



Presidential Commission  
*for the* Study of Bioethical Issues

## TRANSCRIPT

### **Joshua D. Greene, Ph.D.**

John and Ruth Hazel Associate Professor of Psychology  
Director, Moral Cognition Lab  
Harvard University

### **Alfred R. Mele, Ph.D.**

William H. and Lucyle T. Werkmeister Professor of Philosophy  
Florida State University

Meeting 17, Opening Remarks and Session 6  
June 10, 2014  
Atlanta, GA

DR. GUTMANN: Good morning, everybody, and welcome back. I am Amy Gutmann. I am President of the University of Pennsylvania and Chair of the President's Commission for the Study of Bioethical Issues. On behalf of myself and our Vice-Chair, Jim Wagner, who is the President of Emory University and has been a terrific host to us for these two days. I welcome everyone to the second day of our 17th meeting.

Before we continue I want to note the presence of and welcome our designated federal official, Bioethics Commission Executive Director Lisa Lee. Lisa, please stand up. Thank you. We had a very productive day yesterday and we'll continue our discussion of neuroscience and ethics today.

For those of you who weren't with us yesterday, let me just summarize how we take comments. Every member of our commission staff has comment cards. There are also comment cards at the reception desk. And if you want to ask a question or make a comment, please write it down, hand it to a member of our bioethics staff, and they will make sure it gets up to us. If we don't have time to read and answer the questions, we will do so after the meeting.

Will members of the staff just raise their hands, wave their cards. Great. Jim.

DR. WAGNER: Nothing to add. Let's get to work.

DR. GUTMANN: Okay. We are going to get started. Let's ask our panel members for the first panel which is on potential implications of advances in neuroscience research for ethics and moral decision making to come on up. Thank you. So we're going to begin this morning with a panel devoted to the potential impact of neuroscience research on our understanding of ethics and moral decision making. For the most part we have focused our attention thus far on the ethics of neuroscience, the ethics of doing neuroscience research. This panel will flip things

in a sense and focus on the neuroscience of ethics, that is, how if at all does neuroscience change, alter, revolutionize possibly, some people have claimed, our understanding of ethics.

We'll start with Dr. Joshua D. Greene, who is the John and Ruth Hazel Associate Professor of Psychology, a member of the Center for Brain Science faculty, and Director of the Moral Cognition Lab at Harvard University. In just a few weeks he will have a new title: Professor of Psychology. Congratulations.

Dr. Greene studies the psychology and neuroscience of morality, focusing on the interplay between emotion and reasoning in moral judgment. His publications have appeared in "Science", "Nature", and the Proceedings of the National Academy of Science. In 2012 Dr. Greene was awarded the Stanton Prize by the Society for Philosophy and Psychology, and in 2013 he received Harvard's Roslyn Abramson Award for teaching. Welcome.

DR. GREENE: First I just want to say it's a great honor and a pleasure to be here, so thank you very much for having me. There's a lot of excitement about the possibility of neuroscience changing the way we think about issues of right and wrong, and I think that a lot of that excitement is justified. I also think that wherever there's a good reason to be excited, there's always additional unwarranted hype that comes along, and I think that that's true in this case as well.

I'm going to briefly mention I think two general ways and give an example of each way that I think that neuroscience can have an effect on our thinking about right and wrong, and then if there's time, say a few words about how I see neuroscience really as just one piece of a larger story about behavioral science and that we shouldn't get too excited about neuroscience at the expense of other kinds of relevant research.

So the first way I think that neuroscience can be relevant and the more straightforward way is that neuroscience, like other research on human behavior and decision making, can help us understand the processes by which we make moral decisions. And once you understand a process, you understand the strengths and the limitations of that decision-making process. My favorite example for thinking about how human decision making works in general is the example of a digital SLR camera. So the camera I got many years ago, and you probably have one like this, it has two ways of taking photographs. It has automatic settings so you can put it in portrait mode or landscape mode, and those are great, quick, point and shoot, and most of the time you get what you want. The camera also has a manual mode where you can adjust all of the settings by hand.

And you might ask, well, which is better, automatic settings or manual mode, and the answer of course clear in this case is it's not that one is inherently better than the other; they're good for different things. The automatic settings are very efficient, point and shoot, but they're not very flexible; and the manual mode is very flexible but it's not very efficient. And so you want to use them for different things. If you're doing something straightforward, something that the manufacturer had in mind, you're probably better off using your automatic settings; and if you're doing something more complicated, something perhaps that the manufacturer didn't anticipate in detail, that's when you would put the camera in manual mode.

And our automatic settings in making decisions are our gut reactions essentially, our intuitions and often emotional intuitions about things, and they are generally very good most of the time but they are not necessarily good for dealing with complicated, unfamiliar problems. And bioethics is perhaps the best example of a domain in which we're dealing with complicated, unfamiliar problems, unfamiliar in the sense not that people haven't heard of things like abortion,

but unfamiliar in the sense that we don't have a long history of genetic and cultural and personal trial-and-error experience in dealing with things where trial and error is essentially necessary for training one's intuitions.

