

Program Information:

Title: What If the Singularity Does NOT Happen?

Location: San Francisco, CA, Long Now Foundation

Date: Feb 15 2007

Good evening I am Stewart Brand from the Long Now Foundation. This would be the third talk on this thing, you already the first one was Bruce Sterling. That one was called The Singularity: Your Future as a Black Hole. And basically it dinged some of the over enthusiasm of from people counting on the singularity like kind of a technological rapture that was going to happen in their lifetime. And then we had Chris Wild who was basically bet his reputation and career of the singularity is near and we had an evening called Chris Wild's law that was very persuasive on the least so far technological determinism, inevitability of a major transformative and irreversible event happening to humanity in the next few decades. So how amazing - to have the guy who started the idea in the first place. And Vernor Vinge is a science fiction writer of ideas I just heard from Peter Schwartz he actually has an idea box. And out of the box comes, things like his current book Rainbow's End and things like the very early and very influential true names which anticipated a whole lot of what cyber space became. Now did it anticipate it because it thought about what could happen or did the people who made cyber space get influenced by that story? That's the kind of by play that goes very often with his kind of writing and also his kind of speaking which you will hear tonight. Vernor Vinge.

Thanks Stewart it's very great pleasure to be here in and preparing for this talk I had the opportunity to go back over the the talks and the series that are already in place. And it's really a marvelous resource that you in the Long Now have put together. I am certainly very pleased to hear what Alexander is talking about for the reports going now. So my talk here, as the title as you see - What If The Singularity Does Not Happen. Just for the record I should define what I mean what I mean by the technological singularity. And that's it it seems to me possible that with technology, we can in the fairly near future create or become creatures that surpass humans. Events beyond such an event which, what I used the term technological singularity for, are always unimaginable to us as opera is to a flatworm. Well, taking that is to as being what the idea is, it makes almost by definition discussing long term thinking to be in impractical impractical thing. And so when Stewart originally suggested that I come and give a talk about the Long Now I was almost bewildered. But then I thought that that there is something that science fiction writers do and what good scenario planners do. And that is whether or not they believe that a particular scenario was a likely one or not good science fiction writers and good scenario planners think about alternatives to their scenarios.

So it's in that spirit that I am approaching this talk. By the way I am also making one other assumption and that is that we don't figure it out how to get faster than light space travel. Just for the sake of people who are surfing in on this talk perhaps later on the web, out of context, I want to make it clear that I still regard the singularity as the most likely non catastrophic outcome for our near future. There are perhaps ways of stopping the singularity but the only sure far way as I know would be to blow up the world, you know, before we make it happen. There are all sorts of classable catastrophic scenarios about this century. And I recommend Martin Reese's Our Final Hour for some insight into those. For the most part I will try to steer clear of you know, final black out type scenarios here. In their own ways those as unusable for scenario planners as are the scenarios that involve the the singularity although in a much more pessimistic way. So I want to talk about a scenario that hopefully is not

one of the catastrophic ones and that for some reasons does not have the singularity in it. And so I have come up with a the beginning of the scenario is how the lack of the singularity or the singularity failure might appear.

