



Presidential Commission
for the Study of Bioethical Issues

TRANSCRIPT
Current Issues in Neuroscience and Neuroimaging

Judy Illes, Ph.D.

Professor of Neurology; Canada Research Chair in Neuroethics; Director,
National Core for Neuroethics; Faculty, Brain Research Centre,
University of British Columbia

Stephen Morse, J.D., Ph.D.

Ferdinand Wakeman Hubbell Professor of Law and Professor of
Psychology and Law in Psychiatry, University of Pennsylvania Law
School & School of Medicine

Adina Roskies, Ph.D.

Associate Professor, Department of Philosophy, Dartmouth College

Meeting 4, Session 5
February 28, 2011
Washington, D.C.

DR. GUTMANN:

We are going to get started. So, if the people back there can come up and take a seat, I would appreciate it. Thank you. There you go. Thank you very much.

So, we are going to be discussing, again, and with our three experts at the table, some of the current issues in neuroscience and neuroimaging.

Judy Illes is a Professor of Neurology and Canada Research Chair in Neuroethics at the University of British Columbia.

She also directs the National Core for Neuroethics at UBC and is a faculty member in the Brain Research Center in the Vancouver Coastal Health Research Institute.

Dr. Illes is co-founder and Executive Committee Member of the Neuroethics Society and a member of the Dana Alliance for Brain Initiatives.

Her research focuses on ethical, legal, social and policy challenges, specifically at the intersection of the neurosciences and biomedical ethics.

Thank you for joining us today, Judy.

DR. ILLES:

Thank you, Professor Gutmann and Professor Wagner and Members of the Commission. It's a privilege to be your Canadian guest today.

My task is to explore in depth with you, the topic of incidental findings in the brain and brain imaging, and it's my pleasure to do so.

Let me first say that the work I'll show you today, and much of it is empirical work, represents really, the collaborations of many colleagues in Canada, the United States and around the world.

I am extremely grateful to them, to current members of my neuroethics team, as well as to sponsors of our research in Canada and the United States, including the National Institutes of Health, for their commitment to this effort and their support.

So, let me just show you the outline I'd like to follow this afternoon. I'd like to give you a little bit of background. Some of it, you will have heard already, the current landscape of incidental findings in neuroimaging, areas of ongoing discovery and I'd like to conclude with some uncharted territory.

So, you've heard this already today, incidental findings in the brain, potentially bad news that researchers or clinicians stumble upon in the course of research or a clinical examination.

Let's look at the case study to ground our discussion.

MK is a medical student, conducting functional MRI research for his PhD on memory. He has been a mentor to the incoming medical student class.

Two weeks into the program, SH, a transplant from the east coast, maybe Washington, D.C., is a new medical student and enthusiastically enrolls in MK's study.

On the anatomy pre-scan, MK notices an anomaly in SH's pre-frontal cortex. There is no institutional protocol in place for manning this kind of event. What should MK do?

So, we're talking about fundamental issues, real people and real problems, to our research participants, at the heart of the issues of protecting human subjects and research, with potentially significant disorders of the central nervous system.

Psychological and financial costs to them, risk to personal healthcare security, also known as insurance, relevance to third parties and the trust and reciprocity of research participants with the research enterprise. There is a cost to the research enterprise.

We know that there are an increasing number of functional MRI applications and of applications of neuroimaging overall. I'll call your attention to the graphic insert in the right. Simply, just look at the lines that are all increasing in terms of the number of functional MRI studies that are being conducted from 1991, on the left of the slide, to 2002, and the number of journals carrying these studies.

There is about 1,800 of these studies on average per year, from 2002 to 2008, with an increasing percent of about six percent.

So, we're talking about many increasing studies and we're also talking about expanding requirements for biobanks and data sharing, as Professor Wolf's elegant and compelling presentation gave us.

So, let's look at the current landscape. This is what structural anomalies in the brain look like. They may be small vessel disease, as we see here on the cavernous hemangioma. Might be arteriovenous malformation, as shown here, or a tumor, as shown here.

We are less concerned about these kinds of instances, when some of us have a flu or a cold and are entered into an MRI study. These are not of enormous concern, unlike these shown on the top of your slide.

The frequency is approximately 18 to 20 percent, in terms of the consensus of what we see in the retrospective studies of neuroimaging scans, in adults and in children.

We know from the Dutch literature and elsewhere, that they occur — that they're largely vascular and tumors in nature, and that they do seem to occur in one out of every five people scanned. That is some kind of anomaly in the brain.

Those that are clinically significant of the total typically affect two to eight percent, although we do have American data on ultra healthy Air Force pilots, suggesting that the percentage of anomalies in the brain are much, much less even than one percent. However, two to eight percent clinical significance is what we're seeing.

We know that in the older population, the frequency is quite high, in my own data, as much as 50 percent, and these typically require routine referral.

In the young population, these findings tend to be very rare. However, in these individuals, in the 20's and 30's, they tend to be of high clinical significance, requiring urgent referral.

We know that management strategies across neuroimaging laboratories are very variable. We have professionals conducting scans and PIs, and we also currently have, at least in the mid 2000's, graduate students and under-graduate students who are permitted to conduct scans.

In terms of neuroradiological review, we have scans that are reviewed by professional neuroradiologists on suspicious finding only, and in other cases, all of the time. So, quite a bit of variability.

We also have some very interesting data on subjects expectations, per say.

In a study of our own, on the west coast of the United States, we learned from student and — student subjects, adults who had participated in fMRI studies, that they would want to be informed of an anomaly detected in their brain, more than 90 percent of the time, regardless of the level of severity that the finding was thought to be, and even in those that were considered to be benign, more than 70 percent of them reported to us that they would seek evaluation.

So, this is — flies a bit in the face of how much we should really tell our human subjects, based on the severity of the incidental findings.

The challenge maps onto this three dimensional space. Pandora and her box, I really felt this way, when I started to embark on these studies. I opened the box and out came all sorts of things, a little bit like the fire hydrant analogy we heard earlier, an unbelievable amount of problems and issues — and challenges for our community, and this, in fact, is SH's brain, here, on the left — on the right side of the scale.