So let me give you an example I think relevant to bioethics of how I think neuroscience specifically can shed light on our automatic settings and their strengths and limitations. So one thing that people in philosophy and public health and other domains have noticed is that people seem to be strangely insensitive to quantities when you're dealing with large numbers of lives. So you can save one life, that's really important; two lives, that's really important; get up to fifty lives, a hundred lives, a thousand lives, it all just starts to seem like a lot. And from a certain perspective, from I would say a manual mode perspective, this doesn't necessarily make a lot of sense. Why should the 5,000<sup>th</sup> life that you can save be worth any less than the first life that you can save? And yet intuitively people have found time and time again this kind of diminishing returns of saving lives. Why should we be that way?

Let me describe an experiment done with Amitai Shenhav that perhaps can shed some light on this. So we had people make decisions where you can save one person's life for sure. Let's say you're on a rescue boat going to save that person and then you get a message that says you can go in this other direction and you can save 10 people, 20 people, 30 people, with a certain probability of actually succeeding and saving them. And here too we found that as the numbers go up, they matter, but they matter less and less the more numbers that you have. And the purpose of our experiment was to figure **out** which parts of the brain are keeping track of the number of lives that you can save, the probability of saving them, and then putting those two pieces of information together.

Since we are short on time, I won't go into the neuroanatomical details, but the gist of the results, at least for present purposes, is that the system that seems to be keeping track of these variables is the same basic mammalian primate reward system that keeps track of the values of goods for rats and for monkeys. Obviously it's a bit different in humans but it's the same core system. And if you think about it, if you are a monkey, you're deciding between the sure thing you can eat here or the good fruit on the other side that maybe you'll get to, maybe you won't. You want to keep track of those variables. But if you are a monkey, all the goods in your life, they level off, right? You know, you don't have a fridge. You're not going to be able to save this up for the future.

For almost anything in most mammals' lives things level off pretty quickly, but then we humans do these weird things like save people on the other side of the world who we are never going to meet, by the thousands, if things go well. So part of the explanation for why we seem to be insensitive to quantity, as psychologists say, is that we're basically, to put it in the crudest terms, apologies to my neuroscience colleagues, using our monkey brains, using our rat brains, using our automatic settings that are designed for this kind of thing to give us gut feelings about should I go there or should I go there, to think about things like saving large numbers of lives.

Now, you didn't need a neuroscientist to tell you that we have this tendency, and maybe the explanation that I gave is wrong, but maybe it's right. And if it's right, it gives us some insight into why we have this, in a sense, bizarre tendency; and two, why we should perhaps regard it as a problem, why we should think a little more wonkily, if I can coin that word about these things, instead of relying on our gut reaction about what's the best thing to do in cases involving large numbers.

The more general point, and this connects more to Professor Mele's work, is that understanding more generally that the brain is a mechanistic system, that everything that determines our behavior is ultimately a physical process, if that's true, and I think people have speculated about this since antiquity and more so during the enlightenment, but now it's increasingly clear that everything that affects your behavior is ultimately just one neuron making another neuron fire. It's one physical process going all the way back in time back to before you were born.

I think there are a lot of ways in which understanding this doesn't really change our understanding of human behavior. People still make decisions. People still have things they want. People still have things they believe. People are still sensitive to rewards and punishments, and we still have the experience of choosing and doing something by choosing and thinking about what you want to do and then going and acting and planning and doing it. It's different from a knee-jerk response or the kind of thinking that a child might exhibit or someone who is mentally incapacitated in some ways.

But I also think that understanding that the brain is a purely physical thing, that we are people, we make choices, we make decisions, but we are also ultimately physical mechanisms, I think that that does matter. And this is something that I and Jonathan Cohen speculated about ten years ago in a paper. More recently I and other colleagues are just about to publish a paper showing that at least one of the predictions we made seems to be right, which is that if you just expose people to neuroscientific information that gives the people the idea that we're ultimately physical systems either by having them read some newspaper articles about a recent brain imaging study or having them take a cognitive neuroscience class, people become less punitive and more specifically less retributive, that is, seeing punishment just as a good thing in and of

itself. Even if you think we're mechanistic systems, it still makes sense to punish people and hold people responsible, because that's essential for a well-functioning society. But it's also possible to punish just because you think it's inherently good that people who do bad things suffer. This for example was Kant's view and a lot of legal scholars take this view as well.

And our sense, I can give you more details but we are short on time, is that understanding that we are ultimately physical systems determined by forces that are ultimately beyond our control, even if we have immediate control of our behavior, it doesn't make us think, oh, no one makes any choices, no one is responsible at all for what they do, but the idea of punishment as an end in itself I think loses its grip on people.

So to summarize, I think that neuroscience, part of the context of broader behavioral science, can help us understand how we make decisions and the strengths and limitations of those decision-making processes. And understanding that our brains are ultimately physical systems I think can change in some ways who we think we are and how it makes sense to respond to us when we behave in ways that are bad for society.

DR. GUTMANN: Thank you very much. We will come back with questions and you'll have a chance to flesh out some things, but first let's move on to Dr. Alfred Mele. Dr. Mele is the William H. and Lucyle T. Werkmeister Professor of Philosophy at Florida State University. He is the Director of the Philosophy and Science of Self Control Project and Past Director of the Big Questions in Free Will Project. I love the names of these projects. Dr. Mele is the author of ten books including most recently *A Dialogue on Free Will and Science*, and *Free - Why Science Hasn't Disproved Free Will*. His 2009 book, *Effective Intentions*, won the American Philosophical Association's Sanders Book Prize. Welcome.