So I have a bunch of symptoms here which is a common thing that stands in scenario planning. In fact one of the more important things is to think about symptoms that might happen if the scenario was going to go in some particular way. And the particular reason probably not pointed to definitively until after the fact you know 50 years from now when you have to write an essay why it was obvious all along that singularity couldn't happen, that sort of thing. I come up with these perspective symptoms. In this scenario we never really figure out how to put the parts together. We get the hardware power but we never figured out how to put the parts together, if you were the more the mystical bench you might say we never found the soul in the hardware. As the years progressed the symptoms for such a scenario might be the software creation continue to be the province of software engineering. You know using Java to solve or to try to solve ever vaster problems ever vaster ever vaster software projects and finding that software projects are failing in very, very spectacular ways. Already software failures are really strange and interesting things large ones. Many large project failures if you have a government budget you just can give throwing money at it to the point where you can eventually eventually proclaim some sort of victory. There have been software failures that have been so large that even though the entities they were funding them had deep pockets, in the end they really just had to give up and walk away and say yeah, we weren't able to make it work. And that by itself is interesting with a large class of failure. Some of these failures also might be interesting in other ways and that is if the people always claim say these terrible things about the limitations of computer programs, in this world where everything stays software engineering I could imagine some very peculiar failures especially in large automation projects where or large control projects where people attempt the total automation. You could imagine human flight controls that occasionally makes mistakes across aircraft running to each other. I could imagine a bug in a fully automatic air traffic control they can actually cause an end way crash which is you know that's really super human in a certain way. So any way if we if we generally got such an era where the largest of these automation projects failed, one might actually imagine manufacturers backing away from their improvement schedules which is really what makes in an economic sense is the is the step that where the actual slope of Moore's Law for the next year or two is determined. And if that happened we might actually see a failure of Moore's Law that wasn't so much because it couldn't be done but just because some of the economic drive forward being done had been had been removed by the failure of of software engineering. If that happen then one could imagine eventually the basic research would itself also level off, both because of funding lack of funding and also because of failure to have the hardware that could support the type of research that would push hardware project further along.

So what you might get is the situation were hardware improvements continued for the longest time on the simplest forms of hardware, for structures are very regular, things like large memory systems, stuffs like that. And in the end we have some extraordinarily good audio visual entertainment products, nothing pose singular and some very, very, very large databases but without software to properly exploit them. If this was how things actually worked out, most people would probably not be surprised in that case if the power of strong AI the promise of strong AI was not fulfilled. And in the same circumstances are actually other associated things that many of us hoped for like Nanotech General Assemblers and stuff like that that also could probably also elude development. So it's not surprising that all together the early years of this time would come into be called the age of failed dreams.

I want to talk about the characteristics out side of these sort of technical characteristics of this age of

failed dreams. By the way about a year ago I brought this sort of general notion up with Hans Moravec, talking to him about the possibilities that you know What If the Singularity Didn't Happen. He had a very interesting reaction. He was he said, that would be so remarkable. I mean some thing that is obviously so inevitable. If it didn't happen, my god, that might be the most interesting thing of all. To try to figure out why it didn't happen? That some fundamental thing about the universe was not the way we had thought it obviously was. Well I don't go there and this I just say we have those symptoms and eventually you know the people who have been so enthusiastic about this get old and and they are the old farts and it's purely clearly pass $\frac{1}{t^2}$ you know, where is the singularity. On the other hand some of the consequences of the situation might seem comforting. There is a thing called Edelson's Law, it says the number of important of insights that are not being made is increasing exponentially with time. And actually I see this as just evidence of what happens if you are trying to confront the production of information and in some cases some thing like knowledge with a merely human minds. That you will see some thing like Edelson's Law. Well, if these are trends actually slow down then one could imagine that there would be time for us to be end or to catch up. Although in some cases such as just the accumulation of bio science micro biology information, about metabolic path ways and genomics and comparative genomics and stuff like that, the data bases could would probably continue to fill faster and faster and we would just be falling further and further behind which actually would fit with the notion of us remaining eternally clueless about some fundamental understanding of how things go.

There is also the prosperity that might warm the cackles of some software mangers and that is if things slow down on the hardware front then finally we go back and do all that crummy software right. I really don't think that will happen. Although in the long run if we were talking about thousands of years - this is getting ahead of my self, but in the long run I think probably in the fullness of time there would be people who would try at you could imagine if you were given several centuries you might be able to redo all the legacy software. And my prediction is that all that does the set up and set up a new layer in the middle pile upon which you will pile still more garbage and in centuries later they will go back and may be if they are foolish enough think about trying to rationalize that. Less comforting is that I think in this scenario that I am talking about; humanities chances for surviving the century would actually be more dubious than they would be otherwise. Our environmental threats that people are talking about so much, they would still exist and to me, actually more serious are war affair threats. Nowadays we are most of us I think properly so are terrified by the notion of nuclear terrorism against cities. But compare that to what happened say what didn't happened between 1970 and 1990, when there were nations states who were talking about exploding tens or thousands of nuclear weapons, many of them over cities, in the space of a few days, in sort of a cooperative endeavor to wipe everybody out. A returned to Mutually Assured Destruction type schemes, it seems to me is entirely plausible and it is something that goes in hand in hand with whatever external or environmental issues are going on. Environmental problems that might merely cause a lot of misery if they were hooked up with national interests that actually are attached to military machines, of especially of a sort that are involved Mutually Assured Destructions schemes, that's really a plausible civilization killer.

