard Origins of Life Initiative, a multidisciplinary center bridging scientists in the physical and in the life sciences, intent to study the transition from chemistry to life and its place in the context of the Universe.

DIMITAR SASSELOV: I will start the same way, by introducing my background. I am a physicist, just like Freeman and Seth, in background, but my expertise is astrophysics, and more particularly planetary astrophysics. So that means I'm here to try to tell you a little bit of what's new in the big picture, and also to warn you that my background basically means that I'm looking for general relationships — for generalities rather than specific answers to the questions that we are discussing here today.

So, for example, I am personally more interested in the question of the origins of life, rather than the origin of life. What I mean by that is I'm trying to understand what we could learn about pathways to life, or pathways to the complex chemistry that we recognize as life. As opposed to narrowly answering the question of what

is the origin of life on this planet. And that's not to say there is more value in one or the other; it's just the approach that somebody with my background would naturally try to take. And also the approach, which — I would agree to some extent with what was said already — is in need of more research and has some promise.

One of the reasons why I think there are a lot of interesting new things coming from that perspective, that is from the cosmic perspective, or planetary perspective, is because we have a lot more evidence for what is out there in the universe than we did even a few years ago. So to some extent, what I want to tell you here is some of this new evidence and why is it so exciting, in being able to actually inform what we are discussing here.

Basically, in order to explain to you why this is interesting, I want to first of all convince you about three things, which are important to my approach. The first one is that what we are looking for is baryonic in nature. What I mean by that is something of which I don't need to convince you, I believe, but you should bear it in mind because this is a feature of our universe, the one we observe. Baryons are all the particles that make up atoms and all that is around us, including ourselves. But that's not necessarily the most common entity in the universe, as you — I'm sure — know about dark matter and dark energy. I think we have to agree that what we are looking for and would call life is baryonic in nature, and there is good reason to believe that dark matter and dark energy are not capable of that level of complexity in this universe yet — or at all.

The second point which I want to convince you of — or use as my background for

what I'll tell you here — is that we should agree that what we are looking for, what we call life, is a complex chemical process. Basically, the ability of those atoms to combine in non-trivial ways. This is actually my point of departure, where I would be looking at life more from the purely thermodynamic aspect, that is from the point of view which Robert here described and H. Morowitz has been very eloquent in defining and actually done some research on. That is, what is the parameter space in which you can have chemistry which is complex enough to lead to a qualitatively new phenomenon, a phenomenon which we don't see in the rest of the universe. That's actually an important point here.

Do we know enough about the universe that we can have sufficiently good feeling about that parameter space? Obviously we don't have detailed knowledge of most of the observable universe, but the last 50 years have been actually a revolution in that field, in the sense of the ability to get diagnostics of very distant objects and a very large number of objects.

The databases in astronomy up until just a few years ago were larger than what biology had. It's only now that biology — and, I guess, telecommunications companies — have exceeded that. But one aspect of these databases is that you very rarely see unusual, unexplained phenomena. Despite what you all would like to write on the front page of your newspapers and magazines, actually there is a lot of very boring amount of data there, which is hundreds of thousands — already millions — of stars, which have exactly the same isotopic and chemical patterns that are predicted by the theory which is well developed and is called 'stellar evolution' (although it has very little to do with evolution as used in biology).

But it is one of those steps that we now understand as the development of our world that is of our universe, of starting with very simple baryonic structure for that matter, which then becomes more and more complex. Stellar evolution is one of those phenomena that did not exist in the first half billion years of the universe. And this is not a hypothesis; we know it. We actually can observe a lot of it, and we know that there were no stars during the epoch of recombination, which is the cosmic microwave background, with all the structure that we see in it. And then there were stars, and then stars started a new process, which did not exist in the universe before, which is the synthesis of the heavy elements. That is — baryons working together as elementary particles and building a structure — the Mendeleev table, which then would lead to chemistry.

VENTER: How many years ago was this?

SASSELOV: 13.7 billion years ago is where we see the precursor of the microwave background radiation, so that's our first very well studied piece of evidence. Then about half a billion later is the time when the first stars can form, from the gas, and they're mostly made of hydrogen and helium. Then they go through a period where over a time of five billion years they produce enough carbon, nitrogen and oxygen and all the heavy elements, where you start effectively producing planets. And then we come to 4.5 billion years, which is the origin of our own solar system and the Earth. And almost within a half billion years, some complex chemistry which we now see covering entirely and co-opting the geophysical cycles of this planet. So that's to give you a quick idea about the time scales.

In that sense life is an integral part of that global development that we see. And although we know only one example of it, it doesn't seem unusual when you think of it that way — as a progression of complexity that the baryonic aspect of this — baryonic matter — in this universe has actually the propensity to lead to. So the question then is what is this good for understanding the origins of life, or possible pathways? And even more generically, could we design experiments in which we can find out whether all these possible baryonic pathways really merge into one — the one that produces life here on Earth — or are there multiple pathways? Even if you could answer that question, that would be very exciting, because it will tell us something about the general rules of complexity that baryonic chemistry can really lead to.

The question then is, the third aspect which I want to convince you of, is we know quite a bit about the universe, but there are only a few places in the universe where you can think of that complex chemistry being capable to survive over a sufficiently long period of time. And vacuum is not one of them, in the sense of surviving in which you were talking about the origin of life; starting with smaller molecules, which then have enough time to lead to more complex ones. And when I think of vacuum, I don't mean the surface of a comet, but really the inter-stellar medium, with its very low density.

I can imagine life that started on some surface then migrating to live in the inter-stellar medium. But I cannot imagine, as an astrophysicist, from what I know, that there is an environment, which is stable enough over the time scales necessary for that chemistry to take place. So I am a little bit biased in that sense to planets and

planetary systems as the only environment that we know of today, as far as we know in the universe, which has all of those factors put together — that is, stability over long periods of time, but sufficiently low or moderate temperatures. (Stars are very stable over billions of years, but they all have very high temperatures, all throughout.) And basically the overall thermodynamic window that Morowitz is talking about, which allows complex chemistry. That's actually much broader than simply having water.

When people talk about habitable environments, sometimes they would equate that to the existence of water, or the ability of water to be in a liquid form. That's a much broader view of what is available there. But whatever your idea of what could be habitable is, the bottom line is that there are not that many objects, or places, in the observable universe that allow that. In fact, planetary systems are certainly not only the best, but are probably the only ones on which we are certain that complex chemistry can occur.

Then the question is, how much do we know about planetary systems? Up until 12 years ago, essentially we knew only of one: the solar system. That situation is very similar to what we have with life. We only have one example. And that's bad from many points of view, and we — 'we' meaning astronomers — learned it the hard way, because it turned out that what we had theorized about planets was very solar system-centric, and we missed a lot of things that we should not have missed, but that always happens when you have only one example of something.

What planets allow you to do now that we know how many different types of them there are, is you can have a pretty good estimate of what to look for. And

one of the things that we learned - I guess the hard way - is that we do not necessarily have to look for planets just like the Earth. What I mean by this is that although in our solar system we have a very large variety of planets — you have Jupiter, which is very much bigger than the Earth, ten times in size, 300 times in mass; you have Saturn; you have Neptune and Uranus — all giant planets, all made of gas — then you have very small planets: that's the Earth, Venus, and Mars, and Mercury, going smaller — and then comets and asteroids.

There is a very significant gap in masses between 1 Earth mass and 14, where Uranus and Neptune are. That's actually, as we would say in physics, more than an order of magnitude. And it allows for a whole set of phenomena that could happen in that range that we've been missing. And from what we understand now, both from theory and more recently — meaning in the last two years — from observations of such systems, is that the fact that the solar system has no planet like this is just a fluke. It just happened the way the planets were formed that what ended up being the solar system has no planet which is in that mass range. The majority of planets in that mass range will be like the Earth, and for lack of a better term, we ended up calling them Super Earths.

I get a lot of flack for introducing that, but it comes from my bias as an astronomer. We call stars that are bigger than giants, super-giants; we call stellar explosions which are more energetic than novae, super-novae; so it just made sense that if you have a planet which is larger than the Earth but otherwise is in essence similar to the Earth, you would call it super Earth. I guess I didn't grow up with Super Man.

CHURCH: That's not Super-Earth, that's Krypton.