We see the dimensions of our problem space as follows: incidence and significance from low to high on the X-axis. The duty of care, perhaps, the duty of care is lower for PhD investigators versus those for PhD MD investigators, and on this Y — Z-axis, here, we have privacy issues that pertain to human participants right to know about an anomaly in the brain, as well as the right not to know.

So, I'll bring you back to MK and SH in the case study I showed you at the beginning of this presentation.

Through our consensus meetings, and that included the participation of Susan Wolf and others from Stanford and the NIH, we did have a majority consensus that research protocols must anticipate incidental findings.

However, we then agreed that there is range of morally acceptable dimensions, to how to proceed, with the majority agreeing on the path that I'll show you down the middle of the slide.

We did agree that incidental findings should be managed, although a minority felt that it was okay not to manage them because of the blurring of research and clinical medicine.

There was also a minority opinion that subjects deserve the option to decline to be informed.

However, I will assert that that really introduces a tremendous amount of complexity, because it imparts knowledge onto a subject about what a brain finding might be, and it also might create a situation in which there is a clinically significant treatable finding and it is then the researcher duty to aggregate the consent form and return the result to the subject, because there is a potential for treating the severe illness.

I'll bring your attention back down to the middle of the slide, and here, we have some options of how to proceed.

The principal investigator does detect something and flows the information to a physician, qualified to read a scan, or all scans are reviewed neuroradiologically, and finally, in the fourth panel, we have the issue of how to communicate results back to our human participants.

In the best case, no action needed. In the alternative case, an incidental finding must be communicated, needs to be communicated to the research subject or the surrogate, who is then encouraged to initiate follow up.

So, with that said, I'd like to go onto some areas of ongoing discovery in the area of incidental findings that you haven't heard yet this morning, and that I think will bring us further in the discussions of what we need to do, as a research community and what the Commission needs to consider, I believe, in taking up this issue of incidental findings in your deliberations, and move forward.

So, the first is the economic implications of managing incidentally found intro-cranial aneurism.

At the University of British Columbia, in collaboration with members of the Health Economics Department, we did some mathematical modeling of four strategies of different populations of human subjects.

Here were the strategies. No screening at all. No looking. No telling. No further work-up.

MRI read by a researcher not trained in clinical neuroimaging, with a subsequent reading by a specialist, if — upon the detection of a suspicious finding and then, onto MR angiographer, the subject, if it's verified by the specialist.

The third strategy where the MRI's are all — all MRI's are read by the specialist. That is all research MRI's. And the fourth strategy is the full grade clinical work-up for all participants.

So, let me just give you the data on strategy number two, where it is — where a researcher screens the data, and discovers incidental findings, and then passes on the findings — the discovery to a neuroradiologist.

First, let me show you this column here, screening by a researcher. It is never cost effective, according to our data, for a non-trained researcher to be explicitly looking for brain anomalies on research scans, although, that is the case in some institutions today.

No prior screening is justified from a health economics point of view, for men up to the age of 60, who do not have a family history, over the age of 60, when there is a family history.

The situation, however, is different for women and for men between the ages of 18 and 60, where a screening by a specialist is justified from a health economics point of view, at both \$50,000 U.S. dollars and \$100,000 U.S. dollars for quality adjusted life years.

That is, society's willingness to pay for a QALY, and here, you see for women with a positive family history of an inter-cranial aneurism, full clinical follow up is — full clinical pre-screening is actually cost effective.

So, we see here, a recommendation from us, that there may actually be a shift from researcher focus to participant focused decision making for individual, for incidental findings.

Customized strategies are essential and it may be morally acceptable to exclude certain research participants a priori, based on health economics data, as we understand them further.

Let me give you a second example of ongoing discovery, and I think it's here, where the Commission should also perhaps, pay attention, and it is at the interface of combined technologies, imaging and genetics.

In the case of imaging genetics, we do know that we're able now, to associate brain activation patterns with genes, as in Alzheimer's disease, schizophrenia and even pathological social responses, such as fearful stimuli.

In a graphic that was beautifully presented by Josh Rothman from Hoffman & colleagues, we see here, the interface of the laboratory between genes and clinical features, the end of phenotypes and the different kinds of neuroimaging we can use to understand the relationships between genes and clinical features.

We borrowed that cartoon, or that image, and placed some imaging of ethics considerations on it, thinking about how we're going to move imaging genetics into science and society, and here, our analysis gives us both increased discriminative power, afforded by the two technologies together, that is the ability to find more anomalies, more diseases, predict more diseases than either technology individually, and there, as well as the combined technology that the two give us.

Together, we see their application in all areas of social science and society, differentiation of disease, incidental findings at the heart of our discussions, with a particular attention to the fact that genetic analysis

and brain imaging might give us the same or might give us different results and have —

DR. GUTMANN:

One and a half minutes.

DR. ILLES:

That's fine, and how we're going to manage those findings is an open question.

The third area of open discovery is the brain at rest. Bruce Rosen spoke with us about that.

Unlike some of the tasks you've been hearing about, for functional MRI, earlier today, where individual subjects perform a task, now, we're talking about people fully in a resting state, where the resting state default network of the brain may actually be a biomarker for diseases like Alzheimer's disease, schizophrenia and other disorders of the central nervous system.

So, will task dependent functional MRI or task independent resting state functional MRI be the first functional frontier for incidental findings? I think that's an open question.

If it is resting state fMRI, I think there are some central issues we'll have to think about.

There is a current lack of understanding of what the resting state means. The networks connectivities are heritable. These data have to be processed off-line and how they'll be analyzed, and then if they exist in the case of biobanks is an open question, and if it is resting state, the implications for findings for the perception of self and social categories will be profound.

With that said, my last slide is uncharted territory. Economic analysis beyond aneurisms, incidental findings in children, touched upon, but very poorly elaborated, evolving processes for consent, responsibilities for non-clinical for-profit applications for neuroimaging.