I want to talk about envisioning various possibilities for Long Now that would stretch out after the age of failed dreams. So I am going to assume that somehow we managed to survive the the 21st century and we are now in a position to talk about things happening over a very long period of time. Now I want to talk about some some scenarios that actually cover a lot of different possibilities. And I have two or three you know, separate little diagrams for the different scenarios. Typically in diagrams like

this you see people who graph something like population as a function of time. But what I want to do is to graph some aspect of technology as the relationship of it to population size. And here is an example of the sort of path. This is this for the situation in our world up to the present time. Situation so far, all of these pictures are going to be our Long Now on earth. Well, this looks like a lot of pictures I have drawn over the years for the singularity, right? And lot of pictures you have probably seen for the singularity. You notice how steep things were getting to the right. If you look at actually what are the axis, this is not a diagram of the sort that that is advertising some imminent singular sort of event. And so I want to talk about the axis a little bit.

First of all time is not explicit in the axis. The horizontal axis is population size, its not time, its population size. The vertical axis, I wanted it to be some aspect of technology. And so you know I thought about that actually I though about it too much maybe. So maybe I could graph the total effluent or the effluent proportion or something like that have that be the vertical. But there are several reasons for not doing that. One is it would have been too hard for me to research numbers. So in the end I decided just to go for maximum power source and by that I mean almost in a definite way as a discrete power source, what is the maximum power source that that was available to our civilization, when it is size was the size indicated by population number on the horizontal axis. So if I had if I had really good, you know perfect knowledge of the past, the way to generate this to curve would be the just for every year from 50,000 BC on, for that year look out and see what was the maximum discrete power source that humans had usable access to. So that would give me the - the vertical coordinate of a point on the graph and the horizontal coordinate would be what the human population of the earth was at that time.

And as you go along on this graph I you know I spend an afternoon with Wikipedia and Google and looked around and one thing you notice about this sought of graph is, that it actually can talk about a long period of time, and their things on the graph that may be separated by quite a distance on the blue line and actually on that very much time and that we will see more of that later. And there are some places that look quite close that are actually separated by lot time. So the the lower left there is 50,000 BC. I didn't try to chase it back further because, basically 100,000, 200,000, 300,000; we are beginning to wonder whether you are talking about human beings any more. But 50,000 I said okay we have just you know the power of a person. And as they gradually coordinate with each other may be more people. It's sort of matter of human to talk about drooping big rocks on one's enemies partly because it's sort of hard for me to work that out into units of power. But I decided to put in their just or the heck of it. I forget to do over again, I wouldn't have put horses at 6000 BC that would have been oxen at about 8000 BC which I think would have, actually have little bit more horse power than the horses. You will notice that as we go long the graph is rising presumably because we are gating these devices together, but still into manipulatable units. One interesting thing is about how how steady and and robust things seem to be for instance to black death is scarcely a little left ward notch, it actually did decrease in absolute numbers of population. So that's why I illustrated it, actually with a little diversion to the left. I didn't change the vertical coordinates at all. 1850, you begin to get nautical steam engines that could give you a mega watt of power. And then right now or in our era, the largest power stations are about thirteen Giga watts. Now if it isn't clear already let me go to the up right hand corner, where things are beginning to look like they are going straight up.