DR. MELE: Thank you, and thanks for having me. I just got back from a four-day party for my dad's 90th birthday. The old people wore me out.

If I had a title for this talk it would be Let's Be Cautious. What got me interested in this topic really was reading the barrage of news reports that were claiming that scientists, especially neuroscientists, had shown that there was no free will, and I wondered what effect it had on people. Let me just read you a couple of quotations from the news. This is from an article called "Case Closed for Free Will" in Science Now Daily News, 2008: "Your mind might be made up before you know it. Researchers have found patterns of brain activity that predict people's decisions up to ten seconds before they are aware they've made a choice. The result was hard for some to stomach because it suggested that the unconscious brain calls the shots making free will an illusory afterthought."

Here's another one also from 2008. It's called "The Decider." It was from Science News: "Free will is not the defining feature of humanness, modern neuroscience implies, but is rather an illusion that endures only because biochemical complexity conceals the mechanisms of decision making."

So I wondered: How does this affect people? Maybe some people just don't believe it. Maybe some do believe it and are depressed. Who knows. So I had been tracking the scientific work and I didn't believe that it showed these kind of things, so at that point I thought I'd just throw myself into the ring. I want to read you an e-mail I got when I was writing my book *Effective Intentions*. This just came out of the blue. It's sad and cute at the same time, I think. So she wrote: "I recently purchased a DVD by Dr. Stephen Wilensky. He explains from the point of view of neuroscience that there is no such thing as free will as we can only perceive an action after it has already occurred. Can you please help me with this. I can understand that I

don't know what thought will occur next, but that that has already happened is beyond comprehension. Thank you, as I am in a lot of despair." It's sad and cute. I wrote her back. I said, "I'm writing this book. I'll send it to you when I'm done. And the last sentence in it is going to be 'Don't worry. Be happy.'" And I think that was the last sentence.

Okay. So then a little bit later we started getting hard evidence that giving people the news that there's no free will increases misbehavior, and I'll talk just about one of those studies. It was done by Kathleen Vohs and Jonathan Schooler. And what they did is they primed a group of people with the news that there's no free will, actual passages from scientific articles. And there was a pro-free will group and a neutral group. And then the next they were supposed to do is take a math quiz, and they were told that the program was glitchy so that if they didn't press the space bar right after the question showed up, then the answer would show up, in which case of course they could cheat. And then you could measure this just by whether or not they pressed the space bar.

Well, the group that got the no free will news cheated significantly more often than the other groups. The other two groups behaved the same, which is evidence that free will is a kind of a default assumption among people. There is a version of this study in which you get a dollar for every correct answer, so by cheating you're stealing too, and the no free will group stole more often in that one. A friend of mine, Roy Baumeister at FSU, did a study with hot salsa. I won't go through the details, but people behaved more aggressively when they got the no free will news.

So there's actual cause to worry about the news. I don't believe in censorship, but it's important to know whether these studies really do show that there is no free will or make it very

plausible that there's no free will. Somebody will ask me later maybe what free will is. That's complicated but we don't really have to talk about it now.

Okay. So the Libet study, I'm going to do it in two minutes. This is the neuroscience study that got the no free will ball rolling. And the subjects' task is to flex the wrist whenever they want. They're watching a really fast clock. It makes a complete revolution in two and a half seconds. They're taking EEG readings and they're measuring muscle motion. And what they discovered is that under a certain condition you get an EEG ramp up about 550 milliseconds on average before the muscle burst. And subjects are supposed to report on when they first became aware of their urge, intention, decision, or whatever, to flex now after they flex. And the average time there was about 206 milliseconds before the muscle burst. So there's a third of a second gap between EEG ramp up and reported time of first awareness of the intention or decision. And Libet says, "Look, what's happening is that the brain is unconsciously deciding to flex now and the mind becomes aware of it later."

Then there's the generalization claim which gets us to free will. It's that all decisions are made unconsciously, all decisions about what to do are made unconsciously. And then there's the thought that, look, if you're deciding unconsciously, you're not deciding freely. And then if moral responsibility depends on free will, no free will, no moral responsibility. And things are looking pretty bleak in a way because people like to believe that they can control their lives and prove themselves. And if you don't have free will or something like it, self control, you're in trouble.

Then these studies were done, similar studies with newer technology, so there are FMRI versions. One of the quotations I read you from the news is about an FMRI study, a famous one. There are depth electrode studies too done with epilepsy patients.

Now, I don't think it's been shown that these decisions, if they are decisions, are being made unconsciously, and I've written a lot about that. I'm not going to talk about it. I'm going to talk about the generalization move.

One thing to notice is that the experimental design shoves conscious reasoning off to the side, so the idea is be spontaneous. In fact, they are explicitly instructed not to think in advance about when to flex. So be spontaneous, flex whenever you want, report it to us later. And then another point to notice is that these decisions are arbitrary. So there's no reason to prefer this moment to begin flexing to that one, and so on. In all of these studies, the two-button studies, there's no reason to prefer this button to that one. Nothing hangs on it.

So these are decisions made without conscious reasoning and they are arbitrary decisions. So can you generalize to decisions about important moral matters, about what to do now, not about what's right, we're always talking about action, decisions about important moral matters that are preceded by painstaking conscious deliberation? Those decisions, at least on the surface, are so different from the arbitrary ones that aren't preceded by conscious reasoning, that it's hard to see how to build that bridge. A lot of work would need to be done. It's very hard to study what's going on in the head when people are deliberating for weeks, say, about what to do.