How we're going to manage incidental findings in cultures in which ownership of health data and consent are actually shared by communities and finally, evidence based policies in a changing healthcare climate. Thank you.

DR. GUTMANN:

Thank you very much. Our next speaker is Stephen Morse. He is the Ferdinand Wakeman Hubbell Professor of Law and Professor of Psychology and Law in Psychiatry at the University of Pennsylvania.

Dr. Morse is the former co-director of the MacArthur Foundation's Law and Neuroscience Project. He is also a diplomat in forensic psychology of the American Board of Professional Psychology and many more.

But let me just welcome, Stephen. We're happy to have you here and look forward to your comments.

DR. MORSE:

Thank you very much. It's a pleasure to be here.

In 2002, the economists warned, "Genetics may yet threaten privacy, kill autonomy, make society homogenous and gut the concept of human nature, but neuroscience could do all those things first."

My message is that nothing we have learned since 2002, despite the immense advances in neuroscience, suggest that this is true or likely to be, and let me say, virtually everything I'm going to have say from now on about neuroscience applies also to genetics. I think with few exceptions, they can be treated together, as potential genetic — potential ethical problems.

You already have in your briefing book, my 2004 Law and Neuroscience chapter, in which I argue the new neuroscience raises no new problems for ethics, although it will produce data that will inevitably require new applications.

Nothing in years since publication have changed my views. Others, including my friend and colleague, Judy Illes, disagree, but I shall not repeat that argument here.

Instead, I wish to focus on criminal responsibility and moral responsibility, more generally, because this has been the locus of most of the debate within law and ethics, about how the implications — how the implications of the new neuroscience.

My thesis is that criminal responsibility and law, more generally, is an inevitably folk-psychology enterprise, that is, an enterprise that depends fundamentally on mental states, as a partial explanation for human action.

It's an inevitably folk-psychological enterprise that takes people seriously, as acting agents, to whom deserved praise and blame, reward and punishment may properly be ascribed.

Nothing in the new neuroscience remotely suggests that the agentic picture of ourselves, that underlies interpersonal life and its regulation by morality and law is incorrect and needs to radical revision.

Taking each other seriously as acting, deserving agents is central to our moral and social lives and it should not be abandoned without the most compelling reason to do so. The radical challenge should be firmly resisted.

I have a few preliminaries, so that I won't be understood, I'm not a dualist about mind and brain. Our language is shot through with dualism. It's almost impossible not to talk dualistically. I'm not a dualist.

The brain clearly enables the mind, but we do not have a real clue yet, about how this happens or the brain/mind action connection.

Wittgenstein famously asked, in the *Philosophical Investigations*, "What is left over, if I subtract the fact that my arm went up, from the fact that I raised my arm," and the truth is, we are no better able to answer that question today, than we were when *Philosophical Investigations* was first published.

If we discover how the brain enables the mind, then I suspect that there will be a revolution in our understanding of ourselves. It is not clear that we will ever be able to answer that question, and certainly, not in the lifetime of those in this room, I speculate.

I might add that I, unlike the mysterians who believe it's impossible, beyond the resources to answer the question, of how the brain enables the mind, I would not stick on that issue.

Now, despite my physicalist and determinist meta-physics, which rules out God-like contra-causal free will, the ability to act uncaused by anything but yourself, I nonetheless think the robust concept of responsibility is possible on what's known in philosophy, and I don't have the time to do it here, a compatibilist reading.

This is a reading that says even if determinism or something quite like it is true, robust responsibility is nonetheless possible for acting agents like ourselves.

Since this is a non-resolvable meta-physical dispute, the question is, what is the most normatively desirable meta-physics to adopt, since

compatibilism, which believes in the truth of determinism or something like it, is fully consistent with moral, ethical and legal theories of responsibility that we endorse and it's consistent with facts we know about ourselves as human beings, as biological human beings, it is the most normatively desirable meta-physical position to take on the responsibility debate.

Now, the real question, in putting together either genetics or neuroscience for responsibility is this.

Genetics and neuroscience have a completely mechanistic discourse. They are not neuron — neurons aren't normative. Neurons don't think. Genes don't think and the like.

Law is all about acting human agents, as is morality. They're both guides to human action. They help regulate our interpersonal lives together.

The question is, how do you do the translation from this mechanistic discourse, whether in neuroscience or in genetics, to the action discourse of law, and by the way, there is no remote replacement for the action discourse of law and for morality in sight.

Could there be tomorrow, if something comes out of a laboratory? Sure, but today, nothing in sight.

Now, here is why it's an important ethical issue. So much is at stake. Legal cases, legal issues involve people's real acting material lives. Moral evaluation is so crucially important to the way we interpersonally live together. It's crucial that we get it right.

So, my question about the translation always is this. How precisely does the neuroscience evidence or the genetic evidence, how precise, and I want to really underline precisely, how precisely does it help us answer a legal or moral question that we need to have resolved, and it's important to us to resolve?

Here, I want to draw a distinction between what I call rhetorical relevance and real relevance.

There are a lot of people bedazzled by the pictures they see in the newspaper or in the professional journals. They think, "Oh my goodness, razzle-dazzle," they are, you know, sort of neuro-radicals cut — this is going to be the answer to our question.

Well, it's not going to be the answer to our question, because if you actually look, and here is the real relevance point, seldom so far, is the translation successful.

Seldom so far. In other words, if we really understand the criteria of law and morality in the action language they use, seldom does the neuroscience or genetic evidence really help us resolve these questions.

Now, again, this is hostage to what will be coming out of the laboratories tomorrow or 10 years or 15 years from now. Science is ever evolving and moving. It's just, when we think about the translation from a mechanistic correlational, coarse-grained analysis, which is what you get, largely in neuroscience today, especially for cognitive research on neuroscience to law, the translate doesn't — the translation doesn't work well.

This is a plea, again, for neuro-modesty as opposed to neuro-arrogance. We're understanding that there is a lot. The brain is an enormously complicated organ, as we all know, and we should be modest about what we know causally, before we're trying to tell people, "Let's run the railroad differently, from how we run it today."