Let's suppose that we manage to stabilize civilization, which actually looks like we had a shot at it in the next century suppose managed to stabilize situations, stabilize populations. And there was still some modest improvement in the size of power stations that we built. Well in that case this graph, the blue would go straight up, it doesn't mean that anything unknowable has happened. It's just that this is

not a graph of a function. It's just a relationship between the maximum reliable power at a given population size, and you will, in the later diagrams, you will see things, situations were it actually can you know wrap around in various in various ways. So I like this sort of diagram because it's going to allow me to look at the entire human era from, say 50,000 years before now to 50,000 years after now. In other words a whole human Long Now is visible all all at once, which seems sort of appropriate for this sort of a venue.

Now, that was the situation right now. As I say it doesn't look too exiting. But I want to look at three possibilities or scenarios in their own right for this sort of a situation. Here is the first one. I said I wasn't going to kill every body, sorry about that. I want to want to have one where every body gets killed. It illustrates the illustrates the things about this way of looking at the relationship of population and technology, it also illustrates things about where the length of the curve may be misleading. For instance it looks like we got a little bit beyond the present, and then advice about the dangers of Mutual Assured Destruction, strategies were not heated, and we got it back in to such an event and we had a major spasm type MAD war. And so here on the diagram which says that the single afternoon of madness this a this is a few hours or may be a few days that takes the human population down to about 10 million. And then, this is an extinction scenario, it might take a little bit longer to push us back all the way to extinction. It may seem that I have a thing about nuclear warfare. Actually nuclear warfare is one of those tried and not tried in true, but it's one of those one of those things were there are lots of things known about the devices, they can scale way up and the thing about MAD is you actually are getting very large number of people and very large amounts of resources that have been dedicated to figuring out how to cause this sort of extreme damage and there is a sort of magical logic about MAD that keeps raising the threshold, acceptable damage, and to levels of destruction that would be enormously higher than what people thought about what they were going to do begin with. So it's hard to think of any threat where a genius is so explicitly aimed at our own destruction.

What about other forms of technology being brought and to help kill people off. That could be true too. I mean it is it is sort of interesting to imagine what could be done with diseases if your goal was to really make them something that just in particular evaded what the current technology was for curing diseases. I haven't tried to think too much about that but one could imagine that all there. I think Martin Reese made an interesting point in his book which was he was talking about some of these existential threats that some people worry about a lot. Like Grey goo or or other things that involve really a technological tour de force to get within shouting distance of being able to do the threat. And this point was that he didn't think he really have to worry about that. Because basically long before you managed to build the technological house of cards that could support these really cool ways of you know, blowing every thing up and killing everybody off. Long before you get to that level the mundane versions of of destruction will probably by accident or or malignant intend some sort or another. They will cause disasters that will either wipe everybody out or wipe civilization out. So I think that sort of thing is I am concentrating on this one version of it. But that's something to keep in mind that actually we are in a situation where if you put your mind towards you probably could cause an existential threat to civilization. So that was the first one, a return to madness. And anybody who has lived through the era of the fate of the earth and the TTAPS report has has been exposed all sorts of that - stuff like that. I am actually quite skeptical about these two preceding references. But for the reasons that I have I talked about well the slide was up on the board, it's something that I think in principle could be as destructive as as some people have feared in the past.

Ah the next one, the golden age. This is sort of sort of make up for my cheating on I am not well