SASSELOV: Just take it as it is — it's just a term — it's just planets which are larger than the Earth. Now why is that interesting — if you really limit yourself to planets larger than Venus and Earth, but not much larger than Earth, then you're left with very small numbers in the galaxy as a whole and in our part of the galaxy as a whole. If you allow yourself to count super-Earths as part of the inventory that you can tap, then your numbers grow by two orders of magnitude. I'm saying this is because of two lines of evidence.

LLOYD: What is the concentration of the smaller ones? What fraction of solar systems, or stellar systems, has 'sub-Earth' planets?

SASSELOV: Ah, so that's actually a difficult question — what fraction of the planetary systems have planets smaller than the Earth — because they're hard to see. We have some estimates, which go to about the fraction of an Earth mass; well let's just say one Earth mass. We have no technical evidence now for less than that. That's from a technique that is called micro-lensing, by the way.

The evidence for this is in part statistical, but that's quite often the case. You observe many objects and you build statistical cases for all of that. On the one hand we already have detected a number of super-Earths — the current number is actually five. That's a small number for statistics, but it is not a small number statistics when you view it as an effort where a lot of other planets have been detected, and despite the difficulty of detecting smaller and smaller planets, you are detecting an increasing number of those in the planetary systems that you are

observing. In other words, as you go to smaller and smaller masses, below about 12 to15 Earth masses to a planet, the numbers actually rise despite the statistical biases of actually having less of those. This anticipates that as our technology improves, which by the way it is, on a monthly basis, we will be discovering more of those.

There is another line of evidence which is a technique which is called micro-lensing for detection of planets, that is sensitive to the entire mass range of planets, all the way down to one Earth mass, and actually in fact a bit smaller than one Earth mass. This technique is scanning without any prejudice a large number of stars and to this point they have actually detected more super Earths — smaller planets — than larger planets. Which then tells you that if you take the current statistical numbers, which we have already figured out pretty well because we have larger planets in large numbers from the last 12 years of study, you can actually estimate what is the expected number of smaller planets just because of this comparison that you do.

There is a third line of evidence, which being a theorist myself I would not really push too hard, but theoretically if you form large planets you also form small planets, and there is no particular theoretical prejudice that anybody has come up with at this point, that you will somehow create gaps like the one we have in the solar system, where you will have only very small planets and only very big planets.

So the final question here is, are these super Earths actually any good for what we're interested in?

VENTER: Can you actually put a number — what's the number in the universe of super Earths?

SASSELOV: Well, that's a good question. Let's take our galaxy as an example, not the whole universe. We now have a pretty good idea that there are about 10^11 — a few times, 2 or 3 times 10^11 stars in the galaxy. So then we know that of those stars, only about 90 percent live long enough for the kind of complex chemistry that we have in mind, which is half a billion years or longer. However, only about 1/10 of these stars have enough heavy elements so the planets that will form around a star like that will either not form at all, or will have a significant deficiency. In fact we have evidence for that. Then the question is, how much do we know about the number of super Earths? Basically of those left over, where we have ten billion or so, you would say that it's only a fraction which is less than 50 percent but larger than 10 percent from those arguments that I gave to you so far. And then you look where in the planetary system you are — you don't want to be right next to the star and you don't want to be too far from the star, and this is following Morowitz's thermodynamic estimates for the temperature range. The bottom line that you end up with is about a hundred million planets that I would call habitable in the sense that they allow this kind of complex chemistry somewhere near their surface. A hundred million in our galaxy.

VENTER: And how many galaxies are there now?

SASSELOV: Oh, that's a large number, but it's a similar number to the number of stars — 10^11.

The question is — I actually insist on doing it for the galaxy, because I'm interested in the experiment; I'm a theorist, but I really trust the experiment — how many of those environments can we study soon enough (while I am still alive) and with enough detail that we actually can help you guys, the chemists and the molecular biologists, to constrain your experiments into those pathways to life. Basically the estimate is many. Because if you have that many, a hundred million in our galaxy, then only in our vicinity, with the experiments which are already underway, we'll have at least about fifty to a hundred in the next five years. And fifty to a hundred for which we can get some data that will be interesting to inform those questions.

VENTER: So your data set would exclude things like Europa?

SASSELOV: No, not at all — Europa is a great place to look for life. I'm just saying this is the minimum.

VENTER: But I mean size-wise.

SASSELOV: Well, the reason that Europa is viable is because of Jupiter. If Europa was just by itself we may not consider it that viable. In a sense I'm trying to be conservative here, and I can tell you that I can promise you only that many. But there is another reason why I actually would like to make this estimate, and why I talk about the hundred or so that we are going to be able to study. And this is because I do want to be able to study them outside of our solar system. And the question is, how do you study Europa in a planetary system that is 50 light years

away? Very difficult.

But can you study a planet which is five times more massive than the Earth and two times larger than the Earth? Yes. Even much more easily than an Earth-size planet. So the point that I'm making is that the fact that super Earths are viable as planets in the comparison to the Earth is actually great for our ability to do these experiments, because it's much easier to detect and study a planet which is two times bigger than the Earth and is still viable. You can learn a lot from it.

One of the reasons I call these planets viable, and in fact even more viable than the Earth, is because they have the basic characteristics of the Earth, except in a much more robust way. You probably know that there is a big problem in planetary science, which is the comparison between the Earth and Venus. Why does the Earth have an atmosphere which is not very hot, that's sort of understood — not yet, but sort of. Why does the Earth have plate tectonics, while Venus doesn't have plate tectonics, that's not understood — or we are at the verge of starting to understand that. These are questions that are much easier to answer for super Earths.

It turns out that plate tectonics, as understood from Earth, is a process which has been going on theoretically much more easily on a slightly bigger planet. In fact if you do the theory, as best as you can today, the Earth is at the margin of what is viable in terms of plate tectonics. Probably some of you may know that plate tectonics is a very important aspect of the viability of a planet in terms of surface conditions, because it's a good thermostat, it keeps the climate more or less stable over long periods of time, and also allows you to have easy access to the large

reservoir of chemicals and gasses in the mantle of the planet.

In that sense super Earths are as good as the Earth, and I would argue — better. They have more stable and robust surface conditions. So there are more of them, they're as good as the Earth, if not better, and they are easier to study. So we have a very bright future of being able to at least find out what's going on.

VENTER: What role does gravity play in the larger — in the super Earths?

SASSELOV: It's actually a positive role. In the sense that if you take the general amount of out-gassing, fluxes, which interchange between the mantle and the atmosphere of the Earth, the Earth's gravity is very close to marginal — we know Mars is an example where it's definitely sub-marginal, in retaining a sufficient atmosphere, and hence making this thermostat being viable, and really providing you with stable conditions over at least a billion years. So having more gravity is actually better.

VENTER: It increases the odds of having an atmosphere?

SASSELOV: In keeping it. You always have an atmosphere — even Mercury has an atmosphere: there is some helium that is being punched out of the surface of the planet, but it simply cannot retain any of it. It just goes away.

So I'd prefer to answer questions rather than to continue.

SHAPIRO: Which is the closest known super Earth?

SASSELOV: The closest known is called — in fact there are two of them: Gliese 581c and d, and both of them are super Earths, and are just 20 light years away. Wilhelm Gliese was a German astronomer (1915-1993).

CHURCH: When will they arrive here?

SASSELOV: Next week.

CHURCH: Since they're better than us.

SASSELOV: The names are Gliese 581c and d — that's the number of the stars. c and d stands for 'planet c' and 'planet d'. There is also ‘planet b’, which is a bigger, Neptune-like planet. 30 years ago, Gliese made a catalog of all the nearby stars. A lot of them are very faint, they hence were only identified in this catalog, so it's a common practice to call the stars by the name of the author of the catalog with a consecutive number.

PRESS: Can you clarify the ratio that you're seeing from the microlensing studies of Earths to super-Earths? I didn't quite catch that number.

SASSELOV: Ah, it's not Earths to super-Earths, it's Earth-like planets, which is super Earths including Earth mass, to giant planets. And we have a good number for the statistics of giant planets in planetary systems right now from all the ones which were discovered with the Doppler shift technique over the past 12 years. They have 250 of them, so that already gives you a good statistic.

DYSON: How many micro-lensing planets do you have?

SASSELOV: At this point, seven.