So, what I want to avoid, always, is neuro-arrogance. Neuro-modesty, I'm happy with. Neuro-arrogance makes me cross, and I have identified something that I call brain over-claim syndrome. It is, I believe, a curable disease with cognitive jurotherapy, but all right.

Now, here is a very important point that again, gets at the translation issue, actions speak louder than images. It's a wonderful phrase. I wish I had made it up myself. I did not. A writer for Nature did, and what it basically says is, if we are trying to evaluate an acting human being, either morally or legally, and we observe the person's behavior, their thoughts, their mental states, which we, of course, infer, and then we look at images and seemingly, there is a disjunct, what we always have to believe for these purposes, is the behavior with very rare exceptions.

In fact, if someone is to tell you that, "My goodness, that person couldn't have been behaving that way, because this is the way their brain looks," what I want to tell you is, the use of the brain, to answer that question, is invalid. Just surely is, and they have — they're making some kind of mistake.

Now, in the future, I think there are some modest kinds of contributions that neuroscience and genetics can make. Let me just give you some quick examples.

In thinking about criminal responsibility, the laws makes a difference — makes a distinction between — I'm now being very rough — doing things on purpose, knowing what you're doing, being consciously aware of a risk, or not being aware of the risk you're imposing, but you should have

been aware of the risk. These are the mental states we use to individuate culpability.

Well, maybe those aren't the right categories, and what I can certainly imagine happening is, if you will, a reflective scientific equilibrium between the philosophy, the psychology and the neuroscience, where we're constantly refining the categories with each other's help. That, I certainly think is possible.

It could certainly help us with some efficient practices. We use predictions all the time now in criminal law and elsewhere.

If we're going to use predictions, and we already think that they are proper to use, I can't imagine what the argument looks like for doing it less well, as opposed to better.

So, is it possible that genetic evidence or neuroscience evidence in the future might help us make these predictions more efficiently, and the answer is, it's an empirical question. I assume the answer is yes, and then one of the ethical issues, a really important ethical issue will be, given the cost of collecting neuro or genetic evidence, given the other kinds of problems we might have with incidental findings and all sorts of things like that, what is the value added of doing the neuro-scientific or genetic investigation?

How much in addition, are we going to get, given the cost of getting it, and that's a really important issue.

Now, what I want to turn to, finally, is what I call the radical challenge, or sometimes, I call it the disappearing person challenge, and this is an argument that says — and now, I'm borrowing from a particular article from a group of very thoughtful people: You think you're an acting human being. You think your mental states at least partially play a role in explaining what you do. Well, you're sadly mistaken.

You are just a victim of neuronal circumstances. That's all you are, and essentially, what your brain is doing is throwing off all this mental state, epiphenomenally. You've got mental states. No one denies that you have them, but they're not doing anything.

Think of them like your appendix. Somehow, they're a product of evolutionary development, of natural selection, but for reasons we don't understand very well, they just don't do anything. It's neurons, neurons, neurons, all the way down.

Now, this is not a straw person. This is a real position, adopted by thoughtful, intelligent people, who actually have a prescriptive set of

recommendations that they believe follow from this position, which is, they hate deontological approaches to ethics. They hate retributivism in criminal law, although they often misunderstand retributivism in criminal law. They often think it's unnecessarily harsh or cruel, which is not the case, at all.

But they also think that a purely consequential view of law and morality is entailed, if we are just victims of neuronal circumstances.

I don't think anything is entailed, if we are just victims of neuronal circumstances, which I will get to in a moment.

There is a lot of empirical evidence out there that could be taken as indirect or direct evidence for the thesis that it's just neurons all the way down. At present, I don't think this evidence remotely comes close to proving the case of, we are just victims of neuronal circumstances.

Again, I don't know what laboratory work will show tomorrow, but in this area, there has been an immense amount of over-claiming. In some of my other work, I've gone through this evidence in pain-staking detail. I don't have time. It would take me my full 15 minutes, just to go through a part of it.

Let me just wave my hand and say, I think I make a very plausible case, that we're not remotely there yet.

But there are good reasons to reject the victims of neuronal circumstance, and that's what I think is more important for us, here today.

First, it's the implausibility of that kind of reductivist account. The best thinking in the philosophy of biology and the philosophy of science, generally, is that we are likely to have a full account of complex behavior, a multi-field, multi-level description to get a completely adequate causal account.

So, we're going to need biology at different levels, psychology at different levels and sociology at different levels. Common sense.

What would it take to convince you, that your mental states play no role in your lives? There is positive laboratory evidence that intentions and other mental states do, in fact, have a causal effect on our behavior.

There is a plausible theory of mind in the philosophy of mind literature. It's physicalist, but it's non-reductive.

The theory of mind in psychology, the ability to understand that other people have mental states and to guide your conduct by your understanding of other people's mental states. If you develop normally, we all have it. If you don't develop normally, you have let's say, some kind of unfortunate disorder, such as an autism spectrum disorder, in that case, you don't have an adequate theory of mind. You can't "read each other's minds", draw reasonable inferences, and you're not a very successful human being.

Evolutionarily, we know the ability to do this has been around in primates for about 40 million years. It's unlikely that it's just like your appendix, and lastly, and here is the really important point, so far, everything I've said does not suggest that it isn't true that we're victims of neuronal circumstances.

What I'm trying to do is shift the burden of persuasion to those who say we are, and therefore, we should have these radical changes in our moral, legal practices.

It's normatively inert. Here is why. If it's really true, that we're just victims of neuronal circumstances, it's just neurons doing stuff, and our mental states play no role in explaining our behavior, then our reasons play no role in explaining our behavior, and if our reasons play no role in explaining our behavior, then we have no genuine reason to do anything, one thing or another.

So, if you accept victims of neuronal circumstances, it is normatively inert and that's good enough reason to reject it. Thank you.

DR. GUTMANN:

Thank you very much. Adina Roskies is an Associate Professor of Philosophy at Dartmouth. She recently received the Lawrence S. Rockefeller Fellowship at Princeton's University Center for Human Values, and she is a recipient of the William James Prize from the Society of Philosophy and Psychology.