not killing everybody. The golden age was a non singular scenario where I tried to you know, be as nice as I could. And here the first part of the graph is as before. Here you can see it's definitely not a function of population size. I will have some others where it's obviously not a function of you know, not a graph that represents the function of power source. We are at the year of 2000 or so now. And in this in this version here I have what I have labeled a piece full of scent into a golden age which last about a 1000 years you will notice that there is some technological improvement in power sources in that time. And and actually the population falls to about 300 actually to about three billion. So this is a sort of world that I think a lot of people would would kind of hope for think that it might happen. It means that some how we managed to avoid the existential threats that everybody is worrying about. And and so I think that a lot of time and effort of the Long Now at least in the near time for the Long Now is to think how to have this peaceful ascend into a golden age. And I think there actually are a lot of trends that that make that plausible. I see a plasticity in the human psyche that that can actually cause large numbers of people to behave in very different ways over short periods of time. I think the leveling of the population growth in at the turn of the century here has been the example of that that. That when you give people hope and information and communication, its amazing how fast that they start behaving with a wisdom that exceeds the elites. And this is something that I makes me the most optimistic about the present time. That the governments of the largest countries realize that their national wealth depends more than anything on having large numbers of relatively happy, creative, communicating, educated people. That's where the wealth is. These governments may still want to go on and play their usual games, but there is that notion that they are going to have at least create the illusion of of freedom if they are going to maintain these resource. There is only one problem with that and that is for them and that is they wants to hook up millions and hundreds of millions of people that are educated and can talk to each other, they develop a certain understanding of what's going on and there are hundreds of thousands, probably millions of them out there, that are better educated and smarter than the captains of the state and the advisors of the captains of state. There is just a complete mismatch of intellectual horse power there. And so providing the people with the illusion of freedom would probably have to generate in a better imitation of freedom than has ever existed in human history. Just looking at what such I don't really like the term populism much, but I kind of apply here, sort of a new populism. So new populism in which the participants in the populism have self interest but it is such a wide it's such a wide horizon and it's so well informed that it can actually be mistaken for tolerance. And such people working together in large numbers I think can produce miraculous things. In 1995 as somebody had told to describe Wikipedia to me and told me what it has demonstrably done to date, whatever you may think about the future of Wikipedia but it is demonstrably done to date, I would have said, you should take out a pencil and paper and an envelope and do some simple calculations and will you see this is this bull shit. And further more you will see that you don't really understand the destructive impulse of people who like to break things that are beautiful. And so I am just so I am just so happy that to me this is been a miraculous development, how successful this is. And that's why although I regard the golden age as a as probably too optimistic, because on earth there are just dangerous accidents that can happen. It's sort of like having a even if people of good will if they are standing around in a room filled with dynamite of various sorts, its dangerous situation. But I actually regard this golden age one is not it's not an ideal or an impossible dream, it would be something that could happen in figuring out how to make it happen, or how to make deviations from these ideal, how to make those deviations less destructive and certainly how to avoid extinction, that's a very very realistic goal. The part says a long good time that's kind of

my quirkiness coming out there. I just try to imagine what would happen after that; after we get down a population grand of 3 billion and we are all having a great time and all of that, there is a still a Long Now out there. So I understand the we are talking at least 10000 years. Well 15. I you know this is sort of hazy little blob out there at the end of this long good time. I was thinking that you know I was going to label it 50000 because that then that goes from minus 50000 to plus 50000. That's where I really I wanted to see. And I can't any more say what would happen after that than I could say what is really mean before 50000 years ago. No matter how well we do 50000 years from now is far enough in the future just looking at biological time and things like that that even in the happiest scenarios I really wouldn't expect the human race to be around as human race after that point. And if you think about it that's what in the 20s 30s and 40s when people talked about progress in scientific project progress they were quite happy to talk about good things like that happening the good, unimaginably good things happening like that way far away.

It's only in the late 20th century when people started talking about it happening before the audience gets to retirement age; that all of a sudden everybody gets really nervous about the prospect. So, there is no way that I can see that as live as there is no way that I can see that one could expect the long good time for humans on earth to go on forever. So that's where I put that sort of little hazy, we went to become something better, or bloom there at the end of this golden age. And you will notice that I also had the population gradually increasing. Up to actually probably I got my logarithms a little bit wrong. It looks like I got it increasing up to at least a 50 or 60 billion on earth. And I decided to just do that for the fun of it, partly because I've been so impressed by what large population does to things like the Wikipedia where you've you have experts topics of arbitrarily precise and specialize nature, all talking to each other and all being a creative, so I kind of felt that perhaps there is an argument that having a very large population if you can keep the other issues and an high standard living, it might actually be quite a good thing. And so that's what I show there.