Someone had asked if it is possible to launch a spore (panspermia) from here to there and hit something 20 light years away? Is it possible to hit it? The answer is yes, from the physics point of view, it's possible. And I have a colleague who says that if anything is possible from the laws of physics, it happens in the universe. But he's a physicist. Let me qualify that. I think panspermia was very popular, and in a sense originated in the modern sense in the 20th century, when there was a possibility, or even some knowledge, that the universe may be very old. Older than 20 billion years, or maybe eternal.

Fred Hoyle, one of the people who really supported and developed that — sort of a steady-state universe, which is not 13 billion years old. Now this is very significant, because then it IS much more likely that complex chemistry originated somewhere in the universe and then spread — it's a very robust system as we know it on this planet; it can really spread. And there is plenty of time for it to spread — even over the large distances and all the vagaries of high-energy astrophysics. However, the universe is only 13.7 billion years old, and actually you have to subtract 0.7 just in order to have the first generation of stars.

The first generation of stars, they were made of hydrogen and helium — there was no carbon, no oxygen, no metals — were actually very large, and they could not have even sustained protoplanetary disks let alone have planets form. It took a

long time before you could actually start forming small stars, in which the protoplanetary disks would have enough solid particles to coagulate to form planets. And that actually was thought of theoretically but with very large uncertainties, and now we seem to get, somewhat unexpectedly, evidence for this. Very strong evidence, in fact.

There are a lot of searches for planets and planetary systems that are now even targeted towards stars which are slightly older and have less heavy elements than our own sun. You see a very precipitous drop in the detection of such planetary systems. In fact after certain factor, about ten down, in such heavy elements, nobody has been able to detect a single planet, which is kind of strange, especially because there are people who are trying very hard — one of my colleagues has been trying now for eight years and has come up with zero. And for the normal metallicity (that is stars like the sun), the numbers are very high, already 250.

That means that you have to come to that part in the history of the universe when you have stars like that form. Like the sun — or a little bit less rich in metals. Simply because it takes that much time for the previous generations of stars to synthesize through nuclear fusion those elements. And we know actually that number: another of the big successes in the last five years is that you can now go back all the way to that time and see — literally see, by measurements — the increase in heavy elements as you go from one generation of stars to the next and so on.

Basically you are left, I would argue, with about seven billion years. So the first complex chemistry which could have occurred on the surface of a planet would

have started about 7 billion years ago. Now in 7 billion years it's very difficult to bring something from there to here, especially if you want to bring it 4 billion years ago. You only have three.

CHURCH: How hard would it be to hit? At that arc angle, and the radiation damage, and all the rest.

SASSELOV: We don't know that. It's very easy to calculate the cross-section for it hitting. But understanding how it's going to make it there is very difficult.

DYSON: There is some evidence. A radio-astronomer called Jack Baggaley in New Zealand — I don't know if you know him or not — is observing dust grains coming into the upper atmosphere and he claims that first of all a lot of them are extra-solar, they actually come from beyond the solar system, and furthermore they are preponderantly in a certain direction in the sky which comes from a star called Beta Pictoris, which is not very far away, 60 light years or something like that — and it is actually quite a young star but it has a very large dust cloud around it. And so it's plausible that dust is coming from this star and actually hitting the Earth. And if dust grains are coming, there's no reason at all why bigger objects shouldn't also be coming, and they would probably be following similar trajectories. So in principle we know that stuff is arriving from other solar systems. Whether anything is alive on Beta Pic is another question.

SASSELOV: By the way, the stuff which is coming from Beta Pic only started coming to the Earth recently. In the last hundred million years or so. Because Beta Pic didn't exist before that, and it took that long for it to come.

DYSON: Right, but it's easy to get here within the time available.

VENTER: Especially: we had these large asteroid hits hitting the Earth after microbial life existed here — we splashed a hell of a lot of stuff into space from those hits.

SASSELOV: So from here on, Beta Pic's hail will go to some other place.

VENTER: That's why I was trying to push you for a number in the universe, you know, ten to the eighth probability in our own system is pretty high probability. There could have been a million origins, all contributing to panspermic events, you know?

SASSELOV: Right, and then the question is, can you cross from one galaxy to another? And that takes a long time. It takes billions of years. And that's the problem. We can't go from one galaxy to another.

DYSON: Isn't our galaxy big enough?

SASSELOV: That what I was saying, that's why I was making the estimate just for our galaxy.

DYSON: For now.

SASSELOV: Yes.

VENTER: So at comet rates how long would it take to go the 20 light years? Comet speeds.

SASSELOV: It's about a million years. And I'm talking about the really fast comets. The fastest that we've observed.

CHURCH: That would be a lot of radiation damage, in a million years, I would guess.

SASSELOV: Sure. The problem is that we don't yet have evidence? I'm talking about these fast comets — we don't have the evidence for a comet coming from outside the solar system. All of them have close to parabolic orbits — a hyperbolic orbit means that the comet came from outside the solar system and left — so if you take those speeds it's still millions of years.

VENTER: How long would it take a .1 micron object to collect three million rads of radiation, traveling through space.

SASSELOV: You said a ten micron object?

VENTER: No, point one.

SASSELOV: Oh, point one, yes. I don't know.

VENTER: We can find organisms right here that could take 3 million rads of

radiation.

LLOYD: And I bet it wouldn't be that hard to calculate what the velocity distribution of such organic objects would be after an asteroid hit, and roughly how many there would be going out?

SASSELOV: I would prefer a larger object in which your prize collection is embedded. Because then you shield it. So it's shielded from cosmic rays, you shield it from any kind of radiation.

CHURCH: It has to be really large, you know, like meters.

SASSELOV: No, I think if you're just worried about the million years and a certain dose OF rads, it shouldn't be larger than a few centimeters.

VENTER: Well up to recently we've been dumping all the feces from the space station into space, so that's kind of shield, anyway.

SASSELOV: I have a colleague at MIT who calls this the garbage belt of the Earth.

SHAPIRO: I have a different question. Couldn't we get some estimate of the probability of material leaving the Earth, splashing off as they say, by examining the surface of the moon in protected areas? Say those areas that are in permanent SHADOW craters.

VENTER: A couple of hundred kilograms a year.

SASSELOV: Actually several people suggested that as one of the interesting experiments to do when going back to the moon — to look for those. And we know which part of the moon would have most of it — it's not an even distribution over the surface, because of orbital dynamics. So that will be a very interesting thing to do.

SHAPIRO: Look in craters which are protected from the radiation, and so on.

SASSELOV: Yes. I think that's an excellent experiment to do. And if somebody else pays for it...

PRESS: And how would we be able to study the properties of these planets? Is there any way to do it other than looking at transiting planets?

SASSELOV: We've thought a lot about that, because this is where we are spending our money right now. And a lot of people are. There is in fact an Exoplanet task force, which is tasked to think this question over, and this is part of the direction which they are writing into their report right now. The idea is that we want to have two parallel paths.

One is the now old-fashioned one, just a few years old, called Terrestrial Planet Finder (TPF) — direct imaging, which is still viable, but probably will take longer technologically. And by direct imaging, what is meant is you're not imaging the surface of the planet directly, but you are imaging the planet

separately from the star and you are able to get spectroscopic information that way- as well as some surface information (if the planet spins then you see variations which can be interpreted as surface information).

In the meantime, though, technologically it's much more viable to look for transiting planets, and to study transiting planets. Because what transiting planets allow you to do is not only to discover the planets, which is not that important, but once you've discovered them, you have actually the ability to measure their mass and the radius very precisely. And by very precisely I mean to an accuracy of one, two, three percent, which is very precise. So that gives you a mean density of the planet.

It turns out that this mean density can tell you whether the planet is really a small Neptune-like planet, that is hiding as a super Earth but is really a very gas-rich planet without a solid surface anywhere. Or it is an Earth-like Super Earth, which is simply a version of Earth, just bigger. Then once you've passed that measurement, the next thing you can do is you can use measurements both during the time when the planet is in front of the star, which is called a transit, as well as when the planet is behind the star, which is called an eclipse.

In the first case, you measure gasses in the atmosphere through transmission, which is like passing through the atmosphere of the planet. In the second case you actually measure surface features. And the surface features give you a map of the surface: a color map if you do it in the infrared, and an albedo map if you do it in the optical. And this is actually the first such map, that I'm sure many of you have seen, which colleagues in our group at Harvard actually accomplished just a few

months ago. It was published three months ago. Now, this is a thermal map, of a giant planet, just like Jupiter. We're not talking about super-Earth here. But right now technically it is possible to do this for a super-Earth.

