

God Mind and Brain

CiS – Faraday Public Lecture given by Prof. Bill Newsome on 27th April 2010

It is a great privilege also to welcome our guest lecturer this evening, Professor Bill Newsome, who is a Professor of Neurobiology from Stanford University. Well known for his work in cognitive neuro-senses, also a member of the US National Academy of Sciences, has won numerous prizes which I will not even attempt to list this evening because I do not want to take his time. But we are delighted that he is able to be here despite a slightly difficult journey, but no volcanoes this week, thankfully, so it is good that he is here and good that you are here as well.

This evening he is going to be speaking to us on the topic, God, Brain and Mind.

Bill Newsome
Professor of Neurobiology, Stanford University

It is a pleasure for me to be here back at the Faraday Institute after an absence of three years. I certainly enjoyed the last time I was here and I looked forward to this visit. And as Denis says I was only fighting the weather this time, rather than volcanoes so I arrived not when I expected to, but I arrived at about two o'clock this morning and I am in fine shape. So I hope you are too and I hope we have a good time this afternoon.

Science, Religion and the Causal Status of Mind

This is somewhat a pretentious title here, God, Brain and Mind and I think that if I were going to get sub-titled it will tell a little bit more concretely what I intend to be talking about, it would be about this. It would be about science and religion and the causal status of mind. I have to tell you a little bit about the preliminaries just as disclosure about who I am and where I am coming from on this topic.

I am a practicing Christian and I am the first to admit that I don't practice hard enough most of the time. But I am practicing nevertheless and I am a person who has been struggling and going back and forth over these issues as a professional neuroscientist and also as a person of faith trying to understand what the meaning of our resistance is in line with particular for a number of years and a lot of the things that I will present today come from ruminations I have collected over years. But also from new insights or things that I have put together and I hope they are

insights from numerous conversations over the last ten years since I first began speaking publicly on this topic, including some of the conversations I have had with the Faraday Institute.

In a standard talk that I give on this topic I try to hit three basic topics. Interaction between religions and findings of science. It is common these days to think sometimes that science has made religion irrelevant to things that we found out about science, the actual findings of science contradict religion therefore we should throw religion overboard and believe in findings of science. I do not have a lot to say about that today because I think it is fairly well trod ground and people have spoken here at the Faraday who can address this much better than I can. But I will just say that I do not find that to be true. I find that in scientific discoveries like the Big Bang are actually remarkably consonant with religious faith in saying that there is a beginning, a moment of creation to the world. I think that the theory of evolution by natural selection is not a big problem for religious faith as long as we do not take the creation stories literally as a scientific description. Sometimes it is thought that mechanisms of random chance such as in the mutations are inimical to the purpose that we ascribe to the universe from a religious point of view. But I do not think that, I think that we use random events for purposeful ends all the time. We do it in our own laboratory that is why people get paid great sums of money to programme really good random number generators because we put them to useful ends all the time. So I do not really think that there is any big problem with religion and the findings of science.

Religion and the assumptions of science, however, the habits of mind. There is tension there frequently. And I do want to say a few things about that at the beginning of the talk today because this is something that I find comes up over and over again in my own conversations with students, with faculty colleagues, with people, laypeople who are not scientists but still read a lot of science and do a lot of science and think about science a lot. So I will spend the first 15, 20 minutes or so talking about that. Where I really get into the need of the thing as far as neuroscience is concerned is in this last section here, causality, explanation and the brain. When we think about the brain, when we think about human behaviour, what counts as cause? That is going to emerge as the really key question and that will take up the bulk of the talk.

So let us go first to religion and the assumptions of science. And the fundamental observation that I make over the years is that tension can occur between religion and the working assumptions of science, everyday science, and in particular when those assumptions of science are elevated to the status of an all encompassing ideology and their habits.

So there are differences between science and religion as we all know. Science tends to be experiment based. It is not always experiment based but it tends to be. It tends to be precise, it tends to be objective in the sense that the results of science are transferable across communities and cultures. So if I do an experiment well in my laboratory at Stanford and I write it up well it can be replicated at a laboratory in Tokyo or Cambridge or any place else. It is shareable across communities like that. Whereas religion tends to be more holistic, it is a greater dependence on intuition and requires commitment in the absence of proof. Now I do not stand these two things up in absolute contrary distinction to each other. Science requires faith, you have to faith and rationality and orderliness of the universe in order to get out of bed in the morning and do an experiment. And religion is not wholly personal and intuitive. Most religious people practice their faith in their community, and in checks and balances within the community. But still there are real differences here. When I talk to some of my professional colleagues sometimes about these sorts of things, the question that rises is well why do you even go there? Why even go to religion? Why not just stick with science? This comes up in all kinds of ways.

One of my colleagues, we were having a very good conversation actually about these kinds of issues and he is not a religious person at all and finally he looked at me and just pure and simple, "I just do not understand how you get there. You must use a different part of your brain when you do that." That is a neuroscientific question, it all comes back to neuroscience in the end. I also care if people express surprise, sometimes people might laugh who I talk to about such things because I have very high standards as a scientist for what I am willing to believe. And they are amazed that I am willing to believe certain other things based on what is proof when I am handling science? They say it is almost like two minds, I had the postdoc say to me one time "I do not get it Bill, this is so different from your normal way of thinking." And of course when he meant by normal was my scientific way of thinking. I am coming up against this so often that the discussion that I want to have here for a few minutes is regards to this part here. But importantly the religious mode has gone, I

will say religious but it is really a more holistic and in tune about a thought and belief is a normal and necessary mode of evaluation and decision making in real life, for all of us. So I have heard scientists say when topics of faith or religion come up I have heard scientists say no, no, no I am not religious, I am a scientist. And my point here is going to be that the religious mode of thought is a normal and necessary mode of evaluation for all of us.

Well what do I mean? I will say what I mean by that. But I want to contrast this with the scientific mode of thought which is quite peculiar. I am going to argue it is applicable to a narrow range of experience and it is generally practiced by a rather small community of professionals. I would make this assertion that the most important questions in life are not susceptible to solution by the scientific method, when it comes to the most important questions in each of our lives, they generally are not approachable by science. I would be more puckish and provocative in saying that the importance of a question tends to be inversely proportional to the certainty with which it can be answered. Now that is not always true. I think that survivability of life on this planet is an important question and science is going to play a really big part in that. But in general I think that the importance of a question is inversely proportional to the certainty with which it can be answered.

Well what does he mean by that, what do I mean by that? Well here is an example of a question that you cannot go into the laboratory and answer. Is it better to live or to die? That is a real question for people who almost all of you know at one time or the other. And it is an important question. It may be a real question for people in this room. And it is not a question that we go in a laboratory and do science on and get an answer to. That simply is existential for all of us and science has nothing to do with that answer.