Prior to her appointment at Dartmouth, she completed the post-doctoral fellowship in cognitive neuroimaging at Washington University in St. Louis, where she utilized positron emission tomography and the then, the newly developed technique of functional MRI.

We look forward to your input. Thank you very much for being with us, Adina.

DR. ROSKIES:

Thank you very much, Dr. Gutmann. Thank you for inviting me to speak to you.

I was asked to speak about issues at the intersection of functional neuroimaging and moral philosophy.

I will interpret moral philosophy quite broadly, to encompass both practical, ethical issues and deeper, more foundational moral questions.

You've already heard today, a number of excellent presentations about specific aspects of neuroscience and neuroimaging, and as today's last speaker, I will try to summarize what I take to be some of the novel major ethical challenges that neuroimaging raises.

So, neuroimaging is primarily a diagnostic technique that yields descriptive, and to some extent, predictive information about brain function. It's a non-interventional visualization technique, although it's often used to guide interventional techniques, such as use of pharmaceuticals, trans-cranial magnetic stimulation and deep brain stimulation, and the ability to influence or enhance brain function with these interventional techniques raises a host of deep philosophical questions that is beyond my charge today and actually, no one has really spoken about them, so I urge you to consider those.

However, even the prospect of providing information about brain structure and function raises an interesting and challenging set of ethical questions.

So, like the human genome project, neuroimaging makes available, information about the present and future dispositions of persons, information that was unavailable prior to the development of these techniques.

I'd like to suggest that information about the state of the brain has the potential to be even more sensitive and more personal than genetic information. Our brains, quite simply, are more closely linked to who we are psychologically than our genes.

If our genomes can be understood to delineate the space of possibilities for the way our lives might go, our brains may perhaps, map out a space much more closely tied to actuality. In fact, our brains might be who we are.

Brains are the repositories of all our memories, the proximal causes of our behaviors, including our mental states. Thus, information about the brain could potentially provide information about a person's past, present and future behaviors, dispositions and mental states.

So, today, I'll touch briefly on the issues of mental privacy, lie detection, prediction, the ethics of consciousness, freedom and responsibility and public understanding of science. So, a tall order in a few minutes.

To begin, let's consider mental privacy. Each of us is accustomed to his thoughts being accessible only to himself, and with neuroimaging, however, the privacy of our inner-most thoughts is at least, potentially threatened.

The risk is far from immediate, but the recent advances made by novel techniques in discerning mental content from brain scans are impressive.

Although there is good reason to doubt that detailed contents of thoughts will ever be able to be read off of a scan, especially an unwilling subject, perhaps enough semantic information will be available to matter, at least in certain cases.

So, for example, imagine legal situations where self-incrimination is a possibility. Are there instances in which brain scans could be compelled or should be prohibited?

With respect to the law, should scans that reveal the content of mental states be treated more like physical evidence or like testimony?

A number of interesting and potentially important questions bear upon how evidence from neuroimaging squares with, for example, Fifth Amendment rights.

More immediately, there are already substantial efforts to use neuroimaging for lie or truth detection in ways that don't require the identification of mental content. So, various methods have been devised, in general, by measuring correlates of familiarity or arousal.

Although some reports have claimed accuracies of up to 90 percent, when care is taken to look at the rates of false positive and negatives, and to take into account base rates, such tests are clearly far from reliable enough to use in the courts.

Furthermore, careful consideration of the research reveals serious shortcomings in the experiments by which these technologies are developed and tested, and so, so far, I think these fail to establish external and ecological validity.

But despite these caveats, there is clear and continuing incentive to develop neuroimaging techniques to distinguish truth from lies, for the

use in the courtroom, or perhaps, more worryingly, for use in interrogation.

Mounting pressures to detect terrorist plots suggest that we'll continue to see efforts to improve such techniques and to apply them.

Serious thought should be given to the conditions under which it could be ethical to use such techniques. I fear that setting to low a bar may set us on a slippery slope, where brain data will be taken to be more admissible than they should be and far more dispositive than they, in fact, are.

Prediction is another issue, raised by neuroimaging. So, prediction of future behavior and prediction of future disease, like in the genetic arena, it shows promise — it shows promise for predicting future disease.

For example, significant progress is being made in prediction of Alzheimer's.

Because brain diseases and also, their treatments, alter the brain, there is reasonable worry that the disease and also, its treatment, if such is available, will threaten aspects of people's identity, autonomy, self or personality.

The ethical aspects of all of these concepts and the relative value are matters about which there is significant ethical disagreement. Moral questions thus arise, regarding how predictive knowledge of future disease should be handled.

For example, how should one assess courses of action when pitting, for example, health against authenticity?

There are indications that neuroimaging may be useful in predicting future aggression, criminal recidivism or mental illness, and when I say predict here, I mean that in a probabilistic sense.

How should such probabilistic information be understood and how and under what circumstances, is it morally permissible or even obligatory, to use it? What levels of probability warrant action or treatment?

More fundamentally, there are difficulties in realistically assessing the predictive power of our techniques, recalling neuroimaging shows correlations and not causation.

Suppose, for example, that we find that certain patterns of brain activity are correlated with the disposition to behave violently. Adequate

assessment of the predictive power of our data requires knowledge of the prevalence of these activity patterns and behaviors in the general population, something that, in general, we do not know.

Lack of this kind of base rate information severely limits our ability to interpret results in individual cases, and current scientific funding structures, unfortunately, do not encourage acquisition of this kind of knowledge, even though having that kind of knowledge would be a scientific boom and would certainly enable us to interpret our results more realistically.

So, now, I'm going to turn to the ethics of consciousness. Recent breakthroughs in understanding and diagnosing disordered states of consciousness have brought ethical issues about consciousness to the fore.

A few years ago, Owen and colleagues showed with neuroimaging that an unresponsive patient in a minimally conscious state was able to comprehend instructions and intentionally use mental imaging as instructed.

A more recent study showed that a non-trivial percentage of patients diagnosed to be in a vegetative state can do this, as well.