VENTER: What distance is this?

SASSELOV: This one is at a distance of about 45 light years. So if that was a super Earth you could do the same thing with the existing Spitzer telescope. So that's actually where we are putting our money right now and we hope that NASA will put money into the other one, which is TPF (Terrestrial Planet Finder).

PRESS: What is the prospect for actually being able to measure the atmosphere of the super Earth and say, hey, this things looks like it has an atmosphere out of chemical equilibrium, there's oxygen, — there's something there that makes you sit up and say this thing looks like a place that has life. What is the prospect for doing that?

SASSELOV: Five to ten years, where five is more likely at this point, the way things are developing. If we are lucky it can happen even in a year or two. But we have to be lucky. The projects which are going to discover large numbers of them are coming up. Corot is one of them, but NASA's Kepler is much more so, and there are a couple that have just started, or are being built, that will produce enough of those planets that you can then cherry pick and say, ah ha, now I have a few that I can study in detail. But in ten years we will have a whole gallery of them, as opposed to just a milestone.

DYSON: Which molecules will you be looking for?

SASSELOV: Anything that we can see. So basically the idea there is to have enough signal to noise that we can see them all. The resolution is not an issue, because most molecules have broad spectral features, so it's a matter of signal to noise. And we will try to see what we see, and personally that's actually one of the reasons why I'm involved in origins of life research. Origins of life, because I felt that I would be very embarrassed to have this gallery of spectra and maps of super Earths, without being able to answer the question, what do you think is going on on this planet — is it chemistry, or could it be biology?

DYSON: Can you see oxygen and nitrogen?

SASSELOV: Oh, oxygen and nitrogen are actually easier to see. Partly through their proxies which are CO2 and CM. The molecules.

PRESS: If you were looking at Earth from another super Earth somewhere else, is there anything that we've done to the environment that we could actually detect? A large increase in carbon dioxide...

SASSELOV: Yes, people have done this research already — partly in preparation for the Terrestrial Planet Finder. So the strongest indicators there are the existence of ozone — free oxygen and simultaneously amounts of methane — the imbalance is what leads you to believe that something unusual is going on and cannot be reproduced by any of the global planetary cycles that we sort of understand. It is easier to complete the parameter space of global planetary cycles,

like the carbon cycle, the sulfur cycle, and to say we are outside of any of that parameter space, that is you cannot explain that combination of atmospheric gases with any of those cycles operating. So by exclusion, you will see that there is something unusual here. But from that point of view, my estimate of habitable planets, a hundred million in our galaxy, excludes the Earth. The Earth as it is now is not very habitable. It's a very hostile environment for complex chemistry.

LLOYD: Were you hoping that people on this other planet would detect that we were screwing this one up, and would come and rescue us. Is that what you were hoping?

VENTER: Why would they want us?

SASSELOV: I actually mean it — I really mean the large numbers of free oxygen.

DYSON: But if you were looking at the Earth this way, would there be enough CN to detect?

SASSELOV: No. The Earth is actually quite a difficult case in that sense.

VENTER: So if it wasn't for the influence of religion, wouldn't we just logically assume that the extrapolation from life here to the statistical base, you know, that we will find it everywhere.

SASSELOV: Yeah, I would say microbial life — that is, the complex chemistry of that sort — is very likely, and the more important thing is that we'll have some

evidence to say something intelligent about it, rather than just saying it's very likely.

DYSON: Yeah, it could get stuck in any of these phases — I think the phase where you have to invent ribosomes is probably the one you're most likely to get stuck at.

LLOYD: Though it seemed to have happened relatively rapidly on Earth.

SASSELOV: My big question to all of you here is, can we do it by exclusion? Can we develop again this parameter phase of chemistry to such a completeness where I can look at these 50 planets eight years from now, and say, well, I know why all of these have what they have on their surface and atmosphere, but this one has really none of that - it's out of equilibrium, and it cannot be explained simply by physics and chemistry — it must be something which is more complex and is potentially life.

VENTER: Ken Nelson, who is head of the Mars Sample Return had to think of some of these issues a lot and he said the number one thing to look for is the phosphate bond. That's the single greatest signal for biological life as we know it.

DYSON: In looking for life as we know it, or perhaps life is as we don't know it.

VENTER: Might as well start with what we know.

CHURCH: How easy would it be to detect the phosphate bond?

SASSELOV: That would be very difficult. I was thinking the other way around — we understand physics quite a bit, chemistry I hope enough, and so if we say we understand chemistry and physics, and this is neither physics nor chemistry that we see there, we've got biology.

SHAPIRO: The trouble is we're looking for a separate origin, this one has the great philosophical impact — if we discovered that life was on Mars but just a spillover from Earth it would be a curiosity, but it would not turn our view of the universe on its head. On the other hand if we discovered a life that's different enough that it couldn't have originated here, the spread would really validate what he's been saying - that I've been saying — that life is inherent in the universe.

VENTER: Well the two aren't incompatible, it could be identical to what we have here everywhere, the same chemistry, and we find it everywhere.

CHURCH: But he's just saying it's hard to prove that.

SHAPIRO: Hard to prove — a muddled case.

CHURCH: Well, it's not a theoretical argument; you either find it or you don't.

LLOYD: Well that's what's so upsetting about this work. Of course saying 'oh look, there's something we don't understand; must be life' is perhaps not the most

compelling argument in the world. But if there is something weird going on, and it isn't explained by any of the models of life that we have already. Then that would be very interesting.

SETH LLOYD

> *If you program a computer at random, it will start producing other computers, other ways of computing, other more complicated, composite ways of computing. And here is where life shows up. Because the universe is already computing from the very beginning when it starts, starting from the Big Bang, as soon as elementary particles show up. Then it starts exploring — I'm sorry to have to use anthropomorphic language about this, I'm not imputing any kind of actual intent to the universe as a whole, but I have to use it for this to describe it — it starts to explore other ways of computing.*

Seth Lloyd is Professor of Mechanical Engineering at MIT and Director of the W.M. Keck Center for Extreme Quantum Information Theory (xQIT). He works on problems having to do with information and complex systems from the very small — how do atoms process information, how can you make them compute, to

the very large — how does society process information? And how can we understand society in terms of its ability to process information? He is the author of *Programming the Universe: A Quantum Computer Scientist Takes On the Cosmos*.

SETH LLOYD: I'd like to step back from talking about life itself. Instead I'd like to talk about what information processing in the universe can tell us about things like life. There's something rather mysterious about the universe. Not just rather mysterious, extremely mysterious. At bottom, the laws of physics are very simple. You can write them down on the back of a T-shirt: I see them written on the backs of T-shirts at MIT all the time, even in size petite. IN addition to that, the initial state of the universe, from what we can tell from observation, was also extremely simple. It can be described by a very few bits of information.

So we have simple laws and simple initial conditions. Yet if you look around you right now you see a huge amount of complexity. I see a bunch of human beings, each of whom is at least as complex as I am. I see trees and plants, I see cars, and as a mechanical engineer, I have to pay attention to cars. The world is extremely complex.

If you look up at the heavens, the heavens are no longer very uniform. There are clusters of galaxies and galaxies and stars and all sorts of different kinds of planets and super-earths and sub-earths, and super-humans and sub-humans, no doubt. The question is, what in the heck happened? Who ordered that? Where did

this come from? Why is the universe complex? Because normally you would think, okay, I start off with very simple initial conditions and very simple laws, and then I should get something that's simple. In fact, mathematical definitions of complexity like algorithmic information say, simple laws, simple initial conditions, imply the state is always simple. It's kind of bizarre. So what is it about the universe that makes it complex, that makes it spontaneously generate complexity? I'm not going to talk about super-natural explanations. What are natural explanations — scientific explanations of our universe and why it generates complexity, including complex things like life?

I claim that there is a very basic feature of the universe, which makes it natural for it to generate complex systems and complex behaviors. We shouldn't be surprised by this. It's intrinsic in the laws of physics. This is what Craig Venter was asking, what is it about the laws of physics that give us things like life? Not only that, we know what this feature is. Let me tell you what it is, and then I'll tell you what it has to do with life. Because the spontaneous generation of complexity is important for lots of things other than life. Remember, life is overrated. There's plenty of other interesting stuff going on in the universe other than life. Long after we're all dead, and maybe other biological forms — carbon-based forms — of life are dead, I hope that other interesting things will still be going on.