In fact, if we are really talking about 50,000 years in the future, on the scale of this diagram, for all we know these folks actually did some experimentation. Over a period of 4000-5000 years they might drive it down, as they did here to 3 billion. Or may be they pushed it up to like 60 billion or 70 billion and it's, so that's too high. And then a sort of a group concerns it, brought it back to something that was a little bit lower. It's not really possible to see here. But in such a Long Now, there would be some room for that sort of experimentation. I have actually have besides the trends that I have talked about that might make such a golden age possible, I actually have a substitute suggestion that may be a little bit at odds with some of the other talks that I have seen in the series. And here is the suggestion, policy suggestion. Don't want you to take this as special pleading.

The policy suggestion is that all people are good for the future of humanity young-old people are good for the future of humanity. So that prolongevity research may be one of the most important undertakings for the long term safety of the human race. And, when I see most people talk about prolongevity I think actually, Ray Kurzweil has talked about the objections he gets to people talking about prolongevity. Some of the objections are well some of the objections are just very foolish like well, every body would be senile then. Well, you know that's not the that's not what prolongevity is mean here. That's why I said young-old people. And other people say that you know, they were just dead in progress. That actually is very a very common one and well, who knows I mean, we don't have any 500 year people admitting to being 500 years old. We don't have any 500 years old people who actually are still you know young and vibrant. We don't know what they would be like. But actually I think that there is a lot of reason to believe that having a significant percentage of the population that was young and very old would be very, very healthy for the Long Now. One

analogy would it be just the what having very old tribe members did. I think it really helps tribes and the Paleolithic to have members that are very old you know like 40 years old. That this gave a scale of understanding about you know, cycles and stuff personal understanding what things were like doing last dry season.

And I see no reason why the same thing wouldn't be true on longer scale. At the same time, if one talks about self-interest, its one thing to have self-interest in ones great, great, great, great, great, great, great grandchildren and its another thing and it may be a lot more common to some people to have a knowledge of their own self-interest with a notion that they are going to be around after the next election or the next stock market report, or the next 500 years. So if we managed to do that besides having people that have a long perspective on the future a long prospective view when they think about the future, eventually we would have people that actually have a lot of experience about the very far past and this example sort of undercuts my point. But if you know the story, there is this old codge who says you know we try that in 04'. It didn't works then and it's not going to work now. And I could imagine that happening in the future. Only this guy is not talking about 1904 or 2004. He is talking about the 4000's. And he wouldn't be talking like an old codge either in age, you know, he you could look like a fairly young guy and he might be absolutely wrong. But the thing is in the Long Now, having people who actually embody the parts of Long Now it seems to me almost certainly a very positive thing.

Well now we get to the scenario that, if we don't get a singularity, a scenario for the long now on earth that I think is in some form the most likely although as with all of these scenarios I have chosen to make it as extreme you know to make it be an extreme case. And I call it the wheel of time which probably gives it away - but. Now you know what all the white space in the diagram was for. And as I was toiling away with Gibb, you know in trying to make my little bezier curves work on this thing, I was sort of imagining as a science fiction writer what each of those torturous terms must really mean. I mean there is really a separate story or sequence of stories about each of those. Some of them may have been nuclear wars, others may have been some really bad plague or an environmental thing that totally out of hand and some of that would change the actual shape. I generally went for the shape that sort of resonates with the other - rather dramatic loss of population of fall in technology then crawling back up. And so you'll notice that some where in the future we we have here three little turns of the wheel. Now notice those little turns of the wheel percentage wise are much larger than the Black Death was. If these things happened fast they would be extra ordinary disasters. Even these first little three here. And then some thing really bad happened and here you see an excursion that takes us down to a world population of 10 million and and not even not even devices as power as naval no, good naval steam engines. That one I you can see these guys try to climb back up.