Here is another one, less perhaps existential but still important. Should I pursue a professional opportunity elsewhere in the country at the cost of uprooting my entire family, my kid's school, all of whom have their own independent lives? And again we cannot go into a laboratory and do an experiment to get an answer to that question. And the short reason why, for these things which are so momentous in life is we cannot do a controlled experiment. We cannot run that tape back and take alternative decisions and play them out on all the ways that you would like to see it played out and get a rigorous answer to that question of what is the best answer to that question? That just is not happening for most of the important decisions in our

lives. And at the risk of beating at that horse and I find that sometimes I have to beat at that horse to get some of my colleagues at least to think seriously about this.

Here is another, should I marry this particular person? A lot hangs on this decision. All of you that have been through marriage and those of you who have been through a divorce know that this can be one of the momentous and meaningful experiences of your life and it can also bring on great tragedy and great sadness in your life. And in general in making a decision like this we think hard about it, we get the evidence that we can get. Something about your potential spouse's family background tells you a lot, your experience with them during that courtship, trying to make decisions together, going through difficult times together you acquire data. You think as hard about it as you can, you think rationally as much as you can. But in the end if you wait for scientific proof of that kind of quality that this is the person that you should marry or not marry it will never happen, it just will not happen. This is not a decision that is susceptible to that kind of approach.

To put the things simply I would say this is the human condition. It is life and our most consequential decisions in life have little or nothing to do with science. And this is true for everyone, it is true for everyone in this room and it is true for my scientific colleagues. So I get this reaction, this irritational reaction when there are people who I meet and say oh no I do not have anything to do with that, I am a scientist. Because the truth is that all of us are doing this kind of seat of the pants epistemology, making this kind of seat of the pants decision making and a lot is at stake and we are doing it all the time and every day. And I think for all of us the realer question is, rather than parson ourselves in science, religion whatever, the real question for all of us is, is there an ultimate source of meaning and value in the universe? And if there is what is it?

So I would say that the religious quest involves the same sort of reasoning and thinking and acquisition of data as the marriage example. We have sources of evidence, we do not check our minds at door, we have primary experience in the religious community. There is going to be good there but there is going to be some bad, frequently some of each and we have to sort out the wheat from the chaff. There is the testimony of people who have gone before us through centuries. Older people, companions in our own lifetimes and there are the critical reflections and critical thinking of fellow pilgrims or searchers that I meet along the way. But in the end this kind of evidence is not compelling in the scientific sense of the word. Faith

accompanying our commitment is essential and the stakes are really high, your life is at stake here. And the deepest most meaningful things in life are at stake here.

This module is very different from the scientific kind but we try to do one thing and then one thing and then you deduce and you induce and you build this pyramid of knowledge. But this is much more like a web, it is much more like you put your foot out and test the ground and see if it holds your weight and if it holds your weight then you go on to the next step. And this kind of epistemology the writers of the New Testament knew a lot about. So one of my favourite chapters from John's Gospel, John's Sixth, it is at the end of a long chapter and Jesus did some really very difficult teaching about transubstantiation and things. So it is after this many of his disciples drew back and no longer run around with him, this stuff was getting too hairy. Jesus said to the twelve, "Do you always wish to go away?" And Simon and Peter of course answered "Lord, to whom would we go? You have the words of eternal life." I identify with Peter here, I have been banging around the academy for 35 years now and I have not found a better source of words of eternal life and wisdom for living than I find in Jesus' teachings.

But the really epistemic point, the epistemology here that I want to make comes in the very next verse, the sixty ninth verse of this chapter and Peter says "And we have believed and come to know, that you are the Holy One of God." And it is very interesting the order here. There is belief before knowledge and that is the way that it happens so often in our lives, we have to take a step, we have to take a step and believe, we take a gamble if you want to and we come to know down the road whether it is right or whether it is not right. But that belief and that movement based on commitment in faith is necessary first.

So just to re-pictulate this first part of the talk, I do not see deep conflicts between my religious faith and any actual findings of science. In fact they resonate strongly at certain points and if any of you would like to discuss some of those afterwards, by all means let us get into the question and answer and talk about some of them. But in the interest of time I am not going to say more about them now.

The tension can occur between religion and the working assumptions of everyday science, in the assumptions that we start with in the laboratory about how we create knowledge and how we develop belief, there can be tensions if you like my postdoc says "But Bill there is other stuff in your life that is so different from the way you

normally think, your normal is science." Well I would say science is not normal, this other way of thinking is the primary way and it is the science that is abnormal and it is applicable to a slice of reality that is important but not all encompassing and not all important.

Let us talk now about the brain. So what about the brain? Now I am and neuroscientist, I had to get that little peeve off my chest, you can send me your bill and our therapists are probably cheaper in England than they are in the United States so send me your bill but I had to get that off my chest.

What about the brain? So here is half of a brain, we all have one of these things sitting inside our head. And I want to start here with the central dogma of neuroscience. I phrase this very carefully. The central dogma way that I would phrase it is that all of our behaviour in all of mental life including our sense of a conscious continuing self is inextricably linked to the biology of the brain. Now I have chosen my words carefully here. I do not say that it is reducible to biology of the brain, I do not say that it is necessarily completely constrained by the biology of the brain but it is inextricably linked to the biology of the brain. About the business as neuroscientists, we are proceeding along doing our research, taking every bit of human and animal behaviour that we can grab and we produce [unclear-15.36] in the laboratory and we are trying to understand the neuro-mechanisms that underlie that behaviour.

In my own laboratory and the laboratories of a few of the other neuroscientists here in the audience right now we have gotten pretty far. We are up into now from simple sensory kinds of experiences and simple kinds of performance. And I say simple in quotes because when you get inside here the reality is nothing is really simple. We are up to things like decision making and simple forms of decision making and we are studying decisions and it turns out we can record from neural activity in the brain, we can build models that in a simple foraging test can predict about 80-85 percent of the decisions a monkey will make in a foraging test before he actually engages in them and actually reveals the decision. This is pretty amazing stuff. We can actually record neural activity inside the brain that we can listen to or our computers can record that we can know which decision an animal is going to make at the end of the trial several seconds before we allow him to reveal it.

So we are getting into territory now of where we are in cognitive neuroscience, we are actually looking at the neuro basis of pretty high level cognitive functions. We are rowing merrily along with this and sometimes you will find neuroscientists in the lab or over lunch talking about where does this end? What does this say about what it makes, what it means to be human? Not typically, we do not really talk about that kind of stuff very much. It usually takes an undergraduate who hears one of our lectures on another graduate course and pots the question, what about free will? This is not the kind of thing neuroscientists tend to talk about over lunch or dinner, or even here that much.

But it is an important question because this goes to our internal sense of what it means to be human and what kind of control we have over our behaviour and what kind of choices we are really able to make and where does the line get drawn between voluntary behaviour where we have some control and involuntary behaviour which we all know surrounds us all the time. This is a philosopher Woody Allen I think, who said "I am not a fatalist but even if I were what could I do about it?"

So that is one approach to the problem of free will. But there are other approaches as well. And this is important not just for traditional religious understandings of who we are and responsibility before God but it is true for us of citizens of democracies and who we are and to what extent we are responsible before the law. And it is a point that I will make at the end of the talk but this is actually incredibly important for science as well. I will come back to that thing but I will plant that idea in your head.

