Neuroethics
DOI 10.1007/s12152-010-9085-1

COMMENTARY

No Nonsense Neuro-law

Sarah K. Robins · Carl F. Craver

Received: 28 April 2009 / Accepted: 8 June 2009

Abstract In *Minds, Brains, and Norms*, Pardo and Patterson deny that the activities of persons (knowledge, rule-following, interpretation) can be understood exclusively in terms of the brain, and thus conclude that neuroscience is irrelevant to the law, and to the conceptual and philosophical questions that arise in legal contexts. On their view, such appeals to neuroscience are an exercise in nonsense. We agree that understanding persons requires more than understanding brains, but we deny their pessimistic conclusion. Whether neuroscience can be used to address legal issues is an empirical question. Recent work on locked-in syndrome, memory, and lying suggests that neuroscience has potential relevance to the law, and is far from nonsensical. Through discussion of neuroscientific methods and these recent results we show how an understanding of the subpersonal mechanisms that underlie person-level abilities could serve as a valuable and illuminating source of evidence in legal and social contexts. In so doing, we sketch the way forward for a no-nonsense approach to the intersection of law and neuroscience.

Keywords Personal/Subpersonal distinction · Neuroscience · Law

S. K. Robins (✉) · C. F. Craver
Washington University in St. Louis,
St. Louis, MO, USA
e-mail: skrobins@wustl.edu

C. F. Craver
e-mail: ccraver@wustl.edu

Introduction: Sorting Sense from Nonsense

Pardo and Patterson accuse contemporary neuroscientists of engaging in nonsense. It is not that neuroscientists have false or misleading ideas about the brain and its relation to the mind. Rather, their ideas are so thoroughly confused that they are gibberish:

> The upshot of this conclusion is not that claims that the brain 'follows rules,' 'interprets', and 'knows' are false; it is that these claims are lacking in sense (2010).

Pardo and Patterson thus conclude that neuroscience has few, if any, implications for thinking about the law. They borrow from a long philosophical tradition associated with Ludwig Wittgenstein [1] and the "ordinary language philosophers," such as Norman Malcolm [2] and, more recently, Bennett and Hacker [3]. Many of the central arguments for the view can be traced back further to Thomas Reid's 18th century common-sense psychology [4]. The attraction of this tradition is that it promises to dissolve (rather than solve) our problems about the relationship between the mind and the brain. The apparent problems arise only because we (21st century philosophers and neuroscientists) are confused about how person-level phenomena such as beliefs, knowledge, inference, and reason are related to the activities in our brains. In the brain, one finds causes, not inferences. One finds patterns of neural activity, not knowledge. One finds reflexes, not reasons for acting. To relabel these neural phenomena as inferences, knowledge, and reasons is to commit a category mistake.

Published online: 16 July 2010

It is to blend personal and subpersonal descriptions [5] thereby offending the very meaning of our mental terms. The way out of confusion is to stop looking in brains to explain person-level phenomena. Person-level phenomena are features of people operating in social and normative contexts. They are not features of people's parts.

We grant (for now) that neuroscientists and philosophers often speak, think, and write as if brain regions decide, interpret, know, recognize, and understand. We also grant that such thinkers can be deeply confused when they do so. That is, they are not merely adopting a manner of speaking or using temporary filler terms as stand-ins for some-more-adequate-characterization-we-know-not-what. They really believe that brain parts make inferences, know things, store images and maps. Let us grant that this is nonsense.

Our thesis is that this nonsense is irrelevant to the challenges neuroscientists are raising for the law and legal theory. Our reason: Subpersonal neuroscience suffices to raise all the interesting legal and moral problems that Pardo and Patterson hope to dissolve. The nonsense can be brushed away from the empirical bedrock of subpersonal neuroscience with no consequence for the discussion at hand. To see this, one must carefully review the methods and findings of contemporary neuroscience. We begin such a review here and conclude by identifying the central issues that scholars must address if they are to grapple effectively with the legal and moral consequences of our inevitable progress in neuroscience.