How can we know about the amplitude of these cycles? If we look at the diagram, you know I am making statements about how far as the population fallen and in a sense how far the technology fell. But how what would be amplitude? I don't think any body knows although, I think it's worth speculating about. And besides the question of how many people would die and how low technology would go there is the question of how long would it take to track around the cycle. And how high you could get once you came back? You notice in these first little turnarounds that each time the population ended up being quite as big as it was at the maximum before and but we got up to higher levels of power. I don't know that that is true. In the 60's and 70's and 80's there were people like Harrison Brown who wrote "The Challenge of Mans Future", who had this theory that once you went through in an event like this that, what it took to make technology was an infrastructure that once destroyed could not be re built, because we had consumed so much of the resource universe of the world in getting

there to begin with. And so that if we came back at all it would be some crippled state. I think that's very untrue. And John McCarthy is here the audience. And I think has a website that gives all sorts of good arguments while it might not be true that that actually resources can be can be accessed and will exist under all sorts of circumstances. It always seem to me sort of strange talk about a mineral ores being harder to access when you have only cities left over from the last time around. Fossil fuels might be have gone away. In a way that's counter balanced by the fact that I think we probably have left more libraries, more accessible libraries and there are dinosaur fossils that's not sure that's not literally true. But if you look at all the different ways that we have that we have that we have saved our information, in some cases deliberately so it could be accessed in the future and in some case not, but it doesn't matter. A flooded library of paper books that then froze or dried out or something like that, some of them might be totally destroyed and some of them might be eminently resuscitatable. So in this sort of wheel of time scenario, your archeologists at a long haul would be the enduring heroes of civilization. In fact science as such, might except except ,when you get to the top all that things go, new science is not being done except in a very sort of unhappy way. You know new science that's actually urns out to be that the archeologist's just haven't found that page yet. But we actually without running some experiments we don't know. Also we don't know from this diagram really how dreadful this situation is, because we don't know the times. Take what happens after the very bad excursion, we almost looks like somebody was trying for a golden age here. They didn't quite make it. They got up to here and something bad happened to them. I apparently speculated there was a war because I say dropping big rocks on ones enemies, you know asteroid sized rocks in this case. But it may have been that we get hit by a big rock or some other sort of disaster.

In other words life on earth is a dangerous thing and what happened here we may have a very long good run here, may be tens and thousands of years. We eventually came to a almost came to bad end, and you notice a notation here that we we almost lost it on that one. But this is the happy this is wheel of time, so these guys eventually make it back up. So we go round and round like that. And despite the fact that there could be thousands of years that are quite good, there is there are existential threats here and they are existential threats that we really don't we really can't quantify because we are the only experiment. So so as I said we really don't know much about these cycles except that I think is possible that the worst of them could kill every one on earth.

So how to deal with the deadliest of uncertainties. If you look at this talk, I keep saying, who knows, no one knows, I don't think anyone really knows, you know the answers to a lot of these sorts of questions. And I make kind of list here, how dangerous is MAD really. In fact MAD got us to the 20th century alive. You know, it sounds like good thing. We just don't know what how much of the existential threat is posed by environmental change. There are some people who think they know, but I am not sure whether I believe them. How fast we can recover from major catastrophes, how close is technology becoming so good that instead of talking about nation state madness, that just people who are having a bad headache could kill every body. I mean if you want to talk about really high technology, is that sort of thing feasible? And then, the other questions I raise like having what is it a good to have lots of young old people.

So we really there are so many of those things that we don't know because we are essentially running one experiment. And I actually have this tremendous respect for scenario planning. There is another tool that is wonderful, if you have it, and that is simply broad experience. So broad experience means you know, just having other people who have tried out certain things and I can tell you it's a really bad idea. You may even know people who have been killed because they tried some stupid things and that got them killed. At a civilization level, the only examples we have to compare to our examples that are

not comparable, that they had nothing like our technology, they had nothing like our communication abilities in stuff like that, so they are not pure examples. And beyond that there are lots of things about what safe and what's not safe that you really need hundred or thousands of examples. The Framingham heart study is kind of a good example. There are things about diet and exercise that just watching your parents, or even your grant parents is not enough to really get an idea of statistically what how risky certain behaviors are. The Framingham's a Framingham study can provide that. But how do you get that when you are talking about a situation when you have only the one case.