So incredible responsibility. We are seeing more and more in the courts in the United States these days that brain scans, neurobiological evidence is being offered in courts as an exculpatory evidence. They are being introduced as evidence that defendants had diminished or no responsibility so they deserve less punishment. So this is my neurons made me do it argument. It was not me, it was those neurons inside my head, I did not have any choice in the matter. Now this is not a big issue. We have had issues about responsibility before the law for years.

This is one really interesting example. This young man Charles Whitman was an ex-marine core veteran and an undergraduate at the University of Texas in the 1960s and probably none of you have ever heard of him, I would be surprised if anybody has. Any of you Americans, anyone heard of him? He was the Texas Shooter, the

Texas Sniper in the 1960s who climbed to the top of this tower at the University of Texas, way up there, and opened fire with an automatic rifle and he killed 14 people and wounded 38 others before he was stopped. Before he went into the tower he actually killed his wife and his mother with a hunting knife, not even with a gun. The most remarkable thing are the diaries, he was a diary keeper and the most remarkable thing were the diaries that came to light afterwards.

To my mind these are some of the most provocative writings that I have seen in the area of mind, brain and controlled behaviour and I am willing to bet that no-one in this room has read these things. They are available out there on the internet if you want to read them. This is an excerpt from Charles Whitman's notes that I got from Walter Sinnott-Armstrong who is a Professor of Philosophy at Duke. Whitman says several days before climbing the tower, he says "Lately I have been a victim of many unusual and irrational thoughts. I consulted Doctor Cochrum at the University Health Centre and asked him to recommend someone that I could consult with about some psychiatric disorders that I felt I had. I talked to a doctor once for about two hours and tried to convey to him my fears that I felt overcome by overwhelming violent impulses. After one session I never saw the doctor again. Since then I have been fighting my mental turmoil alone and seemingly to no avail. After my death I wish that an autopsy would be performed to see if there is any visible physical disorder."

And then from another entry "It was after much thought that I decided to kill my wife, Kathy, tonight. I love her dearly and she has been as finer wife to me as any man could ever hope to have. I cannot rationally pinpoint any specific reason for doing this." And Charles left another note with the body "If my life insurance policy is valid, please pay of my debts, donate the rest anonymously to a mental health foundation. Maybe research can prevent further tragedies of this type."

So this is amazing because here we see someone acting under compulsions that are clearly coming from within his brain, within the central nervous system. But at the same time another part of his brain is able to stand aside from this, so reflect about what he is doing and wonder about what he is doing. He says "I cannot rationally pinpoint any specific reason for doing this." There is a rational part of his going and then there is this driven part of his brain going.

Now as it turns out after he was killed by security forces, or maybe he killed himself I cannot remember that detail actually. There was an autopsy performed on his brain and when they did the autopsy on his brain they found a tumour. Now the tumour was in the midbrain, sorry not the midbrain it is in the forebrain, but it was in the hypothalamus and it was impinging on this structure right here, probably an amygdala. So it had a strong impingement on the amygdala that was sitting on the hypothalamus. These are two structures that are part of the limbic system, a system within the brain that is known to be involved in the control of emotion and the control of emotional reactions to events in the world. What are the odds of this, that this guy who became a killer and who was struggling with this mental and emotional turmoil, what are the odds that you are going to go in there and find that amygdala with an actual tumour? We don't know that that was the cause of factor but the probability of this seems really vanishingly small, it seems very likely that that tumour was the causal factor. So in the end you ask yourself, would that have been some kind of explanatory evidence?

Of course today we can look for tumours like that with MRI, you do not have to wait for an autopsy. But it raises the question, where is the line between control, where we really have control of our behaviour if indeed we do have any real control of our behaviour, and the line between uncontrollable behaviour? It goes back to that undergraduate question, how do we think about free will? In a day where behaviour is increasingly understood to come from the operation of systems and neurons in the brain, how do we think about choice and how do we think about control or lack of control of behaviour? It is an important question.

We come back to this question. There are a few ways that people think about these things, most of which I do not like and I am grappling with a different way and trying to read some stuff that maybe offers a different way of thinking about it. And I want to tell you about ones that I do not think work too well, at least for me. And then I will tell you about whether I suspect some potential way to talk and think and understand ourselves and our relationship to the mechanisms in the brain. But I warn you ahead of time that I do not have pat answers for this, I do not have an easy solution that I can present to you with a nice bow wrapped round it, but maybe I can offer some ways to think.

So first of all let me talk about some of the ones that I do not like very much. So I do not like bottom-up determinism very much, I do not like very much this point that

we treat individual neurons, these nerve cells in our brains as machines and they are going to do exactly what they are going to do and our behaviour is going to emerge from it and that behaviour is determined by emotions of the atoms and the electrical activity of neurons in the brain. The truth is we do not have anything resembling real or free choice, real or meaningful choice about anything we do. It is all determined in the beginning. I do not like that too much and I will come back and tell you why I do not like that too much.

Some people and some very smart people and I think probably some people in this audience think about the brain in terms of quantum effects and that through quantum effects and the uncertainty and indeterminacy of quantum events that there may be ways to think creatively about sources of human freedom at that level. Well it is not very attractive to me and I am going to tell you why but I am going to do this humbly because I am not a physicist and I will tell you why I do not think this but in the discussion period maybe some real physicists who think about this can tell me why I should be thinking differently.

Here is another thing that I do not believe in. I do not believe that freedom means uncaused. So over and over again when we get into these discussions of free will, what it comes down to is that if you can pin, put your finger to a cause for some behaviour then it is not free, it is caused and it is not free. And the implicit equation there is that freedom equals uncaused. I do not believe that anymore. I think maybe once upon a time I did, but I do not believe that anymore. So what I think we have to think of and where the solution to this problem lies and in fact I do not even like this term anymore, free will, because this equation that freedom means that there is no cause that is so deeply engrained in so much thinking but the truth is I do not even like to use the term free will anymore. This is why I like to use self determination or autonomy.

So it is not a matter of whether behaviour is caused, I think behaviour is caused. It is more a matter of what really counts is what counts as a cause and what cause is active at any given point in time. And what I want to think is that I have control over my behaviour to some extent. I have some autonomy and I have some self determination. There are some behaviours I do not have any control over and I know that, but I can tell that when I lose control, sometimes the verbal parts of brain tries to rationalise it I truly lose control sometimes.

But how do we think about self determination or autonomy? Because I think this is what we really want, we want our choices and our behaviour to be consistent with our past, with our experience, with our tastes that we come to develop, with our values that we have adopted and that is what we really want. When we say freedom we do not really mean a cause, we mean the causes that come organically out of our experience and who we really are and that is what I am going to argue.

The real problem here in most scientific dialogue is that in experimental science we are usually involved in taking things apart into smaller and smaller bits and basically reducing some phenomenon at a high level like memory or like a tension or like detailed perception and we are looking for the mechanisms that make those things happen and in a sense we reduce a high level behavioural phenomena so we a lower level, the next lower level. Now of course that next lower level in science tends to get reduced to another level and another level and another level. So you have this problem of reduction and it is in part important because it tells us the mechanisms and it helps us understand things in our environment and it gives us control over certain things in our environment. It helps us manipulate things in useful ways and that kind of reduction is really important. But that kind of reduction can be dangerous. What is reduction really? And it turns out when you really start looking at this there are lots of different flavours of reduction.