Okay. Another point is that what's really being tested in these studies is how good we are at detecting when we first become conscious of an urge or a decision or intention. And it may be that that's really not very important. We might be a little bit off, like 200, 300 milliseconds off. What may be most important is the conscious reasoning that influences the decision making. When I say conscious reasoning, I'm not talking about anything immaterial or involving a soul or anything of the sort. It's brain activity. It's just brain activity of a certain kind.

And then I think to close what I want to say is that we are going to continue to learn a lot of very important and very useful things about the brain, but that when we're thinking about implications for things that people care deeply about, like free will and moral responsibility, we have to think through the implications very carefully and maybe be careful about what we say to the press. We shouldn't exaggerate the results of our own studies, I think. And that's it.

DR. GUTMANN: Thank you very much. We are open. We'll have questions and give you a chance to elaborate on some of the issues that you raised and some maybe that you haven't raised, and let me begin with Nita.

DR. FARAHANY: Thank you both. Provocative as always and really helpful in illuminating some of the issues. I want to pose a question that I think bridges what both of you are raising, which is your working definition of free will that is motivating your perspective. So starting with Josh, it seems like when you're talking about the decline of blame or retributivism as a notion that can really animate reasons for punishment, that it assumes a free will that kind of ties to mental states rather than ties to freedom of action. And my perception is that law seems to care much more about action choices than it does about mental state, and that earlier in time we cared more about mental state but that that concept has become so thin over time in criminal law that mental state really means nothing more than, you know, did you mean to take that action, rather than what motivated you to take that action. So I'm hoping you can unpack a little bit, because I suspect in these studies that you've done more recently and in the scenarios you've given you have some kind of idea of what people are thinking free will means.

And likewise, for you, you said you'd be happy to unpack what you mean by that. I take it that you are talking a little bit more about freedom of action than you are about freedom of kind

of mental state and preferences. And so it would be great to hear from both of you what your working definitions are that are motivating the way in which you are approaching the problem.

DR. MELE: Philosophers have argued for a couple thousand years about what free will means. What I've always done is not to take a stand on the main point of division in that argument and to give two different sets of sufficient conditions for free will. And the way I think about free will is it's just the ability or power to act freely, which sounds empty, but what that means really is where you want to start is with free action. So what is it to do a thing freely? So here's one set of sufficient conditions, and I'm not claiming these conditions are necessary either. Sufficient. So you can whittle away at them and still get free actions. So if a person is sane and rational and well informed, uncoerced, unmanipulated, and makes a decision on the basis of good information, that I would count as a free decision according to this first set of sufficient conditions. That's enough for the decision to be free. If the person acts on the basis of the decision overtly, the overt action is free. Now, there's a kind of view that requires more than that. It actually requires that the universe be a certain way and that the brain be a certain way. What it requires is this: That in addition to those things, so let's have all those things in place, there's one more thing, and this thing is needed; that at the time at which you make the decision, if you were to roll things back and play them forward again, roll them back just let's say a second and play them forward again, things could turn out differently. And this is the kind of openness in decision making that so-called incompatibilists or libertarians want.

So one way to think about it would be like this: That the actual laws of brain activity that govern brain activity are probabilistic, not the laws we think we know. The actual regularities there are probabilistic rather than deterministic. Or, if you know a bit about quantum mechanics,

what you would need is some kind of activity of that kind in the brain in decision producing streams.

So the second set of sufficient conditions would build this kind of openness in. The first kind of free will I'm very confident that we have. You know, it happens sometimes. The second kind we didn't know yet, but we don't know that we don't have it either.

DR. GUTMANN: Josh, could you address whether you agree with that? Because that's actually a very good colloquial summary of two paths in philosophy, two paths in philosophy that have been debated through the ages with no settlement, but many, many philosophers, and I count myself among them, believe that you can have a deterministic universe, physically deterministic, that there may be -- I put it may be so I'll put it in the agnostic sense. Nothing outside of what exists in the universe and still have this very robust form of free will which is there's nothing spooky about it. In other words, as long as -- where coercion does not mean, you know, physical determinism. It means the ordinary sense of coercion, meaning that there is nobody holding a gun to our head or forcing us into an impossible choice. Is there any -- this would be very helpful for us. Everything I'm going to ask, there's nothing tricky about what I'm going to ask, because I think what we have to do as a commission is very basic here, but the basics are complicated.

Is there anything that Dr. Mele said that you disagree with?

DR. GREENE: So far in his main remarks, and I'd have to do a bit of interpreting on --

DR. GUTMANN: I didn't want to go back to all the -- just in his answer to Nita, the two paths of free will.

DR. GREENE: I don't think I can give a simple yes or no, but if I were forced to, I would say I pretty much agree with everything he said. But let me think where I know that we

ultimately diverge. First, just to get one issue out of the way, going back to the original remarks, I don't for a second think that neuroscience has shown that our decisions are made before we do them. I think those results are wildly interpreted, overinterpreted both by some of the scientists themselves, and certainly by the media. And so the whole idea of your brain decides before you do, everything is decided unconsciously, I think that stuff is just a big mistake. And that's where a lot of this stuff is coming from, and that's a lot of what Professor Mele is reacting to. I don't go in for any of that. And I could say a bit more about how people have really misunderstood the data.