So, neuroimaging can be used to help diagnose patients with brain injury and appears to have predictive power for regaining function, for their prospects to regain function. Thus, neuroimaging might provide novel and superior ways of diagnosing levels of consciousness and predicting future prospects.

Along with such abilities come responsibilities, to address the ethical issues that arise. Many ethical questions arise when we recognize the possibility that those that we thought lacked consciousness may, in fact, have or have the potential for higher brain function.

Should all patients with disorders of consciousness undergo such tests? Suppose we can devise paradigms with which to communicate with patients who are overtly unresponsive? Should patients who can communicate be considered competent to make their own life decisions? What if imaging indicates poor prospects for recovery? Neuroimaging can also inform decisions about ongoing care. Evidence about consciousness raises questions about pain and quality of life. Do we have an ethical obligation to provide pain management for those who are unresponsive?

In part, due to the absence of clear philosophical and scientific understanding of consciousness, these questions are particularly difficult.

In addition, we lack an overarching moral theory that enables us to weigh, for example, considerations of autonomy and personhood with various measures of well being.

Moral perspectives that privilege welfare and those that privilege autonomy serve up, of often serve up, conflicting recommendations.

The moral frameworks that we bring to bear on these problems will affect subsequent policy.

The elephants in the room, of course, are the related issues of freewill and responsibility.

Briefly, there are two main views of freedom, one on which free will requires an absence of causal determinism, that is typically called libertarianism, not the political sense of the word, and the other, which conceives that free will is compatible with determinism or called compatibilism.

I do not believe that neuroimaging will be able to demonstrate that we are free or unfree, by showing that we have or lack libertarian free will. But there is certainly a lot of confusion on this issue in the literature.

However we look at it, neuroimaging results are bound to put pressure on common sense notions of freedom and its relation to responsibility. So, here, I'll argue a little bit with Stephen, although I think we're very much on the same page.

Neuroimaging allows to discern mechanisms of cognition, including mechanisms of choice in normal humans. We as understand more about cognition, we understand in more detail, the causes of our behaviors and recognize them to be physical. This, in and of itself, worries people. How can we be responsible if we were merely physical mechanisms?

Now, as Professor Morse has noted, causation is not compulsion and nor is it an excusing condition. But one doesn't need to believe those fallacies, in order to be bothered by the deliverances of neuroscience.

It seems uncontroversial that when agents lack certain capacities, we do not hold them morally responsible for their actions, or at least less responsible for their actions.

Neuroscience evidence increasingly suggests that brain differences accompany differences in dispositions, for example, aggressive or psycho-pathological behavior, and while we can acknowledge that given a scientific perspective, it would be a shock if it weren't the case.

Neuroimaging, nonetheless, raises the important question of which capacities are essential for moral agency and to what extent we should hold people responsible for actions that issue from abnormal brain function.

Indeed, it raises the vexing question of what differences are merely normal differences and what differences are those that can function as excusing conditions?

A number of philosophers and scientists have suggested, for example, that psychopaths should be exculpated, due to the nature of their deficits. There is increasing indication that there are brain deficiencies, or at least brain — systematic brain differences in psycho-paths, for instance.

Others have gestured towards a slippery slope, whereby no line between the responsible and non-responsible can be anything but arbitrarily drawn, and Professor Morse talked about those.

So, even if we accept a capacity based view of responsibility, recent results in neuroscience and psychology challenge other aspects of our common sense notions of volition and agency and the relationship to responsibility.

One such worry is whether we actually act for the reasons that we think we act for. Again, there are answers to this. I don't believe that no theory of responsibility that respects our best scientific data, is impossible to develop, but I venture to say that there is not theory of moral responsibility on offer, that adequately responds to current challenges and to those foreseeable in the near future.

The “my brain made me do it” defense is clearly a non-starter, but a better theoretical basis will need to be developed to deal with the hard cases.

One final important issue for the Commission to address is the public misunderstanding of data from neuroscience and brain imaging. So, I'm very much in line with Dr. Parens with this.

The public seems to privilege neuroscience data over behavioral and psychological data because it is hard science. What people fail to grasp is

the extent to which the very interpretation of neuroscience data is grounded on behavioral bedrock.

Moreover, as with genetics, there is a deep misunderstanding of biological determinism. People often fail to understand that information about a genome or brain is merely one piece of information in an exquisitely complex and dynamic system.

Not appreciating the many levels of plasticity in the brain may lead to poor inferences about future potential and about responsibility.

There is also a mistaken tendency to view differences as evidence of dysfunction, and finally, the lack of basic understanding about the brain and of neuroimaging methods leads to mistaken beliefs about the implications of neuroimaging data.

For example, people fail to realize that it's often impossible to draw inferences about individuals on the basis of current science.

These problems in understanding are endemic and difficult to manage because neuroimaging is an increasingly complex technology in an increasingly scientifically illiterate society. A concerted effort to educate the public and encourage the media to report science responsibly strikes me as an exceptionally important policy recommendation.

In sum, there are a number of philosophical questions that beg for immediate treatment. Can we develop a viable theory of moral responsibility that comports with our best science? What do we owe to persons with different degrees of disorder consciousness? How should we understand probabilistic predictive data and what are the permissible uses of it? How should we approach issues of mental privacy?

On the policy side, a better public understanding and encouraging the pursuit of information that enables us to best interpret neuroimaging information strike me as important goals. Thank you very much.

DR. GUTMANN:

Thank you, and I want to thank all three of our panelists for a really stimulating talk, and we have time for questions.

If there is anybody in the audience who would like to ask a question, just raise your hand and I'll recognize you. Right there, and then I will see if we have more.

PARTICIPANT:

My question goes back to Dr. Rosen, from Session 2 this morning.

He mentioned that we are still a few years away from cohort brain imaging, but in 2009, early 2009, CNN reported that they were using remote neuro-monitoring in the War on Terror.

Perhaps, you can explain the difference and is it possible that the Defense Department has moved past the medical community for neuroimaging capabilities? Thank you.

DR. GUTMANN:

Does anybody — if you want to answer it.