Now the classical model and I am quoting here from Carl Craver who is a Philosopher of Science at Washington University, St Lewis and I will come back to him later in the talk. And he has this book this challenging book but a really interesting book called *Explaining the Brain* that was published three years ago. Craver is talking here about the classical model of reduction: "According to the classical model of reduction for which most current models descend, reduction is a species of covering law explanation." Covering law has a technical meaning and philosophy of science. "One theory is reduced to another when it is possible to define the theoretical terms of the first with those of the second and to derive the first theory from the second." So we take the terms and at first we might use terms like we are paying attention to the right side of space or the left side of space and we reduce that to the activity of certain neuro systems in the brain and in principle one could derive if you are doing real science according to classical model of reduction, you could then start at the lower level and derive at the higher level. This obviously is inspired by physics. This kind of reduction works and it's advertised at least, that it works well in certain areas

of physics and chemistry. Some physicists would question that but that is the ideal that is held up for all of us who are doing sciences.

Now the metaphysical things that go along with reduction come in very slightly. So deeper and deeper commitment to this reductionist model and here real truth and real explanation lies. Again quoting Craver he talks about the metaphysical fundamentalist. So the fundamentalist here is not about religion it is about fundamentals. Fundamental forces and fundamental courses. He says, "The metaphysical fundamentalist argues that non-fundamental things," that is things like neurons and genes and behaviour and people "have no causal power over and above fundamental things. They believe roughly that everything has cause at a fundamental level and that is a principle of causal completeness of the physical and that nothing has more than one complete cause. This is the principle of non-over-determination. And if so it follows that no non-fundamental things are causes." That is the deterministic view of humans and behaviour, this is what it is, it is reductionism and it is this reductions, fundamentalism that says and my view there is a lot, there is a strong case to be made with this and what it says is that you take these higher level phenomenon, you understand the mechanisms, then you understand the mechanisms and so on, and at the end you get to a fundamental description in terms of physics and that would be quantum mechanics in our day and age, the Schrodinger wave equation. And that is where you have truth, that is where real truth resides. That these other things of convenient epistemological structures for describing behaviour. But when you really have a successful completion of scientific programme and you get down to the fundamental level where real truth resides that is where scientific explanation really happens. And then you are down with the level of forces, electromagnetic forces, gravitational forces, weak and strong nuclear forces. Those forces that we know about through physics and integration and Schrodinger wave equation and that is where real truth lies.

Now here are my problems with fundamentalist reduction and I have several of them. So what it doesn't work for me in real life. It does not give me any help or it does not ascribe that I at least feel like I am doing in real life. I cannot run around and live my life as though I am simply the probabilistic expression and some specific instantiation of assuring the right equation. It does not help me. And something that does not help me in real life is a philosophy that seems rather abstract to me.

Secondly it does not describe what neuroscientists actual do. And the 30, 35 years I have been neuroscientist I have never done classical reduction. I have never actually been able to reduce one series of theoretical entities to another. I have never been able to derive something in a higher level from a lower one. This may be a salutary and may actually describe what happens in physics and certain areas of chemistry but it does not really describe in my opinion what neuroscientists do.

Then there is a regression issue. Whose fundamental level is really fundamental? In neuroscience where I work there is a very powerful lobby centred in the Howard Hughes Medical Institute who pay my salary of genetic fundamentalists. That when you reduce something to the genetic level that is where real explanatory power lies. Genes are what the information of life is about and when you reduce something to genetic explanation then you have really got truth and all this other stuff that people with electrodes and electrical activity were doing it is okay, it is fine but when you really get genes then you have it.

There are people people who are single cell fundamental and they are people who are circuit fundamentalists and there are people who are chemical fundamentalists and the problem with this regression issue is that whose fundamental is really actually fundamental, is indeed proper? If you really want to take the reductionist's philosophy all the way to its core, right down there in their wave equation there is bad news down there at the end. The bad news or thinking about causality is that at the most fundamental level life and science is arguably a causal.

So this point was made by Bertrand Russell, it has been made by lots of people but is most perhaps articulately presented by Bertrand Russell. He wrote this famous essay on the notion of cause published in Proceedings of the Aristotelian Society. He says "In advanced sciences such as gravitational astronomy the word cause never occurs." And he takes issue here with this Doctor James Ward, a philosopher. "Doctor James Ward in this Naturalism and Agnosticism makes this a ground of complaint against physics: the business of those who wish to ascertain the ultimate truth about the world he apparently thinks, should be the discovery of causes, yet physics never even seeks for it." And he goes on, "To me it seems that philosophy ought not to assume such legislative functions and that the reason why why physics has ceased to look for causes is that in fact there are no such things. The law of causality I believe like much that passes muster among philosophers is a relic of a

bygone age, surviving, like the monarchy, only because it is erroneously supposed to do no harm." I wish I could write things like that.

When you get down to the bottom it is not even causal, you write an equation like F equals an A or any other equation you want and there is nothing in that equation of this causing that. You have got to rearrange the terms so that anything you want appears on the left and anything you want appears on the right and you just observe regulatory among variables in the world. There is not really causality there at the deepest level, according to Russell.

My last problem here is the poverty of quantum mechanics. I want to be careful here because I know there are probably physics here and there are people who know a lot more about quantum mechanics than I do. And what I want to say here is that quantum mechanics is not wrong it is just impoverished. So if we undertake this reductionist enterprise and we think that truth really lies at the most fundamental level we come up with a problem. We come up with a wave equation, that is our most fundamental theory of reality in 2010. In principle if we were smart enough and had the resources we could write a wave equation down that would predict the motion of every atom in this room for the next ten minutes. It would be a probabilistic prediction because quantum mechanics is inherent and probabilistic. But if we are smart enough and had the resources we could do this. We could write this equation, everybody predicts the motions of the atoms in this room for the next ten minutes and you say to yourself that is all the knowledge we can obtain, that is ground truth about this room. My problem with that is that the wave equation knows nothing about the people sitting in this room. The wave equation knows nothing about Downing College, it knows nothing about intellectual interest, it knows nothing about hunger, it knows nothing about boredom, it knows nothing about neurons, it knows nothing about genes, it knows nothing about age or people, it knows nothing except atoms and their motions. If that is ground truth for you then the wave equation is what you want. If what you want to know is about atoms and their motions in this room that is what you want. But if you sense that the wave equation is throwing certain things overboard that are important then we have to come back and ask that question, what is it that we have thrown overboard? What is it that is of importance that we have lost when we do the wave equation?

I would say, I would even be sceptical about the predictive power of the wave equation for atoms in the room. If you want to know where the atoms in my body

are going to be tomorrow you do not want to write a wave equation you want to look at this little device here because my calendar is in here. That would be much better, that would have higher predictive value for you than any physics thing you can apply. Now and here we write, I am not saying it is wrong, it's not wrong that would be a stupid thing to say. But it is simply impoverished and when that reductionist fundamentalism leads us to that level if we follow that that program. Those are my problems with bottom-up determinism.