Okay. What is this feature that is responsible for generating complexity? I would say that it is the universe's intrinsic ability to register and process information at its most microscopic levels. When we build quantum computers, it's one electron: one bit, to paraphrase the Supreme Court. Because of quantum mechanics, the world is intrinsically digital. That's what the 'quantum' in quantum mechanics

means: it says the world comes in chunks. It's discrete. And this discreteness implies that elementary particles register bits. Their state can be described by a certain number of bits. In the case of the electron spin, one bit. In the case of photon polarization, one bit of information. Bits are intrinsic to the way the universe is. It's digital. And this digitality at the level of elementary particles gives rise to a very digital nature for chemistry, because chemistry arises out of quantum mechanics together with the masses of the elementary particles and the coupling constants of nature and the electro-magnetic force, et cetera.

Quantum mechanics means that there are only a discrete number of species of chemicals. You can only put together two hydrogens and an oxygen to make a molecule in one way that I know of. This means that we can catalog chemicals in a discrete list — chemical number one, chemical number two, chemical number three — you can order it any way you want according to your favorite chemicals. But it's discrete. This digital nature of the universe actually infects everything, in particular life. It's been known since the structure of DNA was elucidated that DNA is very digital. There are four possible base pairs per site, two bits per site, three and a half billion sites, seven billion bits of information in the human DNA. There's a very recognizable digital code of the kind that electrical engineers rediscovered in the 1950s that maps the codes for sequences of DNA onto expressions of proteins. There's a digital nature to the universe, and quantum mechanics makes this happen.

But the digital nature of the universe doesn't immediately tell you why the universe is complicated, and why something like life should spontaneously arise. The fact that we're here doesn't tell us anything about the probability that life

exists elsewhere in the universe. Because we're here, and so we have to be here in order to contemplate this question, this tells us nothing about the probability of life except that it can exist. That's why this kind of question that Dimitar is trying to answer by looking for planets and signatures of life elsewhere is so important. We really don't know how likely it is that life should arise.

So why does complex behavior arise? Well, the universe is computing at its most microscopic scales. Two electrons, two bits of information, every time they collide, those bits flip. It's just these natural interaction and information processing that we use when we build quantum computers. Now I claim — and I can claim this because this is a mathematical theorem, which is different from just mere observational evidence — that when you have something that is computing and you program it at random, just tossing IN little random bits of programming, that it necessarily generates complex behavior.

Einstein said, God doesn't play dice with the universe. Well, it's not true. Einstein famously was wrong about this. It was his schoolboy howler. He believed the universe was deterministic, but in fact it's not. Quantum mechanics is inherently probabilistic: that's just the way quantum mechanics works. Quantum mechanics is constantly injecting random bits of information into the universe. Now, if you take something that can compute, and you program it at random, then you find is that it will spontaneously start to generate all possible computable things. Why? Because you're generating all possible programs for the computer as you toss in information at random.

In fact the universe is computing. I know this, because we build quantum

computers — in addition, I can see a computer over there, so the universe clearly supports computation. And if you program it at random to start exploring different computations, if you go out into the infinite universe, (observational evidence suggests the universe is infinite), then somewhere out there every possible computation is being played out. Every possible way of processing information is occurring somewhere out there.

Okay? I don't think this is controversial, but in some funny way it seems to get people's dander up. The fact that the universe is at bottom computing, or is processing information, was actually established in the scientific sense, back in the late 19th century by Maxwell, Boltzmann, and Gibbs who showed that all atoms register bits of information. When they bounce off each other, these bits flip. That's actually where the first measures of information came up, because Maxwell, Boltzmann, and Gibbs were trying to define entropy, which is the real measure of information.

What happens when you have a computer being programmed at random? The computer generates all possible mathematical structures, and one of the most important things it does is to generate other computers amongst these structures. As first proposed by Alan Turing in the 1930s, a universal computer is a device that can simulate any other computer. It can be programmed to simulate any other computer, in a simple fashion. Including itself.

If you program a computer at random, it will start producing other computers, other ways of computing, other more complicated, composite ways of computing. And here is where life shows up. Because the universe is already computing from

the very beginning when it starts, starting from the Big Bang, as soon as elementary particles show up. Then it starts exploring — I'm sorry to have to use anthropomorphic language about this, I'm not imputing any kind of actual intent to the universe as a whole, but I have to use it for this to describe it — it starts to explore other ways of computing.

Now remember, chemicals are digital. There are only certain chemicals that can exist and the laws of chemistry are set catalogs of chemical reactions, potentially infinite in extent because the total number of possible chemicals can be extended as much as you want. You can make polymers longer and longer and longer — you can think of the laws of chemistry, which are actually in some sense simple being implied by quantum mechanics as being a catalog of this huge set of possible reactions, where if I produce chemical A, and chemical B, I put them together, then that produces chemical C in abundance. Or if chemical A and chemical B are there and chemical D is also there, then chemical C is not produced.

Now you can see the relationship of these kinds of reactions to logic, right — if A and B, then C — if A and B and D, then not C. I'm simplifying chemistry, of course, because there there are temporal dynamics as well. But those dynamics' if-then statements, the digital statements that lie at the bottom of computation, are an intrinsic part of chemistry.

The digital logic inherent in chemical reactions is extremely important in biology of course, because this is how the metabolism of a cell works. I receive this chemical and this other chemical; therefore I'm going to open this switch over

there and turn up on this other chemical pathway. Chemistry has this computational nature embedded in it, which it inherited from the underlying computation that's going on in quantum mechanics in general. Chemistry itself, then, explores out there in the universe all possible combinations that are out there in the universe. Chemistry explores all possible computations, all possible things that could happen — including all other things that a computer can do.

Let's produce this self-reproducing structure, and then see what happens. Or let's see what happens when we produce this structure and this other structure and they react with each other — let's see what they produce. We don't know exactly what went on in proto-life, but we do know the sorts of things that go on in proto-life, even without knowing the exact chemical reactions that took place. It is not surprising that chemistry should produce more and more complicated structures, which then interact in more and more complicated ways, and go on to fill out more and more of the set of all possible chemical reactions, and then produce further computationally complicated structures, like, say, bacteria, or human beings, or computers.

Because there is an intrinsic capacity built into the laws of nature: this ability to process information in an open-ended fashion. And once things start doing that then they're very hard to stop. I call such things "complexors" — because they generate complexity automatically. From the mathematical or physical perspective, complexors are actually rather simple, because all they are is something that can compute, which is systematically exploring a wide variety of, or all, possible computations. Once you have such a thing, once such a thing gets popped into existence, set into motion, then it will produce complexity, whether

you want it to or not.

We actually already know that at its most microscopic level the universe possesses this computational capacity, because we're building quantum computers every day. In these quantum computers, we store bits of information on individual atoms, we use the laws of electro-dynamics to process information in a complicated fashion, and then we get even more interesting complicated behavior like chemistry. We shouldn't be surprised at this complexity. This ability to produce complexity infects the universe at ever higher and higher levels.

What are the implications of this intrinsic capacity of the universe to generate complexity? There are a bunch of concrete implications. Let's start by testing hypotheses for the origins of life. The first thing that this capacity suggests is that since we know to a very high degree of accuracy a large fraction of reactions for simple chemicals, we can explore the consequences of those reactions. As Bob was just telling us, we don't know a lot of these reactions when we start to include interactions of various minerals. And that's true: we don't necessarily know what all the key reactions are, and I don't think that we should hope right away to be able to show how life exactly started on Earth, or elsewhere if it started elsewhere.

But we have a good chance of showing that something like life should start. If we start with the set of chemical reactions that we know (and we could guess where we don't know what they are), and we try to drive them in different ways, we would expect to see, from this computational ability, that we would start out with a very simple set of chemical reactions, then they start to produce more complicated chemical species, which then auto-catalyze, or catalyze sets of more

complicated reactions, so would you see these species turning themselves on and then turning themselves off as they get consumed by later chemical reactions.