The Empirical Bedrock All parties agree that neuroscience has discovered and will continue to discover mechanisms in the brain that are *necessary for* person-level abilities and states. Thus, Pardo and Patterson:

> We in no way take issue with the claim that the brain is intimately related with mental life. We recognize that the feelings and activities that we associate with the mind depend upon a (properly functioning) brain. Particular neurological states, in others words, may be a necessary condition for various mental activities. Finally, we do not contest that neuroscience may illuminate how these activities depend upon the brain and how damage or defects in the brain may affect one's mental activities, including whether one has the capacity to engage in these activities at all (Pardo and Patterson 2010)[1]

Let's be more precise. First, everyone grants that there are correlations between person-level goings on and brain-level goings on. Neuroscientists discover such correlations with a host of techniques (including, for example, EEG, fPET, fMRI, multiunit recording, and single unit recording). In experiments using these techniques, one engages a subject in a task and measures changes in the brain that occur as the task is being performed. Neuroscientists measure blood flow, patterns of activity in neural populations, single cell activity, and changes in molecular events, such as protein synthesis or the concentrations of neurotransmitters. One looks for correlations between these subpersonal measures of the brain and performance on the person-level task. When two variables are correlated, they are interdependent in the sense that the value of one variable can be used to make predictions about the value of the other. That is, the value of one variable *indicates* (perhaps fallibly) the value of the other.

Second, to say that brains are necessary for minds is to say, minimally, that people would not have minds if they did not have brains (*ceteris paribus*). But we can say more than that, of course. Neuroscientists have known for nearly two centuries that specific person-level skills (such as remembering, planning, and speaking) require specific activities in more or less circumscribed regions of the brain. This knowledge is grounded in clinical facts about deficits in patients with focal brain damage, and it has been

[1] Likewise, Bennett and Hacker write: "Neuroscience can investigate the neural conditions and concomitants of the acquisition, possession, and exercise of sentient powers by animals. It can discover the neural preconditions for the possibility of the exercise of distinctively human power of thought and reasoning, of articulate memory and imagination of emotion and volition…what it cannot do is replace the wide range of ordinary psychological explanations of human activities in terms of reasons, intentions, purposes, goals, values, rules and conventions by neurological explanations. And it cannot explain how an animal perceives or thinks by reference to the brain's or some part of the brain's perceiving or thinking" ([3], p. 3). See also [6] and [7].

confirmed with a variety of reversible lesion methods (such as cooling, pharmacological intervention, and transcranial magnetic stimulation). On the basis of such evidence, it is possible to infer what a person can and cannot do from what one knows about the functional and structural integrity of their brains. A person with extensive hippocampal damage, for example, will likely not remember what they have done or what they are told. A person with extensive damage to Wernicke's area cannot be expected to give informed consent on the basis of verbal instruction.

Finally, neuroscience has found and will continue to find ways of intervening to change brains in ways that are *sufficient in the circumstances* to bring about changes in how certain person-level tasks are performed. One can stimulate brain regions and cells (with, for example, electrodes or pharmacological agents) to produce visual experiences [8], coordinated movements [9], and to evoke apparent memories [10]. In many areas of neuroscience, such interventions are at best a distant possibility, but neuroscience is young.

Pardo and Patterson are prepared to grant all of this. They insist, however, that this fails to show person-level activities or states just *are* brain activities or states. This is true, but irrelevant. The fact that subpersonal mechanisms are necessary for (or even simply correlated with) person-level abilities and states is sufficient to raise legitimate concerns about how these findings can, should, and will be applied in civil, criminal, and military contexts.

Reading Minds Pardo and Patterson object to the idea that knowledge is in the brain. (They also object to the idea that rule-following and interpretation are in the brain, but the issue of knowledge is more fundamental. See their note 7). In particular, they object to the idea that one can tell what subjects know by looking in their brains. This idea, they say, "depends on a confused conception of knowledge. To know something—knowledge that propositions about a crime are true, for example—is not located in the brain. As a conceptual matter, neural states of the brain do not fit the criteria for ascriptions of knowledge" (2010). Because it is nonsense to think that knowledge is the brain, it is folly to try to tell what someone knows by looking in his or her brain.

Two clarificatory comments are required before we make our case. First, Pardo and Patterson's use of the term "knowledge" is a distraction. Knowledge is (at least) true belief. Neural facts might indicate what a person believes, but neural facts cannot indicate whether the beliefs thus discovered are true (unless, of course, the beliefs are about neural states and related facts). Thus, one cannot attribute knowledge to subjects by looking at their brains. One has to look outside of the subject, at the relevant portion of the world, to determine whether the beliefs thus discovered are true. However, the most interesting and troubling applications of neuroscience to the law rely only on the ability to tell what a person believes, true or false.

Consider the case of Prosecutor trying to establish that Defendant killed Victim with the motive of collecting Victim's life insurance. In that case, Prosecutor needs to establish that Defendant believed she would receive the money. Prosecutor does not need to establish that Defendant would *in fact* receive the money. The motive can exist even if Defendant is not the beneficiary and even if there never was an insurance policy. What matters is whether Defendant believed there was an insurance policy and believed that she would be the beneficiary. The important question, then, is not whether neuroscience can discover Defendant's knowledge but whether it can discover facts relevant to Defendant's "state of mind" (including Defendants abilities, beliefs, and motives).