The theory that quantum mechanics and quantum and deterministic leaves place in the brain for non-causal functions and freedom of will to enter is also problematic to me. It is problematic for two reasons, one of them is scientific, one of them is intuitive. The scientific reason which you should probably take more seriously is that I have talked to many biophysicists over the last several years trying to get a handle on this and trying to find out whether they can, that's physicists themselves, people who know a lot more about physics than I do. Whether they can see and people who know about synapses and the fundamental electrical properties of nerve cells. To find out whether there is a way for quantum and determinacy to play a serious role in synapse transmission or electrical signalling of the brain. The answer I have gotten from every single one I have talked to has been no. The notion is that the ion cells that are the fundamental units that open in the pore to allow electrical current to flow into the cell and out of the cell and they are fundamental to synapses and they are fundamental to action potentials, these ion channels are large protein molecules and they are too large for quantum events to be a basic organising principle. So there is noise in the brain but the noise is thermal noise, it is heat noise, it is not quantum mechanical noise. Now there are ways that quantum mechanics influences the nervous system. So vision is a good example. The absorption of little light photons is like the absorption of any photons, quantum, mechanical, probabilistic event. The absorption of gamma radiation by cells in my skin that can lead to skin cancer and change my life are quantum chemical events. But as far as I know, as far as I have been able to determine, the macro molecules that govern normal signal, transit normal signalling mechanisms in the nervous system are too large, those molecules are too large for quantum mechanical events to have, to play a serious role in it. That is my scientific reason for doubting that quantum mechanics offers us any way to think about the mind-brain problem I have got free will.

The personal reason is that quantum mechanics is fundamentally probabilistic. It says there are things that are probabilistic to happen that tend not in principle to be predicted. And I do not like that idea too much. When I go out to cross the street afterwards and I look both ways to see a car, and I better look both ways in this country because I never know which way I am supposed to look. When I cross that street I want a mechanism in my brain that detects those cars reliably every single time. I do not want some quantum chemical thing that say this time you missed it dude because that is going to be the end of my life. I do not want to be subjected to quantum mechanical randomness or randomness of any kind if I can help it. I want the mechanisms in my brain to be robust for that kind of randomness as well as quantum mechanical so the control systems of the brain are built to damp out on this to some extent. That is a scientific reason for not being real fond of quantum mechanics.

The intuitive reason is the one that I just gave you and I have a couple of quotes here from Pat Churchland that I just communicated the gist of those to you so we will just skip that.

I have already said that I am not really attracted to this freedom, it was uncaused thing. A lot of it is because of this central dogma of neuroscience that our behaviour is inextricably linked to the brain and there are causes inside the brain. So I am not wedded to that and what we really want to get to is this self determination or autonomy. Well what the hell does that mean? How can we possibly think about that?

Here is where I just want to leave some thoughts about possible ways of thinking about this in the future. The key issue in my mind, this is critical, this goes to the heart of the matter, is what counts as a cause? Are causes only legitimate to think about at a fundamental level in terms of forces of physics or are there that causes, legitimate ways to think about cause throughout a hierarchy of mechanisms? This is something that I take just the obvious, I do not take this as a controversial statement and it is not an answer in and of itself, but I take it as obvious that wholes are more than sum of their parts. They are more than the sum of the parts by virtue of the organisation of causally interacting components. Wholes have causal powers that parts do not. I have a causal power to kill somebody in this audience, that is real causal power. A lion has the causal power to kill any of us. A table has the casual ability to support the weight of any amount of food we want to put on it.

Those are real causal powers and they are causal powers that belong to the organised system, not to any parts. So you can take my body apart and do into organs or into its cells and it loses the causal power. The organisation is the critical thing that endows causal power. It is the parts and their organisation.

This kind of organisation is important and I do not see it emphasised enough in this literature. Organisation is actually critical to me. It is the logical sequence. If you think about any kind of piece of music that you are familiar with like, and I know this is not a culturally appropriate choice, but like the Star-Spangled Banner. The Star-spangled Banner can be realised at any given instant in time, it can be played on the violin, it can be played on the piano, it can be played on a horn, it can be represented as notes on a sheet of music, it can be represented as electrical disk files on some disk somewhere, it has a physical instantiation. But the Star-Spangled Banner when you come down to it, it is not its actual physical instantiation on any piece of paper or on any piece of music or on anybody's memory, it is a logical system of organisation of notes and timing between the notes. If you want to get to the essence of what the Star-Spangled Banner is I think it is that logical organisation. This organisation, that organisation is very, very difficult to reduce the physics or to ensure it in your wave equation. That organisation is our higher level entity that has causal powers because the Star-Spangled Banner written on a piece of music can guide the arms and motions of lots and lots of musicians playing in an orchestra. So it has causal power, it can orchestrate the musicians rather than the musicians orchestrating it but that is the kind of causal power that comes into level when you have organised complex systems. So this organisation, and in the abstract what that really comes down to is information is reducible in individual instances. So at any one moment in time it might be reducible to either notes this person is playing on the piano, or notes this person is playing on the horn or this sheet music in front of me, but not in the general case. In the general case the Star-Spangled Banner is information, it is a relationship of notes and timing.

So fundamental laws like in physics constrain but do not determine the behaviour of complex systems. So this programme that is running on my Macintosh right now is PowerPoint and that programme is orchestrating the activity of lots of transistors and things inside that computer. Now that programme cannot make the computer do something that is fundamentally against the laws of physics, the currents flowing on those circuits have to obey Kirchoff's laws, they have to obey all the principles of

electronic circuits so the fundamental laws of physics constrain what this computer do but they do not determine what the computer do. What determines that the computer is doing is at any given moment in time is the software that is running in the computer, it is a higher level organisation that can orchestrate its lower level parts, like Star-Spangled Banner itself can orchestrate lower level parts. This is like much more similar to what neuroscientists really do in explaining the brain. I would argue that neuroscientific explanation is intrinsically multi-level like that.

This is laid out in this book that I have alluded to already by Carl Craver, this philosopher who also did a masters degree in neuroscience in the middle of his PhD in philosophy of science. Craver argues cogently to me that multilevel mechanistic explanation is characteristic of neuroscience. These are a couple of quotes. "The systems tradition construes explanation as a matter of decomposing systems into their parts and showing how those parts are organised together in such a way as to exhibit the explanandum phenomenon." This is what we do as neuroscience. If want to study intention we start decomposing the brain system and the parts and trying to understand how the neural circuits in the brain exhibit the explanandum potential behaviour. "Systems explanations involve showing how something works rather than showing that its behaviour can be derived from more fundamental laws." This is critical. This is a different notion of reduction. This is not the current law model where we are trying to take theories at some level and reduce them to theories at another level and then show that we can derive a higher level from the lower level. That kind of reduction is an argument, this kind of reduction is what scientists do everyday which is to discover the mechanisms. It is not an argument about derivation but it is about showing how something works. Craver is a philosopher but he is a philosopher who is very congenial to scientific points of view in my mind. He has this concept that I find very clever and again I find it very consistent with what I do in the laboratory, this concept of mutual manipulability. So "Mutual manipulability," he says "a part is a component in a mechanism if one can change the behaviour of the mechanism as a whole by intervening to change the component" So that's a bottom-up kind of intervention, you change the component and you change the behaviour, "and one can change the behaviour of the competent by intervening to change the behaviour of the mechanism as a whole." So he is saying the causality runs in two directions here. I find this very congenial.