But the more general question of does free will require that we are more than just physical systems, I think the answer is to some extent yes and to some extent no. That is, I think -- I don't think Professor Farahany asked for a definition. I think that to give a definition would be to just prejudge the question. What we really want to know is, what we really care about is who is ultimately responsible and why. And when we ask that question it leads us to this question of, well, did you really choose it, right, and what do we mean by really choose it, and that's where we get into this question of free will and free choice.

I don't think we come in -- if we just supply a definition, then we're just answering the question by fiat. So what we really have to ask is what is implicit in our ordinary thinking about these issues. And I think that part of our ordinary thinking is along the compatibilist line that Professor Mele outlined, which is as long as there's no coercion, as long as you're not obviously insane, as long as there's no special reason to think that you were constrained in some way or damaged in some way, then the choice is free, it's free will, you're fully responsible, et cetera.

I don't think that it's quite as simple as that, because I think that part of the intuitive experience of free will is spontaneity, that is, that I truly have this choice. I mean, to take a

trivial example, you're deciding do I want the soup or do I want the salad. And up until that moment you feel like it really could go either way. The past history of the universe is not determining it, and that at some point, bam, you say soup, and there's your answer, right? And you really, given history just as it was, could have gone either direction. I think that that's part of the intuitive experience of choosing, that it really --

DR. GUTMANN: So do you have evidence that there is one intuitive feeling of choosing? Because all of the evidence that I know of suggests people have very different, not infinitely different, but different intuitions about choosing. In other words, I don't know whether you are speaking now as a scientist or as a philosopher about people's intuitions about choosing. I mean, most people in the United States, if you ask them on a public poll, you'll get intuitions about whether there's angels in the universe, you know, all kinds of things. So are you saying scientifically speaking people's intuitions about choosing requires the incompatibilist's view or is there a division there scientifically speaking about what people's intuitions are? I'm trying to figure out you see where your neuroscience -- yes.

DR. FARAHANY: I gather that the answer is something about preferences and desires, right? Which is your intuitions and your preferences and desires are not something that actually from your perspective and from some of the neuroscience is necessarily within your control. And if you look back to the beginning of the universe until now the fact that I have a genetic predisposition to increased preference for sugar versus salad necessarily, things like that, those are things that go into the choosing of the action that I experience freely but in fact are outside of my control. So I gather that for you to say you need more than freedom of action, it's because preferences and desires that we might experience as free are in fact determined from the beginning of the universe until now; whereas, you are

looking more precisely at the action choice, which is yes, but still I am sorting between these preferences and desires and coming up with an action choice as to which path I'm going to follow.

DR. GREENE: There's a lot on the table there.

DR. GUTMANN: That's actually, Nita, not my question, but it's okay.

DR. GREENE: To answer your question, I think there is some evidence. It's not neuroscience evidence, it's behavioral evidence, that at least part of people's intuitive conception of choice is this kind of spontaneity. So, for example, I may be misremembering, in the experiment I believe by Shaun Nichols and Joshua Knobe where they asked people about a pot of boiling water, and you say if the heat is at this level and the temperature is at this level, could it not boil. And then I ask a parallel question about if the person is in a store and this person may or may not shoplift as far as we know, and everything in their brain is exactly the same up until that point, is it possible that the person could go one way or the other? And there people say the boiling pot is different from the person.

Now, we all agree that there are important differences between people and pots of water, but the question is, when you hold all of history constant right up until the moment of choice, if you really believe that the universe is ultimately a physical system, maybe there's some randomness in there but that doesn't seem to be what people are getting at when it comes to the choice, people should say that for that question those are the same. This is a hotly debated issue among people who are studying this. I could give you other examples, but it's not just me saying, "Hey, this is what it feels like to me." There's some evidence that people do have this. But I think it's not --

DR. GUTMANN: Right. That's overinterpreting an experiment, right?

DR. GREENE: You're saying you disagree with my interpretation of that experiment?

DR. GUTMANN: Yeah. I mean, that is -- not what you've just said about the experiment, but the leap from that to if people say there's something different about when a pot boils and when their decision is -- who knows what they believe. There are so many things they could be thinking that are different.

DR. GREENE: They are not just saying there is a difference here. The question is more specific than that. I will gladly grant you that the interpretation of these results is contentious.

DR. GUTMANN: Jim.

DR. WAGNER: This is fascinating. I comment as someone outside of the realm of expertise in moral philosophy. Part of what we are charged to do, of course, is to look at in this session for these past couple of days what we would advise about the ethics of neuroscience research and neuroscience application. Dr. Greene, I hear you talk about the physical system of the brain, a DSLR camera. Sure it has an automatic mode that nobody has to really think about, it just sort of happens, and then the deliberative manual mode. To the degree that we could make an argument that the brain actually can be reduced to a physical system, how does this advise the appropriateness of our recommendation about seeking to fix the machine, to modify moral behavior, and to encourage those kind of applications of neuroscience research?

DR. GREENE: I'm not sure that it actually changes that. Maybe I don't understand what you have in mind. But you can believe that people are purely physical systems and that you can make the world better by fixing problems. Everyone is fine with that. You've got a problem with your brain and let's fix it. There are interesting questions about is there --

DR. WAGNER: I do have to insert that I don't think you would say everyone is fine with that.