Second, neuroscience should not be expected to deliver *incontrovertible* evidence about Defendant's state of mind. To demand certainty from neuroscientific evidence would be to demand more of it than any empirical evidence can provide [11]. Pardo and Patterson are attacking a straw man when they claim that:

> neuroscientific evidence might reveal that certain brain activity is inductively well-correlated with this behavior, or that damage to certain brain areas makes one incapable of engaging in this behavior, but it cannot establish conclusively that one's brain is engaged in lies or deception or that an intent to deceive or a lie is located in the brain (2010).

Evidence is a finding that changes the probability that a hypothesis is true (relative to its rivals). DNA, eyewitness testimony, and fingerprinting all count as sources of evidence despite the fact that they can yield false positives and false negatives. If brain states are correlated with person-level states, then brain states can function as evidence for or against the attribution of person-level states. Given that people act falsely, lie, and malinger, measures of defendants' brain states might, in some cases, provide relevant evidence about their true state of mind. Why, we ask, do Pardo and Patterson think that correlations and causal relations (rather than identity claims) are insufficient a) to supply new forms of evidence to be weighed among others, and b) to meet the standards of evidence accepted in the use of DNA, eyewitness testimony, and fingerprinting?

With these clarifications in mind, let us now consider some uses to which neuroscience might be put, starting with the most obvious and least controversial.

Malingering

Suppose Helen claims to be blind. You are charged by an insurance company to discover whether she is malingering. One kind of evidence is person-level. Perhaps you can catch her playing ping-pong or reading a (non-brail) book by the fire. But suppose her behavior is reasonably consistent with the claim that she is blind. You might inspect her eyes to see if they are present, unobstructed, and in apparent working order. You might check her pupilary reflexes, rotate an optokinetic drum or a mirror in front of her eyes to see if it induces nystagmus, or test her galvanic skin responses to intense light shown directly into her eye. Moving into the brain, you might check for damage to her retina (using an electroretinogram), her optic nerves, or her primary visual cortex using imaging techniques. You might see if a light stimulus changes the basal occipital rhythm of her EEG. The point is simple and uncontroversial. If you are confronted with a good and consistent actor then subpersonal states and activities offer independent evidence for assessing their mental lives.

One might object: Seeing is a person-level phenomenon. Eyes, nerves, and visual cortices do not see. Seeing is something the whole animal does in a social and normative context. Such an objection would miss the point that even if seeing is a person-level activity, one can nonetheless learn about that person-level activity by studying subpersonal mechanisms. Subpersonal facts provide evidence about whether the person can see, and so about whether they are malingering. None of the evidence is incontrovertible, but no inductive evidence is incontrovertible. The subpersonal bedrock is sufficient to raise the legal problem.

A side comment: This is enough to keep legal scholars busy. We know that brain damage often leaves patients with deficits to some cognitive functions and not others. If certain cognitive capacities are necessary for one to engage in certain person-level tasks (such as the ability to imagine the future, to compare subjective values, to assess counterfactuals, to understand the emotions of others, and so on), and if those person-level tasks can be degraded independently of one another with focal kinds of brain damage, then one can begin a nuanced neurolegal evaluation of how competency, agency, and responsibility are affected by damage to specific areas of the brain. Were such a "clinical moral psychology" available, subpersonal facts might be perfectly relevant to whether a person can consent, act for reasons, and be held responsible for his or her actions. Such a project will require genuine collaboration among philosophers, legal scholars, neuroscientists, and clinicians. It should be a focus of research in neurolaw. It could not even be taken seriously if Pardo and Patterson are right.

Locked-in Syndrome

Pardo and Patterson, following Wittgenstein, believe that knowledge (or better, belief) is an ability. It is manifest in what we say and what we do. They discuss ventral pontine syndrome (VPS; commonly known as "locked-in syndrome") as a potential counterexample. People with this syndrome experience global or near-global paralysis as a result of damage to the motor output of the brain through the ventral pons (a brainstem structure) while the rest of their brain remains in more or less proper working order.

Pardo and Patterson bite the bullet: It is only because the person with near-global VPS learns to communicate by blinking her eyes that one is justified in thinking that she still has beliefs, thoughts, reasons, and the like. In the case of *total* paralysis, they opine: "If such a patient were not conscious of their knowledge *in any way*, and could not manifest it in any way, on what basis would we ascribe knowledge to them? We would not" (2010). But Pardo and Patterson cannot appeal to consciousness, Wittgenstein's beetle in the box, to save their view. It is supposed to be behavior, not consciousness, that determines whether there are any person-level activities going on. And a patient with total locked-in syndrome cannot behave, at least not on the standard interpretation of what counts as behavior.