We can change the behaviour of an animal in a laboratory, have them to do one intentional task versus another intentional task but pay attention to this part of space or we can train them to pay attention to this part of the space and it changes the components, the components inside that are mechanistically linked to the intentional behaviour. But we can also go in with microelectrodes and we can electrically stimulate structures in the brain that we believe are involved in intention and we can change the animal's intentional behaviour. So we can go back and forth between these levels in a bidirectional way. And Craver summarises this in something that should be worn on a scientist's heart which is "making a difference". Its mechanisms, you should know your own mechanism when you can make a difference, when you can manipulate something and actually make a difference and those currents of manipulations run in multiple directions up and down explanatory hierarchies.

Here is the example that he favours and I think it is a pretty good one. He chose the phenomena of long-term spatial memory. At this point it is probably the best example we have of a cognitive phenomena that is as broad as outlines at least are understood at the neurobiological level. We trace this down through several levels here, Craver traces, I have added on one to his. At the beginning when we study long-term spatial memory it is frequently studied in mice who are navigating a water maze. So a water maze is a pool of water that is made opaque by pouring milk or some kind of food colouring in it and the mouse just swims around in a pool of water until he finds a little pedestal that is submerged underneath the water. He finds the pedestal and he gets up on it and he rests and after doing this for several times the mouse when you drop him into the thing he just goes straight to the pedestal. He has learned, he has a long-term spatial memory about where that pedestal is. And this is a classic behavioural thing for us to study in the lab. Now we know that underlying the mouse navigating that there is a structure called the hippocampus and that hippocampus generates a spatial map and we can talk to you a lot about hippocampus and the kinds of cells we find there and why we think it is involved in a spatial map but you would still want to know, well how does the mouse learn this particular thing?

We noticed some things about the next lower level, about mechanisms underlying that learning, that there was connections between cells called synapses and these synapses change strength. A particular experimental paradigm that is studied a lot

by neuroscientists was changing strength, that is the deep mechanism underlying memory is called long-term potentiation, LTP for short, it is just a way for cells to change the strength and synaptic connections behind them. We know a lot about how LTP is actually accomplished at molecular level, we know about these different receptors that are critical for the LTP phenomena and then deep at the bottom we have genes that actually produce more of these receptor molecules that get inserted into the synapse that actual result in the synapse having a stronger connection, or a weaker connection between cells.

So here we have got at least five levels and as neuroscientists we go up and down this hierarchy all the time and we manipulate at this level. We teach the mouse a new paradigm or something and we can show changes at these levels. We can manipulate things at the gene level down here. We have knock out mice where we can knock out certain genes that we believe are important for LTP and we can show that we have behavioural deficits when we knock out those particular genes and they are selected for learning. We started at the top and we looked downward for mechanisms but many times we start in the middle somewhere like LTP. LTP has been studied now for 20 or 30 years in neuroscience and we are still trying to figure out whether it is actually related and exactly how it is related to learning behaviours up at the top. So we are looking up as well as looking down and we are going up and down this hierarchy all the time. This feels right to me, this mutual manipulability thing feels right to me and the causes come down and causes go up. We can cause changes down here by doing manipulations up here and we can cause changes up here by doing the manipulations at the bottom.

So what does all this mean? I lost track of what is at stake here. And if I lost track I am sure some of you lost track too. Like I said I think the key issue here for neuroscientific explanation, for understanding of people and behaviour and how to think about ourselves as biological beings, those beings with some autonomy is this issue of self-determination and autonomy and responsibility within the law for those of us who are religious, responsibility before God and the key issue is what counts as a cause? That is the really deep thing.

If we can find a way to talk meaningfully about non-fundamental causation so resisting that natural tendency to tumble down the causal chain and believe somehow that when we arrive at the bottom that is where we have got truth, the fundamentals is real truth. If we can start thinking, because I think neuroscientists

do implicitly, that causes lie at all those levels for any particular phenomenon. Then we can take mental causation seriously and mental causation is the key to taking responsibility seriously.

Now this is not to say, I want you to hear me very clearly on this, this is not to say that bottom-up causes are unimportant; explanatory relevance runs in both directions. So there is nothing that I do as a person that I could not do without my body and my body instantiates certain laws and biology and chemistry that constrain the freedom that I have. I am never going to have the freedom to play basketball like Libron James, probably another culturally inappropriate reference. I do not have that freedom but the physics and biology in my body do not permit that. But there are some kinds of freedoms that they do. So explanatory relevance runs both upward and downward.

I will use this quote here from Pat, she says "The absolutely crucial point is that not all kinds of causes are consistent with free choice." Now again I don't like that word free, I would rather have self-determination or autonomy in there. "Not all kinds of causes are equal before the tribunal of responsibility. Some causes excuse us from culpability; others make us culpable because they are part of the story of voluntary action. The important question is what are the relevant differences among the causes of behaviour such that some kinds play a role in free choice" or autonomous choice "and others play a role in forced choices." So I think she is onto something.

Here is another example of this bottom-up and top-down thing. Richard Dawkins is one of the fundamentalists from a generic side. I do not know what he does about the problem of that his fundamental level has other fundamental levels. But here is the way he describes humans and their relationship to the genes inside them. "Genes swarm in huge colonies, safe inside gigantic lumbering robots," that is us by the way "sealed off from the outside world, communicating with it by tortuous indirect routes, manipulating it by remote control. They are in you and me; they created us, body and mind; and their preservation is the ultimate rationale for our existence." That's a very bottom-up looking down, the genetic fundamentalist.

Now here is another biologist down there at Oxford, that un-nameable place down the road, Denis Noble who wrote *Music of Life: Biology Beyond the Geneome*. Noble looks at genes differently. I am sure this is a parody on Dawkins, I think it is, "Genes are trapped in huge colonies, they are locked inside highly intelligent beings,

they are moulded by the outside world, communicating by complex processes through which blindly as if by magic function emerges. They are in you and me, we are the system that allows their code to be read and their preservation is totally dependent on the joy that we experience in reproducing ourselves. We are the ultimate rationale for their existence.

These are two very different ways of looking at this relationship and I think they are both right, they both have some truth. I am not here to argue primarily for bottom-up or top-down but to day both are important, I just want room for top-down.