DR. GREENE: Like removing a tumor to improve someone's neural function? I mean, there are people who say you shouldn't do that?

DR. WAGNER: I think we could both find examples for one another where it would be obvious that, yes, we should go in and try to do something, and others where we shouldn't. And even the tumor removal is something that some folks would opt not to do. I'm just saying we have to be careful when someone says everyone would agree with that. My antennae go way up, and I'll bet yours would too if I had said it.

DR. GREENE: Fair enough. It's less contentious that there are things that we can do to people's brains, with their consent of course and with good evidence that it will work and not harm them and all of the things we expect from medical interventions, and there's a debate about where's the line between fixing a problem versus enhancement and what are the ethical issues that come with that, and I think that's interesting and tricky. But whether you have a full-blown metaphysical view of human nature where we're bodies and souls joined somehow, everybody agrees that there's a body.

And so I am inclined to think that these are important questions but that they are orthogonal questions. Not that everything isn't related in interesting ways, but I think that whether or not you think, when it comes to the proximate causes of behavior, just brains, or whether you think that we are brains that are in some sense being animated by minds or souls that are distinct from brains, brains are still what most immediately cause behavior. And I think it's relatively uncontroversial that there are things we can and should do to people's brains to help them, and where you cross

the line from that into more dangerous territory, that's an interesting question. But I think those issues are orthogonal.

DR. GUTMANN: Dr. Mele.

DR. MELE: Can I tell you something about how people think about choice? We've done a lot of control studies. Can I quickly do that? So my friend, Eddy Nahmias, years ago told a story in which there was a super computer that has complete knowledge of the laws of nature and of the state of the universe at a given time, and it predicts that John will do such and such on a certain day at a certain time. And John does it, and Eddy asks did John do it freely. And the great majority of people say yes. And if it's a bad thing, he asks does John deserve to be blamed for doing it, and a great majority of people say yes. So these are lay folk, great majority, saying something that's consistent with compatibalism, with the idea that determinism and free will can go together, and determinism and moral responsibility.

There are lots of studies like this. There are some that go the other way. It's a little complicated. I myself did a study about whether souls are required for free will. I had about 280 subjects, and the story was this, the cover story: Scientists have finally discovered that the universe is entirely physical, there's nothing nonphysical, decisions and intentions are brain events and they're caused by other brain events. Bill sees a \$20 bill fall out of the pocket of the person in front of him. He thinks about returning it but he decides to keep it. Did Bill have free will when he made his decision? In that condition, 73 percent of people say yes. And a lot of those subjects actually believe in souls; they just don't think that souls are required for free will.

So even lay folk, I think, a lot of them seem to be compatibilists. They don't seem to be requiring this fancy stuff. Why does that matter? Well, philosophers have argued for way too long about this, and we know what we think but it's nice to have an anchor out there in the world.

I'm not an ethicist, and so I'm definitely not a bioethicist. I love Josh's answer to your question because I could not do as well.

DR. GUTMANN: Dan.

DR. SULMASY: Thanks. Again, a very fascinating discussion. I was wondering for Dr. Greene whether either you have an opinion or perhaps there's evidence that the sort of intuitive automatic system you're talking about is itself malleable. Is it possible to change what one's intuitive responses are? Does it depend upon how one is brought up? Is this sort of genetically fixed as you seem to have been suggesting? I want to make sure. If that's not the case, that's good, so say that explicitly; and then if that is the case that it is malleable, how it is that you think we decide what are the good intuitions to have.

DR. GREENE: So as a general answer to your question, I think that our intuitions, our gut reactions, what I've called metaphorically automatic settings, are very much malleable. They are influenced by our genes, but they are very much influenced by our cultural experiences and by our personal experiences. So the idea that intuitive means hard wired or genetically determined, I want to be very clear, and thanks for giving me the opportunity to clarify that, doesn't mean that it's generic or hard wired.

Now, when it comes to particular domains, take for example visual illusions, we've all seen things where people come up with something where -- you know that these two lines, for examples, are the same length but one looks longer than the other. That's essentially an intuitive response. That's the automatic functioning of your visual system giving that false interpretation.

And you can spend 20 years convincing yourself, not that you need that long, that the two lines are exactly the same length and you're always going to see it that way. That would be a case where the basic wiring of the visual system is giving you an intuition, an automatic response or way of seeing it that's not going to go away.

When it comes to the kinds of moral and political issues that we deal with here, they are of course people's personal experiences, and of course people's cultural experiences play enormous roles in their gut reactions to these things. So we're talking about a very broad category of psychological responding. It's really most of the brain that does its work outside of the influence of conscious control.

DR. SULMASY: Then the second question: How do you decide which are the morally correct intuitions to have and to cultivate and to, as you're suggesting, that we have -- if you suggest that we can be mistaken in our intuitions, how do you decide that?

DR. GREENE: That is the big foundational question of moral and political philosophy. It begins as a moral question, not as a psychological question. I think psychology can help but I don't think --

DR. SULMASY: Psychology as a science cannot answer that.

DR. GREENE: Correct.

DR. GUTMANN: Is there any disagreement on that? Because that is actually very fundamental to something we have to opine on. So here let me just state what I think you said and let me ask both of you if you agree with that. Whatever neuroscience discovers about the determination of our actions, it cannot tell us what is right and what is wrong.