What you would hope to see is as this effective computation proceeds is that it would become more and more complex as time went on, and eventually more stable sets of reactions, for instance the citric acid cycle that Harold Morowitz is so fond of, would establish themselves as the dominant modes of operation. If we saw that happening, that would be very powerful evidence for how life occurred. You would not expect — it helps a lot that Bob and I just talked about this question — to reproduce the exact origins of life because (A) there are many possible sets of initial conditions, (B) the set of reactions could be driven in many different ways, © we don't know what these conditions are, (D) there's a huge number of possible ones, (E) because these interactions are non-linear and hence (F) chaotic in lots of cases, so that (G) they can be very sensitive to these initial conditions. You'd have to get very lucky to find the right ones right away. But you could establish that things like life could occur.

Just as important, you might also be able to establish no-go theorems. If we only involve a certain set of chemical reactions, it's not large enough to be computationally universal. It can only extend a certain amount and then it's just going to produce uninteresting things, such as ABABABABABABABAB — that's all it will produce. It will never produce a varied and intricate set of outcomes. And that you can analyze by looking at the set of reactions and saying these reactions alone are not enough. Hence if you look at your planets and say, hey look, this is what's going on on this planet, then we could say, okay, sorry — no life there.

There are lots of interesting things for life itself that we could look at. One interesting consequence — there are both good and bad consequences here — is that something like life, or the things that come afterwards, you're number 7, and — there could very easily be 7, 8, 9, 10, et cetera, keeping on going forever, if the laws of physics allow it. That's a good consequence. BUT it's not clear right now, given the way the universe is, that something like life could continue forever. If the dark energy persists at the same level that it is right now, then in not a very long time — a hundred billion years or something like that is the number that springs to mind — we're all screwed and nothing can still exist, simply because all matter will have been pushed outside every other piece of matter's horizon, and it can't communicate with anything else, and that's bad.

But it might also be that this dark energy level is continually decreasing, in which case the universe could survive forever. A chapter in Freeman's book from 1982 was very influential for me in thinking about this. He pointed out that if you're willing to slow down and get very large — so you have to slow down and get fat, essentially — then you can still collect free energy essentially forever, and keep on metabolizing and growing. But it would require different technology than just ordinary biological life.

That's the good news. The bad news, at least from the standpoint of a scientist, is the very feature that makes complex behavior arise spontaneously. The fact that something capable of computation will spontaneously generate complex behavior means that it's also not in general possible to calculate, a) whether it will do so in a particular circumstance, or b) how likely it is to do so.

In fact trying to figure out the possibilities for events early on in proto-life, just given the information we have today, is intrinsically probably a very hard problem. If we're lucky and the pathway is not so long, we could figure it out. But if the pathway is very long, and frankly I have to say that given the complexity of ribosomes and the way that life is organized right now, it smacks of being something which is the process of a long and complicated and arduous process of evolution at the metabolic level, prior to the individual level. And that means it could be very hard to figure out what happened. That's potentially bad. On the other hand there is a good thing, which is there is a way to find out what's going to happen in life's future — that is you wait and see. I suggest we do that.

That's all I have to say. I can tell you why there's probably not life in dark energy. Or why there's not life in the first fraction of a second of the universe. But that would not be very interesting.

SASSELOV: Probably the limiting case of much later than the first second of the Big Bang. But three hundred thousand years later, but I bet that there are not enough chemical reactions, which can allow you to do the complex computation.

LLOYD: There's not a lot of free energy in the matter at that point.

SASSELOV: You can actually see that universe — it's observable.

LLOYD: The bizarre thing about the universe is that we understand the origins of the universe much better than we understand the origins of life. It's a simple

system — we've nailed down that 13.7 billion years ago, first this happened, and then this other thing happened, and this happened, and this happened. That's why Dimitar can speak so confidently about how stars behave; it's really really well known. Whereas with regard to the first set of chemical reactions that started life, we just don't know what they are.

SASSELOV: What you said about when you set up these experiments, conditions where you lose — those are familiar to chemists — for example, when everything turns into black tar. Also things just go to equilibrium and not perturbed by any further energy, they just turn it into heat but the chemicals stay the same — that's another 'you lose' scenario.

LLOYD: Even though every atom carries information around with it, in the Big Bang most of that computation that's going on is pretty uninteresting: it's just a bunch of stuff that thermal equilibrium bouncing off of each other. To get interesting things to happen you need the source of free energy. For that gravitation has to kick in and take things out of thermal dynamic equilibrium.

DYSON: Yes. One of the laws of physics which is absolutely crucial, which you didn't mention, is the fact that objects bound together by gravity have negative specific heat.

LLOYD: That's certainly important.

DYSON: That is absolutely crucial. If everything has positive specific heat, as the 19th century scientists believed, then it means that hot objects then lose energy to

cold objects. You are constantly losing free energy, and as hot objects lose energy they become cooler, and cold objects gaining energy become warmer. Everything goes into a uniform temperature and the universe dies and life cannot persist. That was talked about a great deal in the 19th century. They called it the 'heat death', when everything goes to thermal equilibrium so life couldn't persist. But it happens that gravity has the opposite effect; that if you have an object like the sun that's held together by gravitation, that in fact the more energy you give it, the cooler it gets. And the more it loses energy, the hotter it gets.

LLOYD: Yes. If you look at star clusters, they occasionally will kick out a star, and the star will escape to infinity. And if you then look at the other stars, they're huddled together more and they're moving faster. They've gotten hotter, effectively.

DYSON: It means that in fact energy flows from cold objects to hot objects, if they're bound together by gravitation, so that you get further and further from equilibrium. That's the basic reason why the laws of physics favor heterogeneity rather than homogeneity.

LLOYD: Yes, absolutely. That's extremely important. And indeed, it's not clear how far that will go, with this historic dark energy out there. It could be that dark energy is quite useful. We just haven't figured out what to use it for yet. Of course that's the key if you want life to survive forever; you have to do some tricky stuff to harvest energy from further and further away. If you take things and move them closer together, then you can take the energy out of them as you move them closer together. Of course if you do that too much, then they form black holes, and

they're not as useful.

DYSON: Black holes are essential because they are sinks of entropy; you can throw entropy into black holes and it disappears.

LLOYD: That's the cosmic garbage problem we were talking about before is — the ultimate in recycling.

DYSON: You definitely need black holes.

TING WU: Earlier you were talking about 'if-then' statements. One of the things I find life so extraordinary at is self-correction — of chemical reactions as moving around certain pathways that are fairly predictable, they go a certain way — not that this would define life, but it's part of many lives, which is it will go down a pathway and it can sell-correct.

The most dramatic one is when DNA errors are corrected. There's a directionality there that isn't easy explained just by a chemical reaction. I don't like to anthropomorphize either, but it is as if life has a behavior — I shouldn't say a direction — but it's moving along a direction that may not be easily explained by 'if-then'. I was wondering if you could comment on self-correction or self-righting behavior — as a chemical reaction or not — which reminds me that as we try to define life, or figure out the features of life, probably the most puzzling part of life, which we don't have a grasp on, is behavior, and so maybe we're missing one of the key aspects — I know that as a biologist, behavior is almost a complete mystery right now. So we're trying to find life by many things except perhaps one

of the most mysterious things life does, which is behavior.

LLOYD: Interestingly, this DNA correction mechanism which you allude to lies at the very beginning of my own field of quantum computing. In the 1970s Charlie Bennett looked at the thermodynamics of this DNA correcting mechanism, and when you are correcting errors you have throw information away because afterwards you want the DNA to be in the right state, independent of whatever error happened before.

TING WU: Whatever "right" is.

LLOYD: Whatever "right" is. In this case the DNA correction mechanism is detecting to see okay, do these two strands match, for instance; or are they complementary to each other — and if they're not, then you go back and you try to re-write them. Then the information about the error goes away, and it turns out that this has to generate entropy, because the laws of physics at bottom are reversible. They're only irreversible in the macroscopic sense — and that means you can never throw information away for good. So if I throw information about the DNA away, that information has got to go somewhere else. And so these interactions are entropy generating: you have to supply them with a source of free energy and drive them along. In fact if you supply them with too little free energy they'll go back in the other direction and they'll generate errors. So an error-correcting mechanism, if it runs in the wrong direction, is an error-generating mechanism, which is actually also — not to anthropomorphize it — kind of human behavior. The ability to operate in a stable robust fashion in the presence of noise and errors is a key aspect of life, and is not so easy to effect. Particularly

at the level of individual quantum, may I say.