Here is another example, most of you have not heard of the Lasker Awards. These are rewards that are given out for biomedical research in the United States every year that are considered a pre-Nobel prize for biological and clinical research in medicine. And it really caught my eye a few years ago when a psychiatrist, this man Aaron Beck here was chosen to receive a Lasker Award. Aaron Beck is the guy who invented cognitive behavioural therapy which has been adopted enthusiastically by the Euro National Health Service system here in this country. The reason why Aaron Beck won this prize is that scientific studies have shown (I have the references if you want them,) that the treating of really important cases of depression, cognitive behavioural therapy in combination with drugs, with meds, with anti-depressants is more effective than either alone. That is an important observation. Because psychiatry used to be about therapy but in the 1980s, 1970s it is biological revolution where we treat mental illness with drugs. Psychiatrists became little more than prescription writers. A model of a guy in my lab who is in psychiatric training to be a psychiatrist, his model was leave no receptor unoccupied. That was his answer. But here we come along into the 2000s and we are demonstrating that the therapeutic approaches in culmination with the meds are more effective than either alone. Now this is bottom-up and top-down. The bottom-up intervention is giving the prescription, giving the drug, you modify the neuro-transmitters and you modify the receptors very deliberately and very intentionally you are going in there and mucking with the low level mechanisms and it has real positive effective. But what we know now is that top-down intervention is part of the equation to a real treatment. And in cognitive behavioural therapy what you really try to do is change the patient's beliefs about the world and their patterns of interaction with the world and that you call this cognitive restructuring. You are going into the system that is a human being and you are operating at a very high level, you are trying to change

the beliefs and in short, beliefs matter. Beliefs have strong causal influences, they are very powerful things in our culture and in science.

Finally a guy who knew this very well, this systems approach to causality is this guy, Leo Tolstoy and he wrote this book which some of you have read, War and Peace and in one of the climactic scenes in War and Peace the protagonist Pierre facing a firing squad in Moscow. He's sitting there, all these vile thoughts going through his mind before he is about to be shot and he is trying to think who was it that had actually sentenced him to death? He was trying to find the locus of causality here. And he was convinced it was not those guys in front of him holding a gun because they did not know him, probably they did not want to be there, they just as soon be doing something else than shooting him and he found it hard to hold the locus of causality in the fingers of the men who were going to pull the triggers. He was trying to think through the long chain of events that got him to this point and what he concluded, what Tolstoy writes is that it was a system, the causality was in a system. It was the apparent organisation that was set up, it was not any little piece of it, it was not any little sub-component that mechanism, it was a system, it was this concurrence of circumstances.

I would conclude here by saying that understanding the nature of human freedom or human autonomy if we want to use that word, is the most important problem facing the neural behavioural scientists. Others are up there. Finding a cure a Alzheimer's is up there. But I think the one that has the most long-term implications for our society and for our understanding of ourselves is this problem right here, understanding the nature of human autonomy.

Here is the last little twist I am going to put in here. I have talked about how this is important for religious understanding of humans and our responsibility, I have talked about how it is important for our legal system. But here is the thing, that it is also important for science itself. I do not see how we can take science seriously unless we can find some way to think creatively among non-fundamental causes in our time. I illustrate this with this quote from JBS Haldane who was a mid century British geneticist and a very famous one, Haldane says "If my mental processes are determined wholly by the motions of the atoms in my brain," that's the fundamentalist point of view "I have no reason to suppose that by beliefs are true and hence I have no reason for supposing my brain to be composed of atoms."

His point is, and I think it is a good point, how do we even think about truth, we scientists who pride ourselves on searching for truth in dealing the fundamental reality of things, how can we even think about truth if everything we are doing is bottom-up determined about blind forces and notions of atoms. And that is the kind of thought that I want to leave you with and leave you with this notion that for neuroscience and for us integrating neuroscience into our view of who we are as people, as people within society with each other, living in community with each other and living as responsible beings to each other and to God, this is an important problem to solve. And so I would just close by, well that was the close, that was punch line. That this is an incredibly important problem to solve that I do not as I said pack solutions but maybe some of these ways of thinking are ways that we can use going forward. So I will stop there and we will have some question and answers but thank you very much for your attention and for your patience.

Question and Answer Session

Newsome: Now for some of you who know a lot more about some of this than I do, let us hear your points and counter points and questions.

Moran: Cannot you argue that all the elements of the system can be formalised with equations or even feedback loops that link so the outcome of the external cause to the minor event whatever it is, and all of that gets formalised and ultimately it is just our ignorance that means we do not know how to formalise them and it is fallacious to be as polite as possible to quote God in place of that ignorance?

Newsome: I certainly think it is fallacious to put God in place of ignorance. I hope that nothing I have said today comes across as using God for an explanation for something that is probably a scientific explanation. I am explicitly not doing that. I do not use God to explain scientific events. I need God for meaning in my life. So let us make that very clear upfront.

But your main point that you made about higher levels of interaction such as feedback loops, I think those are incredibly important. And feedback loops is one of those things that I do not think can be reduced to fundamental laws of physics ultimately. We cannot write the Schrodinger wave equation that captures the meaning of the very

term feedback loop. So feedback loop is one of those elements of organisation. We have tons of feedback in the nervous system all up and down those causal chains that we go on all the time. And that is an element of the organisation that makes the hole more than some of the parts and gives its cause, gives the hole causal powers that parts individually do not have. So those feedback loops are incredibly important and we can write equations at their level. But when we try to go down level after level after level in search of fundamental truth at the bottom, we lose the whole concept of feedback I think. It is just atoms and their motions. So somehow, strange as it seems we have to find ways to take that from loops and their quantitative expression in equations and what they do. We have to try and take them seriously and give them and grant them the causal status that they deserve. That is a key part of my whole argument here.

Q2: Hi, Barbara Sahakian, very nice talk, thank you very much. I wanted to ask you if you would not mind commenting on the very interesting work of John-Dylan Haynes with the FMRI experiments, I am very interested in your ethics. It is quite a challenge because as you know, there is this experiment that has been picked up in the press where they were able to demonstrate that when they asked people to have an intention as to whether they were going to add or subtract a couple of numbers, so it is a rather kind of indication that one can imagine would go to a much more important view of intention. They were actually able to determine using the FMRI procedure and obviously some nice modelling. Whether what the intention was in advance. Do you think that puts a challenge to some of the issues that you raised and I just wondered if you could comment?

Newsome: yeah thank you for the question. So for those of you who do not know this, there is this tradition or body of literature of contemporary neuroscience that was really started in some ways by a neurosurgeon at University of California, San Francisco named Benjamin Libet and has been recently followed up by John-Dylan Haynes using FMRI technology. And Libet did a simple experiment that is probably the best known neuroscience experiment among philosophers in the world.

It has a close rival now with mirror neurones. But Libet's experiment was the following, he asked human subjects to sit under voluntary control and decide to move their fingers, with no cues, no instruction just under pure voluntary control. There was this little clock sitting here with the second hand ground around the clock and he asked the subjects to just move their finger at some point in the next sixty seconds and then report the time that they decided to move their fingers. So the subject moved their finger and they spoke verbally at the time that they decided to. And what Libet found doing EEG recordings from the head is that there is this potential well known to neuro psychologists called readiness potential, I forget the term and name of it that actually preceded the subject's identified time of the conscious decision by a few hundred milliseconds. John-Dylan Haynes in some of his recent literature says that he can predict some of these things minutes ahead of time, which he better be able to due if he is using FMRI because FMRI does not have the time resolution.