DR. GREENE: I agree, but I don't -- you can take that to mean two things: One is that the science is irrelevant, and the other --

DR. GUTMANN: I do not take it to mean that at all. I think we should talk a little bit about that before the end of the session, but let me just say that that is a total non sequitur to say that the neuroscience would be irrelevant.

DR. GREENE: Okay. I think it's -- It doesn't determine it by itself.

DR. MELE: I agree.

DR. GUTMANN: These are fundamental, and it's important. Steve.

DR. HAUSER: I'd like to follow up on Dan's question from the other side, the side of pathology. You mentioned, Joshua, that you get the caveat of assuming that there's no damage when you were speaking about the automatic systems. But let's take your, for example, anecdote about John, and let's add to the anecdote that John had suffered multiple concussions and that one could image noninvasive aggregate, tissue-destroying aggregate in some area of the brain or the prefrontal cortex that is relevant to automatic decision making. How would people feel about his free will in that situation?

DR. MELE: It's undetermined. It's a bit of randomness. I think that's a great question. You know, we could actually test it, but we'd have to make it clear to people what the deficit is. So we can't just do it in neural language. We'd have to talk about the psychological consequences, sort of the high-level consequences. But if the consequence was that he couldn't resist the temptation to keep the money, if we made that clear, I think most people would say he didn't steal it freely. I run these manipulation stories by people, brainwashing stories, and so on, and that really lowers the free will response. It comes down. Some people will say free will no matter what. It's about a 20 percent floor. I myself, I'd need to know more details. So what is the brain deficit associated with a level up? Inability to resist temptation? Then the guy is off the hook, I think.

DR. GUTMANN: I want to do the flip side, because we focused an awful lot on this big fundamental question of free will, and we've established an agreement that the question of what is right and wrong or should we hold people ethically or legally responsible, which is the practical implication of that, that broad question, is not determined by neuroscience. But let me just ask a question about what neuroscience does contribute, because otherwise it's not clear why we are taking this up in neuroscience.

So to what extent can neuroscience help us figure out how difficult it may be for people to do what we might think is right or wrong, how easy or difficult it might be? Is that something neuroscience can help us with? Or if that's not it, maybe you could tell us what it is, because I think there are things that neuroscience is beginning to tell us about the difficulty or ease of people making decisions that are held legally or by ethical standards to be right or wrong.

DR. GREENE: So I think that what neuroscience can do is correlate different states of the brain with different patterns of behavior. So if someone has this kind of damage, they are very likely to have a hard time resisting these kinds of impulses. In a straightforward way it can explain where the behavior is coming from. I'm glad this came up, because where I think these issues really matter most are when it comes to holding people responsible, and where I think it really matters is punishing people as an end in itself, not because it serves a practical purpose but because we just think it's inherently just. I may have been brought in as, okay, we'll get the guy who thinks there's no free will and that neuroscience answers everything, but really what I think all this comes down to is, when you connect all those dots which we haven't connected very thoroughly, the problem is really with retribution. The problem is with punishment as an end in itself as opposed to one that serves a social purpose.

But let me bring this back to the immediate discussion at hand. When you say things like how difficult is it, we already know from your example the person failed, let's say, to resist the impulse. They did the thing that they shouldn't have done. And I think, and there is some evidence to support this view about how people are thinking, some evidence against it, that what people want to know is, okay, the person didn't but could the person have. And that "could" is what we really I think have strong feelings about and it really matters to us. We know you didn't do it. You know you didn't return the money, right? And we know that statistically, let's say, someone with this kind of brain damage is more likely to do that than someone who doesn't have this kind of brain damage. But I think that what we want to know is there's a difference between the person who did the bad thing and could have done otherwise versus the person who did the bad thing and could not have done otherwise, the person who could have resisted, or degrees of difficulty. It was you did it but it would have been easy for you to resist; you did it but it would have been hard for you to resist. Well, if you're just a physicist, there is no, you know, here it was easy, here it was hard, but was it easy or hard for this ball to roll down the hill.

It's all ultimately mechanical interactions, and I don't think neuroscience is going to give us the kind of answer we want about that "could" because all we're going to find there are neurons making other neurons fire and ultimately making muscles flex, and that's what we call behavior. So the fact that we are even asking that question, we know he did it but could he have resisted, or how hard would it have been to have resisted, a satisfying answer is not just going to be a table of statistical correlations. The fact that we are asking that question to me suggests that we think there's some kind of element of freedom or choice that goes above and beyond the physical knocking of one neuron into another.

DR. GUTMANN: Nothing I said rationally speaking should cause you to believe that.

DR. GREENE: To believe what?

DR. GUTMANN: The fact that I asked how easy or hard it should be must mean that --

DR. GREENE: No must.

DR. GUTMANN: What?

DR. GREENE: A lot of people ask that question and -- I'm not critiquing your question.

DR. GUTMANN: So it doesn't matter what I ask. It's the fact that a lot of people ask that. Do you understand?

DR. GREENE: I think compatibilism is defensible, so I don't think that neuroscience has proven compatibilism wrong, but I also think that there are aspects of a lot of people's thinking, maybe not everybody's thinking --

DR. GUTMANN: I don't disagree with --

DR. GREENE: -- that are challenged by understanding that it's all ultimately mechanism inside the skull.