Let's now look at this question of behavior. This computational issue, the fact that things are computing, are processing information, according to things like 'if-then' statements, can be thought of as the origins of inscrutability of behavior, either of chemical reactions or of human beings. Let me phrase this in terms of computers, because then again I'm on safe ground because this is a theory I can prove. There's a famous theorem in computer science called the halting problem, which Alan Turing first proposed. He pointed out, just from the very fact a universal computer can simulate itself — remember we talked about universal computers being able to simulate itself or other computers — that you can construct self-contradictory statements. As a result, certain questions can't be answered by a computer. One such question is, If I change this one bit in this computer program, then it will it stop and give an output? This problem is called the halting problem; it means there's no way to compute what's going to happen when you set a computation in motion, other than actually waiting and seeing what happens. There aren't any shortcuts, is another way of saying it. If something's going through a complex computation, there's no logical shortcut that allows you to figure out what it's going to do, other than going through the computation and seeing what's going to happen.

What this means is that computers are intrinsically inscrutable. When you press 'return' today, everything looks exactly the same as when you did yesterday — today you press 'return' and your computer crashes. Right? Yesterday it printed out your manuscript, today it crashes and takes your manuscript with it. Has that ever happened to anybody here? It certainly happened to me.

This is a necessary part of digital computation. There's no way, in general, if a computer is performing complicated computations — and those computers are performing pretty complicated computations — to figure out what's going to happen except to do it. This also holds for chemical reactions. Because these chemical reactions have the same sort of 'if-then' quality that computations have. That's a simplified version of chemical reactions, of course, but the more complex version is at least that complex. It's at least as inscrutable. Even in the simplified 'if-then' picture of chemical reactions, the outcome of a complex set of chemical reactions is by necessity inscrutable. The only way to figure out what's going to happen, in general, is to let it go and see.

This is why if we're going to figure out what the origins of life are, we're going to need either to do some pretty major experiments and/or burn a whole bunch of super-computer power, because the only way to figure out what they're going to do is see what happens. And if it's true of computers and chemical reactions, it's certainly true of human beings. If I think of the thing that makes other people inscrutable — I'll just speak for myself; maybe you find me completely transparent — I find most interactions with other people inscrutable. Or even interactions with myself.

If I want to see what I'm going to do tomorrow, I'm a free agent, and I'm the only one that's going to determine what that is. But the only way for me to actually figure it out is to go through the thinking process and then to figure it out. The inscrutability of my own actions comes from in part from this essential logical feature that the only way to figure out what's going to happen in a computing

system is to go through the computation. And certainly for other human beings who are at least as complicated as I am, I cannot model what's going on inside their heads, and even if I could, the only way to figure out what they were going to say or do would be to go through the complete thought process they're going through. Which I just can't do.

I would say that computers and chemical reactions share with human beings the feature of inscrutability of their behavior, and there's nothing to do about it. There are things you can try: you can get more familiar with them, you can try to model then better, but you're never going to eliminate the uncertainty and essential inscrutability, because it just is the nature of anything that's behaving in a logical fashion. Bizarrely enough it's like Spock: the Vulcan code makes him more strange and hard to understand than if you were actually irrational. It's rationality that makes us inscrutable rather than irrationality.

PRESS: You mean this metaphor of the computer very literally — you can literally envision the universe as sort of going through a set of procedures that you could trace back.

LLOYD: Yes, I don't even mean it as a metaphor.

PRESS: How do you avoid the Gödel trap, in the sense that there are things that exist that you can't possibly explain the origin of?

LLOYD: Exactly. The halting problem and Gödel's theorem are essentially the same problem — they're very closely related, and Turing knew about Gödel's

work when he came up with the halting problem. In fact he came with the halting problem and the Turing machine because he wanted to write about Gödel's work.

Gödel's theorem is basically the Cretan-liar paradox, which comes from St. Paul's letter to Timothy, which says who's going out to preach to Cretans and St. Paul says, watch out for those Cretans — one of their own philosophers says all Cretans are big bellies, gluttons, and lascivious liars.

The question is, how do you treat someone who says, 'I am lying no matter what I say'. And in the logical sense this becomes a statement. Probably the best one — which is what Gödel used — is to construct a statement which effectively says, this statement cannot be proved to be true. So it's a logical statement within a set of axioms. And there are two possibilities: either the statement is true, or it's false.

Let's say it's false, if it's false then it can be proved to be true — but now you've proved a false statement to be true, and that's really bad because if one false statement can be true, then you can prove all false statements could be true. As my children demonstrate to me all the time. Dad, you just said — therefore you're unreliable in all ways. The only other alternative is that the statement is true, but it cannot be proved.

Such statement is one that is, as it were, inscrutable to the logical structure of the theory. It cannot be proved from the theory, but the only choice you have is to adjoin it to the set of axioms of the theory as an addition axiom. And once you've done that then there are more statements like this — Gödel's incompleteness theorem says that no self-consistent logical theory of beyond a certain complexity,

basically complexity which allows it to compute, is complete. The theory can always be extended in a whole variety of different ways.

PRESS: It means that there have to be things in this universe which are not the result of the series of computations — in other words, they're true, because the truth in this example is something that's produced by these calculations, but you can't find the origins, you can't trace them.

LLOYD: I agree. In fact they can't be derived from those laws because quantum mechanics say the universe is not really a universe but a multi-verse: there are different branches to the universe, in which different possibilities are explored. I would say even in some branches you could say the false possibilities are explored, where the universe is inconsistent and then ceases to exist.

PRESS: It's possible then that life could be one of those things that you cannot trace the origins of.

LLOYD: It's conceivable — there has to be a kind of infinity built into the problem. Life presumably originated in some finite context, so it could conceivably be discovered. But the kinds of finite problems that are analogous to these Gödelian problems are things like NP-incomplete problems where there's a huge number of different possibilities, and you'd have to explore each one to find the answer.

SHAPIRO: I just want to emphasize, lest it slip away, one point which was in the middle of the conversation, which is basically that we may never be able to

capture the actual circumstances that led to the beginning of our life here on earth, because environments may have been destroyed or circumstances changed of which there's no record, but there's every opportunity by experiment for searching elsewhere, to find what are the general principles involved in generating life.

LLOYD: Right. Suppose we start to do these experiments, both real and computational experiments, to say, okay, here' s chemistry, it's doing these funny autocatalytic interactions that are a computation, in the strict definition of a computation, and we're going to explore what happens. Then suppose that as we start doing that, we find things that give rise to complicated behavior. That certainly fits your definitions, Freeman, of proto-life — steps number one and two — and maybe even things that are like step three but what we get are totally different from ribosomes in step three. If this happens then that I would say is very strong evidence that we should expect to have life in all sorts of places, involving all kinds of different ways of living other than having ribosomes.

SASSELOV: That's the question of multiple vs. simple pathways to life. Just answering that question would be essential.

LLOYD: And that it is quite possible — even if it's too hard to figure out exactly how life originated on earth. This is a much easier question I think to answer than the question of how did life exactly originate on earth. Because there you have to figure out the exact initial conditions for this complicated set of chemical reactions and that's going to be hard.

SHAPIRO: And the other point of view has been very much pushed over the ages;

I think George Wald once said that if you study your biochemistry text on earth you can pass examinations on Auctorus. Which is a star somewhere out there — and this is essentially saying the opposite.

DYSON: I did an experiment to demonstrate that genes don't determine personality — I have a pair of identical twin grandsons and it's a remarkable fact: these kids have exactly the same genes and exactly the same environment, and still they're totally different. One is George and the other is Donald, and they know the difference.

SHAPIRO: One may have had a more preferable location in the womb than the other.

DYSON: But anyway, it is a fact that the brain is random in its development, even when the genes are given.

LLOYD: Right, that's right — and also their experiences are different, so they're being, as it were, programmed by slightly different sets of experience — that could have a radically different effect on how they behave.

DYSON: But the microstructure of the brain is different, even quite apart from the experience.

LLOYD: Absolutely. Right. Yes, genes are overrated, too. Not just intelligence, life, consciousness.

ABOUT THE EDITOR

John Brockman is a cultural impresario whose career has encompassed the avant-garde art world, science, books, software, and the Internet. In the 1960s he coined the word "Intermedia" and pioneered "Intermedia kinetic environments" in art, theatre, and commerce, while also consulting for clients such as General Electric, Columbia Pictures, Scott Paper, The Pentagon, and the White House.