DR. FARAHANY:

I do want to answer it.

DR. GUTMANN:

Okay, go ahead.

DR. FARAHANY:

So, the CNN report talked about this detection technology, supposedly remote detection technology, something which I have spent a good bit of time trying to get to the bottom of, and the technology that they reported, I am — I think they've mis-reported.

They reported about some technology that is thermal imaging technology, not remote brain technology.

Thermal imaging technology, which is still also in its infancy, looks at blood flow changes in the face and micro-facial changes. It doesn't get passed the face. It doesn't go into the brain. It doesn't look at the mind. It doesn't do any sort of mind reading. It doesn't do much differently than we already have done for years, in detection of micro-facial changes in individuals.

So, that CNN story, which I read and I know many people did, as well, I think is not really relevant to our discussions here today.

DR. GUTMANN:

Okay, but it's relevant in clearing up the misunderstanding. So, thank you.

DR. FARAHANY:

Absolutely, on the neuroimaging side, I think it does.

DR. GUTMANN:

Thank you for doing that. Yes? Hank? Would you mind standing up? Thank you.

PARTICIPANT:

What I believe is that law enforcement and the Military have weapons that —

DR. GUTMANN:

That's not Hank, okay.

PARTICIPANT:

— remotely influence people. If you're going to make guidelines that a brain image is what's going to determine whether they're guilty or innocent, they can be attacked, even remotely. There are weapons — one that's called Medusa.

It's a microwave attack. It gives microwave auditory effect. They could hear threats, violence. If you read their brain, they would appear to be delusional. They might become anti-social.

But the truth would be they are under an influence. Just by taking a reading of their brain, I don't think is something that can actually determine their state. There is a cause for that state, unless that's understood, I don't think this will work. I'd like to ask your opinion about it.

DR. GUTMANN:

Does anyone have a response to that? No? Other questions from the — do the Commission members have any questions? Martha?

DR. WAGNER:

She has a response, maybe.

DR. GUTMANN:

Martha has either a response or a question. So, Martha.

DR. WAGNER:

We need the microphone up here.

DR. GUTMANN:

Okay, a question.

DR. FARAH:

It goes mainly to Stephen and Adina.

DR. GUTMANN:

Okay.

DR. FARAH:

Although, perhaps Judy will weigh in, as well, and it has to do with two kinds of claims that could be viewed as neuro-arrogant.

One kind of claim, that I am definitely willing to join Stephen and say, that is neuro-arrogance, in fact, it's neuro-silliness, is the idea that, you know, once we really understand how the brain makes behavior happen, the notion of personal responsibility, of moral agency, of acting for reasons will be undermined.

But in this session, and the earlier sessions, we were kind of moving back and forth between that claim and I think, a much less bold claim, which is that — and I guess I want to know, do these panelists think that the following claim is neuro-arrogant, that brain imaging is — has already begun to give us an unprecedented ability, in terms of just predictive power, to understand a range of human behaviors and predict a range of human behaviors, not just in the medical realm, but in, you know, economics, business, marketing, you know, education, learning, lying and all the rest.

Like Stephen, do you want to tar us all with the neuro-arrogance brush, or just the people who are undermining the notion of moral responsibility?

DR. GUTMANN:

So, let me just repeat the question, because I think it does — it's an important question to answer, and it sounds, as I've heard it, as an empirical question.

Is it the case that brain imaging has begun to give us an unprecedented ability to understand and predict a wide range of human behavior, and I would just add, as we answer it, if you answered in the affirmative, it would be very helpful for the Commission, to know the behavior that brain imaging is enabling us to understand and predict, in an unprecedented way, yes.

DR. MORSE:

Well, understanding and predicting are two very different kinds of questions.

DR. GUTMANN:

Yes.

DR. MORSE:

Understanding suggests we have some sort of causal understanding of what is going on, and there, I would suggest we have a much less ability.

In terms of prediction, we already have all sorts of behavioral measures for predicting a wide range of fields, economics to marketing and the like, and one of the questions is, how much value added or how much additional marginal predictive accuracy does the brain imaging produce, and especially, given its cost. That's an empirical question. I'm willing to look at, case-by-case, and I just think we would need to decide that, case-by-case. My guess is, given the cost of getting it, given how good our behavioral techniques are, and given the behavior is still the gold standard, we neuroscientists don't go on fishing expeditions. They already have a targeted, well addressed, well characterized behavior, and they want to predict better or understand better.

Then, the question is, case-by-case, how much better do we do with the neuroimaging, and that is, I think, an open question and I don't think it's a bold claim to say, we might, in some cases, do it better. Let's just see the evidence case-by-case.

DR. GUTMANN:
Yes, Adina?

DR. MORSE:
I think that there are cases where we clearly can predict things that we couldn't otherwise predict.

DR. GUTMANN:
Can you give us an example?

DR. MORSE:
Well, there are some examples in the neuroimaging literature where you can, for instance, predict just better than chance, whether somebody is going to add or subtract some numbers.

I think that without neuroimaging, we didn't have any way of weighing on those things.

Now, but what that means is very different. I think what it indicates is that there are maybe factors unknown to us, as choosers, that have impact on our decision making and that is not a surprising thing to say.

I mean, it wouldn't be surprising, with the physicalist stand anyway, to say that completely independent of neuroimaging. But now, there are proofs that that is the case, and more than actually undermining any — by evidence, undermining any theory of responsibility that people might have, I think that what imaging does is call their attention to tensions in their folk-psychological understanding, their folk-psychology, I think implicitly has a sort of dualist bend to it, and I think all brain sciences

move people away from that kind of view, or put pressure on that view, which I think, requires them to rethink things.

And while I don't think that imaging can unseat our notions of responsibility or agency, I think it might change them. I think it might make people question things that they took to be unquestionably true, and that's where I think neuroscience puts pressure on our ideas about responsibility.

DR. GUTMANN:

Judy?