The implications of this the philosophers care about is that what it means is that consciousness, the conscious decision by the human subject does not have causal power about when to move the hand, that the brain has already started working on this decision and the brain has already made this decision because these brain events precede the conscious report by several hundred milliseconds and the notion is in these that the consciousness is along as a rider on top of the casual elements inside the brain. That is why it is so much discussed in philosophical circumstance.

I would say two things about that. Number one in Libet's experiments, there are things that do not satisfy me yet in their modern reincarnations do not satisfy me yet that they are procedurally rigorous. So before a subject actually looks at the clock and reports the time in their finger they actually have to reorder attention, pay attention to the clock and that certainly takes time, we know that from psychological states. And there are other issues like that. But it may be that even once you do the best controlled experiments possible that this observation stands. And if this observation stands it seems to say

at least for some kinds of choices the brain is already working on that before you are consciously aware. In principle do I have a problem with that? Maybe, maybe not.

Some of the decisions that I make I am sure my brain is working on before I become consciously aware of it. And I do not take that as damning in view of consciousness. Is consciousness always coming along later, several hundred milliseconds after a bunch of neural events? I don't know that. And I do not know, you will notice in this talk I carefully avoided using the word consciousness because I do not understand consciousness very well and I do not understand how consciousness systems of brain can become conscious health, bunches of neurones can be come conscious. I once talked with Dan Schechter, the famous psychologist of memory just exactly about this notion, he said. "I do not care if I arrive on the scene a little late as long as I get to play once I am there."

So I think these are things to be taken seriously and I think they are things that can be done better on the experiments I think it is a potentially important observation and I think it may be that a lot of what we think about this conscious experience is that the nervous system is arriving at that before we do. It still can fit under autonomy and mechanisms the way that I have laid them out, the way that Carl Craven does, for any of Craven's arguments and either of these things I have said about autonomy it is important that the decisions have come out of my brain be based on my own experience and on my own values. It may not be necessarily that I have been conscious of it before the decision actually gets made. So I have deliberately let that consciousness thing ambiguous, and that is a little bit of a dodge of your question but it is where I feel like I am, so I am being honest. Do you have, what do you think? Do you have? This is my thing, there are people in this audience who have thought a lot about these things.

Q2: Well I have thought quite a lot about it. At the moment as you say obviously we cannot do it with the accuracy at the moment that we would really like to see and also the rather trivial intentions as you

pointed out. But it does set the scene for what might happen in the future [it does] and it is, I agree with you it is rather challenging.

Newsome: Yes. I can tell you one thing about Hayne's experience that is procedurally or somewhat problematic. When you choose things, even if you ask a person to choose randomly between two buttons for example it is very hard to generate random numbers and what you actually do is fall into patterns. If you fall into a pattern it means that the current choice, like if you are going one, two, three on one button and one, two three on another button you are going two and two, what that means is that my choice on this current trial depends somewhat on the choice that I made on the previous trial. And that choice on the previous trial there is a record of it inside the brain, we know that. So if your choices are not random it makes a lot of sense that they are going to be variables inside the brain that are correlated with the choice you are making on the upcoming trial. Those kinds of correlations are not necessary deterministic, but they are in your brain. So all this kind of thing, this needs to be thought about, reasoned about very carefully and it requires real serious hardcore experiments but they are important ones. This is one of my worst faults as a speaker, is long winded answers to a fairly simple question.

Q3: I would like to know how what you have been saying interacts with some traditional religious commitments? So you have had this dogma of neuroscience where the sense of self and all that goes with that inextricably linked to the action of the brain.

Newsome: The biology of the brain right.

Q3: But then I guess there is the commitment that maybe a person dies or you die, the self continues? Or also that there is some kind of personal self in some sense, maybe God who doesn't rely, or is not inextricably linked with a particular brain?

Newsome: And there is also the issue of a soul, traditional religious understandings are that there is a soul and in certain traditions especially medieval Christianity the soul was conceived as something separate from the body. What do we do what that if all our minds are

linked the brain? And what I do with that is that I still resonate to the religious concept of the soul in the sense that there is something central at the core of my being, what I think it is it is not like a non-material ghost lying in there just waiting to get out. But the point is it is the organisation. I am literally the sum of my experiences and the sum of my aspirations and the sum of my background – I need careful about that word sum but it is the organisation. It is really the organisation.

So let us do this thought experiment, if I could take one of the neurons out of my brain and replace it with a piece of silicone, a silicon shield that mimicked all of the connections that that neuron had with all the other connections and mimicked its input, output functions I would still be me, I wouldn't recognise it. I would have to find a way to power the silicon chip. Then you do that experiment ad infinitum where I have gotten all the neurons out of my brain and it is just replaced with all these chips, am I still me? I think I am. And what has happened is I have replaced all the parts but I have maintained that organisation and the information that is stored in that organisation. So I think just like the Star-Spangled Banner that a whole bunch of different physical instantiations and just like PowerPoint and run on a whole bunch of different computers I think that there is going to be logical reason why the stuff that is really me there cannot be transferred to some other kind of medium. The futurists raise this all the time, that one day we are going to be free of these bodies, that intelligence is going to be downloadable into a different medium that is much more durable than these kinds of bodies. It may be still be recognisably you. So I would change the language around, I would not use soul in the mediaeval sense but oddly, this sense and the word soul where it is deeply embodied in the connections of your brain to the rest of your body that is more fatal to traditional Hebrew understanding of personhood than to the medieval scholastic Christian understanding. So I think that these things are ultimately straightforward, thank God.

MS: Can we just make this the last question?

Q4: Just thinking about what you have just said about the soul, does it not follow inexorably - I am not trying to be controversial - When the body dies the soul must?

Newsome: Unless that information is transferred to some other medium which is the Christian doctrine of resurrection. If it does, if I believe that if all the extant copies of the Star-spangled Banner were to disappear, it gets removed from all sheet music, for all computers, from all people's memories, it has no physical substantiation then it goes away. It has to have some kind of substantiation in order to continue to exist. I think it is more than a single substantiation but it is pitted on some substantiation. So if we are to exist after death if there is anything beyond these lives that we do not know, something about that organisation that is really us has to get passed to, who am I say to passed to, one way to think about it would be to say that the organisation that is really us gets translated to another kind of being. I do not think that is inconsistent with the Pauline notion of spiritual and body for example. So I think you can think it up that way. In fact if you want to push the metaphors, it is dangerous to push metaphors but if you want to push the metaphor you can think about resurrection, the promise of resurrection in the Christian tradition is not that it will have you reproduced exactly like we are now. But that we are actually going to be in a different and better state than we are now. And the Star-Spangled Banner or any other music, can be moved from a minor key to a major key and yet still recognisably be the same piece of music. So maybe one way to hope about this is that you, the organisation which is recognisably you, move from this minor key that you are in now to a major key - that's resurrection. Now some of you like minor keys so that (laughter) let us not be too chauvinistic about this but that is just a way to think about it.