Now I think everybody who has been looking at these diagrams and the titles, you notice that the upper right hand corner of each of the diagrams, it says the Long Now our Long Now on earth. Well, the obvious thing that I am leading up to here is that self sufficient of earth settlements are our best hope for long term survival, not of life on earth. Although I will argue that it actually would improve the prospects of our having the golden age, but for the survival of the human race. And I am so pleased that people like impressive people like Hawking and Dyson and Reese are talking about brought that back to the centre of discussion. And people have been doing that for decades. Stewart Brand had his a book on "The Space Colonies", in the 70's. And having this come back to be at the centre of discussion I think is actually quite important and I think its actually very, very important to in particular the Long Now. I also want to say that again this is not the way to avoid the singularity. I think the singularity can happen, it is going to happen and it would happen in the singularity situation too. But along with other sorts of other things we don't know about.

What are the some objections to having talking about really serious self sufficient settlements in space. The first one is which I get very seriously from very serious people that chasing after safety in space would distract us from solving our life and death problems on earth. I you know, I think the situation on earth is sufficiently dangerous that a moral stand on this issue is a you know against space colonies is very, very dubious. Chasing after space assumes the real estate is not already taken. That's actually a possibility. That and the singularity are two of the most important practical mysterious that are to be face. A real space program would be too dangerous in the short term, that's what I don't hear very often. I think it's actually some virtue in that, you remember in that last slide, things blew up because people are dropping big rocks on their enemies. A real space program meaning cheap access to space is very close to providing a WMD capability. This may be one of those very rare circumstances. Normally when people talk about things they want to do, they are dangerous in the long term, but there are some gratification in the short term. Here you have something that might conceive it well that I think is very good for the long term and might add somewhat to risk on the short term. That is going to be a peculiar situation; I don't think the existential risk is that great. But I wanted to put it on the list. And then there are the practical objections. There is no other place in the solar system to support human civilization, and the stars or too far.

Well being a child of 1950s science fiction, I could go on about this another hour or so, but let me just say, the stars are not too far, even at relatively low speeds if you are a Long Now type of person. And in fact I thought this was cool, Robert Heinlein had this book, "Time for the Stars". He said, the long range foundation is a nonprofit organization that funds expensive long term projects for the benefit of mankind. So Danny and Stewart are still out there and I think this is just a perfect fit. Now, I really in some ways as far as propaganda, the real thing here is, I think I am talking about a real space program. And that's not what we have. I've we started with launch cost \$5 to \$10000 a kilogram. As far as I can tell, the vision for space exploration, which is currently the manned space flight initiative of the administration, still and is talking about \$5 to \$10000 per kilogram, just to low earth orbit. This is this makes any talk of humans in space a doubly gold plated sham for two reasons. One of

course, that it proposes pitiful limitations on the delivered payloads. And the other is that it means the payloads themselves have to be so reliable and so compact that they are enormously expensive. Now, I am addressing this to humans in space but actually because I am not you know, I am not trying to alienate astrophysicists and other forms of astronomers and stuff like that. But I really the the astrophysicists, the cosmologist, they have also been mugged by these two gold plated reasons. First of all, that they spend they have to spend enormous amounts of money for their projects in space which then may be de-funded, you know at a political whim. And the other thing is that if the launch cost were lower, just think if you are involved in these programs, what you could do with 500 dollars or pound to low earth orbit or \$500,000 a kilogram to low earth orbit, what you could do is that? I think in many cases an argument could be made that even with those still relatively high expenses, you could do things like make space telescopes that are cheaper, large ones, that are cheaper than what you could do on the ground. Because, at those prices what you could put in space is stuff that could use things like laser positioning and synthetic aperture that would probably be orders of magnitudes simpler and cheaper than trying to do that on the ground.

So, I think its something even the robot or space people should think about, especially the scientists. I am talking about here just for the a human based program in space. And that's why I think people all through the 20th century; average people understood how important it was. And you know got way how important it was to have to humanity into space. And I think their enthusiasm has been abused and because I think they are getting human kind into space is so important to long term survival, I would urge that we reject any major humans in space initiative that does not have as a prerequisite, the goal for much cheaper access to space.

Well, I think I'll end there. I think this is where the long central to serious Long Now survival of the human race, thanks.