DR. ILLES:

So, I'd like to answer that important question, by submitting that neuroimaging, I believe, has profoundly advanced our understanding of neurologic and psychiatric disease, and the knowledge that we've gained through research has directly improved the lives of people with mental illness, as well as move towards society's understanding of what mental illness means and attempts to mitigate stigma and discrimination and so forth.

I do think prediction is a place of — is the cutting edge. We certainly have neuroimaging techniques that can predict certain diseases, as well as non-clinical behaviors, and I think of the old days of EEGs, in which EEGs could actually give us direct measurements of anticipatory behavior.

What neuroimaging will not give us, I don't think in our lifetimes, are measures of motivation and intention. There is a real complexity that comes together.

DR. GUTMANN:

What is striking about the answer, because I've repeated Martha's question, about understanding, put that aside, predicting a wide range of behavior, is how fuzzy the answer is to that.

Where is, if you ask that question about, you know, in — if you ask a different question, to what extent does — can genetics now, tell us what the increase or decrease in the probabilities are, of certain diseases, we get a very specific — very specific answers to that, you know, based on the mutations or not, of different genes.

So, there is — and this is just the development. I mean, as I think you all have agreed, it varies with the development of the science.

So, the science is not sufficiently developed and I'm putting out a hypothesis here, for those of you who study this directly, to agree or disagree.

The science is not sufficiently developed, to tell us, in very specific ways, what the mapping of a brain scan and human — certain human behaviors will be, except in the extremes, I mean, where you already see the behaviors, and so you're not predicting.

DR. ROSEN:

I think, all right, for example, all right, the Alzheimer's disease data, when you're pin positive, it's much more predictive your likelihood to convert to Alzheimer's, even when you're cognitively normal.

DR. GUTMANN:

Okay, good, good.

DR. ROSEN:

And there is increasing data that the same is going to be true in say, schizophrenia and in autism.

DR. GUTMANN:

I'm not saying it's not going to be true. You've given a good example, where Alzheimer's, it is true.

DR. ROSEN:

Well, and the —

DR. GUTMANN:

That's good.

DR. ROSEN:

Already, the predictive value, given the variability with the genetic data, right, the predictive value of the presence of certain genes is very, very small for almost all behavioral diseases.

You know, Huntington's disease is a great example. You know, it's yes or no, but there is no gene for schizophrenia, autism, etc., though these are highly genetically pre-determined diseases, where the propensity is any greater than a few percent.

The neuroimaging data is already significantly better than that. Is it 100 percent — even in AD? Not yet, but it's getting there.

DR. GUTMANN:

Okay.

DR. ROSEN:

So, I think you'll see a spectrum in both, and I think in that regard, the neuroimaging, in some ways, is ahead, in others perhaps, perhaps not.

DR. GUTMANN:

Okay, so, now, the wide range of human behavior. We're said Alzheimer's and we've mentioned some other diseases.

The wide range of human behavior, the — what a lot of members of our society did, lots of them, we won't need to name — to cause the great recession.

You said, you know, wide range of human behavior, are we anywhere close to that? I mean, I'm just trying to map out the parameters, here.

PARTICIPANT:

Okay, good evening —

DR. GUTMANN:

Because you did — there is a sweeping question about the wide range of human behavior.

DR. MORSE:

I actually agree with the impulse that Martha and Judy and Adina have put forward, that we're learning things.

But here are a couple of cautionary notes. So, for example, Professor Roskies said, "Brain images are showing us that we don't know all the causes of our behavior." In 1979, in a wildly influential psychological review article — so, we're now talking 32 years ago — Nisbett and Wilson wrote an article called Telling More Than We Could Know, at which point, they went through a wide range of empirical psychological research, showing beyond per-adventure, that we often do not know all the causal background of why we do what we do, and we often mis-attribute.

Second of all, so, let's take psychiatric disorders, which as we know, can be crippling.

When the American Psychiatric Association began its run-up to DSM-V, as probably the people in this room know, they were quite sure that by the time DSM-V was published, we'd have good neuro-markers for psychiatric disorders.

Well, DSM-V is about to be published. We don't have one yet, which is not to say that we haven't found lots of good neuro and genetic studies, showing differences between people with certain disorders and controls,

but none of them is sufficiently sensitive or specific to be used diagnostically.

Where we have the clear evidence, and remember, we're doing these studies on people who are clearly defined as suffering from these disorders to begin with, we didn't need the neuroscience. And it's precisely where we need the science, in the unclear cases, where we have the least good neuroscience or genetic evidence, because you've got a clear-cut problem.

DR. GUTMANN:

Okay, I'm going to take one more question and then we're going to convene, as a Panel. May I ask you to — and identify yourself, then ask your question.

MS. SNYDER:

My name is Miriam Snyder. I believe that this Commission was launched because of some of the horrific medical experiments that have been done. Is that why the President initiated this?

DR. GUTMANN:

No, we were launched, actually, well before those revelations came, and we will talk about that charge, by the President, tomorrow.

MS. SNYDER:

Right, but with respect to neuroimaging, I have found a lot of research based information regarding injection-induced seizures through and aligned with neuroimaging.

I would like to know what protections are going to be put in the place, with respect to this new neuroimaging, to protect people from becoming injected with, whether it's a microchip, whether whatever it may be, that leads to seizures, with respect to the advancement of neuroimaging, so to speak? That's part one.

Part two, yes, with the history of these horrific medical experiments, can we — can at some point in time, we focus in on how — what is going to be done, to make sure this process, neuroimaging, induced seizures — what is going to be done, to make sure these type of horrific experiments are going to be stopped, because I think that when you launch a committee like this, instead of sitting here, talking this and that, that's really not what I know the public is here for.

We have a history. We have problem. There are people being killed, behind human experimentation, that is unregulated, unsupervised and above all, the research shows that injections are being used to induce

many of these diseases under the study of neuroimaging. What's going to be done to stop this? Thank you.

DR. GUTMANN:

Thank you. We are going to take up the topic tomorrow, of human subjects research. So, we will come back to a series of issues that that raises.

I'm going to ask all of us to thank our three speakers, again, for a really stimulating time, and if I could ask all the previous presenters to join us at the table.