DR. GUTMANN: For sure there are elements of people thinking that, but that's not the only thinking there can be, right? Raju.

DR. KUCHERLAPATI: Thank you very much. I want to follow up on Jim's question. I'm not a philosopher and some of this discussion is way above my head, but I want to try to simplify at least my understanding and see if it is accurate or not. The question that we are trying to address is whether an individual always has a choice about taking one action versus another. I guess that would be considered free will. I'll give you an example, and I want to ask your opinion about that.

A few years ago in Netherlands a young man was caught by the police, and he was caught by the police because he was trying to set fire to a building. So they took him to the police station and booked him. And they found out when they went to the police station that this was not the first time this has happened; that he was caught a couple times before for trying to set fires to buildings. And then they also found out during that process that a couple of members of his family were also arrested for trying to set fires to buildings. And any time something like that happens, there is a family history of particular types of events and this particular case of behavior, then you try to study what the basis for that is. To make a long story short, it turns out that there is a genetic chain in this family and they have a deficiency in an enzyme called monoaminoxidase, and it turns out that that is what causes this particular type of behavior.

So the question is does this person who has this monoaminoxidase deficiency, did he have free will?

DR. MELE: That's a great question. One thing you'd want to know is did he have to do it or could he have resisted the urge to do it. And it may be -- I mean, free will isn't an all-or-nothing thing, and within people it may come in degrees. So he may have more freedom about what he eats than about whether he burns down my house, let's say. The condition may have made it very hard for him to resist but not impossible. If that's how it was, he's got a kind of diminished responsibility and diminished freedom. If it made it impossible for him to resist, I think he's off the hook. So that's the kind of thing you'd want to get more evidence about.

About punishment, I don't like punishment just because people deserve it either. I'm with Josh on that. I don't even want to be talking about punishment now. It's just free will. It could be a little free as opposed to a lot free.

Also, when we think about these things, what's characteristic of free actions? Well, people can generate options, possibilities of action. They can assess consequences. They can resist temptation to act contrary to their better judgment. And if any of those things gets damaged, there's less freedom. And I figure that in principle you can study this from a neuroscientific point of view. So anticipation of consequences, generation of options, ability to resist temptation, all of those things can be studied and they are studied in different branches of psychology.

So evidence that a person is damaged in one of those ways would be evidence of at least diminished responsibility, I would think, and diminished freedom, maybe goes to zero sometimes.

DR. GUTMANN: If doing more experiments and continuing to show what you've shown leads more and more people to believe that retribution is a bad thing, and I'm happy for Dr. Mele to -- I'm struggling with this, because let's just stipulate that I think retribution is a bad thing prior to being exposed to any neuroscience experiments or findings. Let's just stipulate it that I think it's a bad thing.

And you found the more you expose people to the findings of neuroscience they find that retribution is a bad thing. Does that mean retribution is a bad thing, I mean, punishing people? That's what I'm struggling with. I just don't see how neuroscience can establish that retribution in and of itself is a bad thing. An eye for an eye, a tooth for a tooth in the Bible I don't think depended just -- I don't think the view, having the view depends on necessarily -- it may correlate but I just don't see how it depends necessarily on any picture, internal picture of human action. I think you could believe -- and this is where, Raju, I think it is more complicated. I think you could believe that if people do something horrendous, you should show society that it's

a bad thing by punishing them even if it doesn't have a deterrent effect. I may not believe that, but I don't think it depends on neuroscience.

DR. GREENE: I don't think it logically depends on it either, which is to say you can hold the neuroscientific view fixed and have any moral view you want. But I also don't think that -- I could say more but it sounds like you want to move on.

DR. GUTMANN: No, no. I'm agreeing with you. Keep going. It does make a difference in how people feel about retribution because a number of people who believe in retribution believe in it because of a picture that neuroscience evidence starts to erode. Is that correct?

DR. GREENE: I do think that's right.

DR. GUTMANN: I ask these questions in an attempt to be constructive, trying to sort out what it can and can't do. Okay. Christine.

DR. GRADY: I have a perhaps simplistic follow-up on that line of thinking, and that is I heard both of you say that the sort of power of these data have -- that these data have the power to affect behavior too, you know, thinking about retribution after people have seen the data, the aggressiveness after some of the -- the despair. And I think, Dr. Mele, you said earlier we shouldn't overinterpret our data. Well, of course, we shouldn't overinterpret our data, but I guess the question is, is there more -- in this particular case should we be more concerned about the power of the data in terms of how the data itself and the neuroscientific findings affect behavior? And if so, what should we do about that?

DR. MELE: So should we be more cautious about what gets reported in the news, for example? I've never been a fan of censorship but -- again, I'm not an ethicist so maybe censorship is right and I just don't know it, but I doubt it. I think we just have to carry on the

way we're going. I do think too that if some neuroscientists were better informed about what free will might mean or what lay folk mean by it, and so on, then they wouldn't be making some of the claims they make about free will. This was actually part of the point of the Big Questions in Free Will Project. We had neuroscientists and social psychologists and philosophers all involved doing research together and enriching one another.

So my thought would be, you know, it's not a matter of policy or anything like that, but it is a matter of educating one another.

DR. GUTMANN: Thank you. My impulse and emotion tells me I want to go on, but my watch tells me that we are over already. And I really want to thank you both for a very edifying panel. Thank you.