In 1973, he formed Brockman, Inc., the international literary and software agency specializing in serious nonfiction. He is the founder of the nonprofit Edge Foundation, Inc. and editor of Edge (www.edge.org), the highly acclaimed website devoted to discussions of cutting edge science by many of the world's leading thinkers, the leaders of what he has termed "the third culture".

Included in his works as author and/or editor are "By the Late John Brockman"; "The Third Culture"; "Digerati: Encounters with the Cyber Elite"; and "The New Humanists: Science at the Edge". In addition, he is the editor of a series of books based on the Edge Annual Question: "What We Believe but Cannot Prove"; "What Is Your Dangerous Idea?"; and "What Are You Optimistic About?"

Brockman has the distinction of being the only person to have been profiled on Page One of both The New York Sunday Times "Arts & Leisure" (1966), and The New York Times "Science Times" (1997).

Edge
IN THE NEWS

SUNDAY TIMES
"Challenging notions put forward by some of the world's sharpest minds."

THE GUARDIAN
"Titillating."

SUNDAY HERALD
"A selection of the most explosive ideas of our age."

THE GUARDIAN
"Scientific pipedreams at their very best."

THE TIMES
"Wonderful reading."

TELEGRAPH
"Strangely addictive."

BBC RADIO 4
"Fantastically stimulating... It's like the crack cocaine of the thinking world.... Once you start, you can't stop thinking about that question."

THE OBSERVER
"Brockman's cross-fertilising club, the most rarefied of chatrooms, has its premises on his website www.edge.org. Eavesdropping is fun. Ian McEwan, one of the few novelists who has contributed to Edge's ongoing debates, suggests that the project is not so far removed from the 'old Enlightenment dream of a unified body of knowledge, when biologists and economists draw on each other's concepts and molecular biologists stray into the poorly defended territory of chemists and physicists'."

THE SUNDAY TIMES
"Brilliant! Stimulating reading."

THE NEW YORK TIMES
"One of the most interesting stopping places on the Web."

THE GUARDIAN
"A stellar cast of thinkers tackles the really big questions facing scientists."

THE TIMES HIGHER EDUCATION SUPPLEMENT
"It is like having a front-row seat at the ultimate scientific seminar series."

KANSAS CITY.COM
"A fascinating site."

THE COURIER-MAIL
"Fascinating...a lot of fun."

AMERICAN SCIENTIST
"Fascinating and thought-provoking ...wonderful, intelligent."

THE NEW YORK TIMES
"Today's visions of science tomorrow."

THE WALL STREET JOURNAL
"You can improve your own science education at www.edge.org."

SONNTAGS-ZEITUNG.CH
"Clever minds debate on Edge about God and the world: what life is, what will result from global warming, or what the most recent discoveries in immunology research tell us. It is almost as colorful as the days of Louis XVI, when philosophers, writers, and political thinkers disputed one another in Parisian living rooms - and prepared the way for revolution."

WIRED
"Awesome indie newsletter with brilliant contributors."

FRANKFURTER ALLGEMEINE ZEITUNG
"Everything is permitted, and nothing is excluded from this intellectual game."

THE SUNDAY TIMES
"Websites of the year...Inspired Arena...the world's foremost scientific thinkers."

ARTS & LETTERS DAILY
"Deliciously creative... the variety astonishes... intellectual skyrockets of stunning brilliance. Nobody in the world is doing what Edge is doing."

PROSPECT
"A marvellous showcase for the Internet, it comes very highly recommended."

MERCURY NEWS
"Profound, esoteric and outright entertaining."

THE GUARDIAN
"A terrific, thought provoking site."

THE NEW YORK TIMES COLUMNS
".... a fascinating survey of intellectual and creative wonders of the world... Thoughtful and often surprising... reminds me of how wondrous our world is." (Bill Gates)

WIRED
"One of the Net's most prestigious, invitation-only free trade zones for the exchange of potent ideas."

SALON.COM
"An enjoyable read."

WIRED
"A-list: Dorothy Parker's Vicious Circle without the food and alcohol ... a brilliant format."

NEW SCIENTIST
"Big, deep and ambitious questions... breathtaking in scope."

THE ATLANTIC MONTHLY
"Has raised electronic discourse on the Web to a whole new level."

THE INDEPENDENT
"Lively, sometimes obscure and almost always ambitious."

Edge
IN THE NEWS

SAN FRANCISCO CHRONICLE
"The world's finest minds have responded with some of the most insightful, humbling, fascinating confessions and anecdotes, an intellectual treasure trove.... Best three or four hours of intense, enlightening reading you can do for the new year. Read it now."

THE GUARDIAN
"They are the intellectual elite, the brains the rest of us rely on to make sense of the universe and answer the big questions. But in a refreshing show of new year humility, the world's best thinkers have admitted that from time to time even they are forced to change their minds."

THE INDEPENDENT
"As fascinating and weighty as one would imagine."

EL MUNDO
"A great event in the Anglo-Saxon culture."

THE INDEPENDENT
"The splendidly enlightened Edge website (www.edge.org) ... I strongly recommend a visit."

THE TELEGRAPH
"Provocative ideas put forward today by leading figures."

NATIONAL REVIEW
"For an exceptionally high quotient of interesting ideas to words, this is hard to beat. ...What a feast of egg-head opinionating!"

THE GLOBE AND MAIL
"A jolt of fresh thinking... This is the intellectual equivalent of a New Year's dip in the lake — bracing, possibly shriek-inducing, and bound to wake you up."

THE NEWS AND OBSERVER
"These world-class thinkers have responded to impossibly open-ended questions with erudition, imagination and clarity."

O, THE OPRAH MAGAZINE
"Persuasively upbeat."

SEED
"Our greatest minds provide nutshell insights on how science will help forge a better world ahead."

THE MAIL ON SUNDAY
"Uplifting... enthralling."

THE SCOTSMAN
"Praised by everyone from the Guardian, Prospect magazine, Wired, the New York Times and BBC Radio 4, Edge is an online collective of deep thinkers. Their contributors aren't on the frontier, they are the frontier."

THE CANBERRA TIMES
"There is much in many of these brief essays to astonish, to be appalled at, to mull over or to wish for... Most of them are vitally engaging to anyone with an ounce of interest in matters such as being or whatever."

SCIENTIFIC AMERICAN
"What's the big idea?... When the lightbulb above your head is truly incendiary."

THE GUARDIAN
"...fascinating and provocative reading."

THE TIMES
"If you think the web is full of trivial rubbish, you will find the intellectual badinage of edge.org to be a blessed counterpoint."

THE HINDU
"Recommended read to detox a tired mind."

DISCOVER
"... reads like an intriguing dinner party conversation among great minds in science. Don't expect to find answers here. Brockman will have you asking more questions than when you started-and may even change your mind about the ideas you've always been convinced are right."

THE SUNDAY TIMES
"Brilliant... a eureka moment at the edge of knowledge, as scientists ponder the imponderable.... Visiting Edge will make pseudo-scientists feel cleverer, and the rest of us more than usually stupid, as we discover, with a jolt of pleasure, how little we really know about the world."

THE INDEPENDENT
"[Ian McEwan] loves the spirited playfulness evident in places such as John Brockman's celebrated website Edge, where "neuroscientists might talk to mathematicians, biologists to computer-modelling experts", and in an accessible, discipline-crossing language that lets us all eavesdrop. 'In order to talk to each other, they just have to use plain English. That's where the rest of us benefit.' "

THE IRISH TIMES
"www.edge.org...has established itself as a major force on the intellectual scene in the US and as required reading for humanities heads who want to keep up to speed with the latest in science and technology."

SKEPTICAL INQUIRER
"Intellectual and creative magnificence."

TELEGRAPH
"Open-minded, free ranging, intellectually playful ...an unadorned pleasure in curiosity, a collective expression of wonder at the living and inanimate world ... an ongoing and thrilling colloquium." (Ian McEwan)

THE BOSTON GLOBE
"Astounding reading."

TORONTO STAR
"...the fascinating website edge.org."

SEED
"An unprecedented roster of brilliant minds, the sum of which is nothing short of visionary."

THE EVENING STANDARD
"Danger - brilliant minds at work...exhilarating, hilarious, and chilling."

THE INDEPENDENT
"Provocative"