We urge Pardo and Patterson to forego the bullet. Recently neuroscientists have shown that some patients with total locked-in syndrome can deliberately alter their brain activity in ways that allow them to correctly answer questions posed to them about their families and homes [12]. Furthermore, they can begin and end their responses at the arbitrary times requested by the experimenter. Why, we wonder, should brain activity be any better or worse than eyes, hands, or (for that matter) tongues as a vehicle of communication for these patients? Why cannot a person use their brain states to indicate what they want or what they know? What matters is whether the patients use the artificial language we teach them in ways that conform to standards of truth and intelligibility, not which parts of their body they use to communicate.

Let us place ourselves in this situation, wondering if our loved one is in a locked-in state, with beliefs, experiences, and desires hidden behind the mask of a totally paralyzed body. Suppose that we learn from anatomical brain scans that (like other people with VPS) our loved one has damage to the ventral pons but that the rest of the brain is more or less structurally sound. Suppose further that functional brain scans such as EEG and fMRI show typical responses to sensory stimuli and to more complex cognitive tasks. Suppose she can answer your questions accurately and intelligibly by modulating the activities of her brains. Now, consider some of the choices before you. Should you remove the ventilator or perhaps pursue active euthanasia? Should you continue to visit the person, to talk with her, to read her books, to play music for her, to update her on events in family life? Should you try to find a way to use brain-machine interface systems to enhance her range of available actions?

While we acknowledge that these are complex questions, we submit that subpersonal evidence (the only evidence you have) should be weighed very heavily. Given that neuroscience has shown that various activities in the brain are correlated with, necessary for, and productive of certain mental states and abilities, evidence of brain function is crucial for deciding whether or not our loved one is "still in there" and whether we should invest the time and energy to keep her comfortable and, with the advances in brain-machine interface, try to "get her out." Thus far, very few patients with total locked-in syndrome have been shown to communicate via brain activity. But these discoveries justifiably bring hope to people whose loved ones find themselves in this situation.

The case of the locked-in syndrome is a bit like the case of the malingerer. In both cases, we suspect that the person's behavior is an inadequate indicator of the person's state of mind. We then use neuroscientific techniques to discover whether the person's subpersonal mechanisms continue to function as they do in normal perception and cognition. It is not nonsense to use the best evidence at your disposal.

Willful Lying

Pardo and Patterson object to the very idea that neuroscience could be used to build lie detectors. Again, they press Wittgenstein's conceptual point: "Neurological states do not fit the criteria for ascriptions of lies or deceptions" (2010). Brains don't lie. People do.

This line of attack confuses two questions: 1) Can brain regions lie? and 2) Can one detect willful lying of persons by looking at subpersonal indicators in their brains? We have already granted a negative answer to the first question, but this has no implications for the second. Even if brain regions do not lie, they might nonetheless indicate to us that someone is lying. The question is whether the brain does something different when people lie than it does when they tell the truth. This is an empirical question, not a conceptual or linguistic one.

Consider a more familiar form of lie detection. George is a terrible liar. Every time he lies, he raises his eyebrows and widens his eyes, ever so slightly. People who know George well can spot it right away. George's betraying eyes are reliable subpersonal indicators of his person-level fibbing. It will not do to complain that George sometimes raises his eyebrows when he's not lying or that he sometimes lies without raising his eyebrows. Useful evidence need not be incontrovertible. Nor is it productive to object that eyes, as parts of George, cannot lie. Eyes do not have to lie to be indicators of lying. Similar remarks apply to the use of blood pressure, galvanic skin response, muscular contractions, and pulse as more familiar lie detectors.

We are not arguing that any of these methods actually work. Rather, we are arguing only that it isn't nonsense to suggest that they might. Perhaps people get nervous when they lie. Perhaps their blood pressure goes up. Perhaps their eyes move in characteristic ways. Perhaps they activate different brain regions, or activate them differently. Perhaps not. Whatever the answer, this is an empirical matter, not something that can be derived from the ordinary use of the terms "lie" and "deception."

Memories and Stored Beliefs

Let us now consider whether one can use facts about the brain to infer facts about what a person does or does not remember. Pardo and Patterson object to this possibility on the grounds that memories are not literally stored in the brain. Neuroscientists' persistent illusion that memories must be stored in the brain is grounded in a mistaken view of memory. Memory, they claim, is simply retention of knowledge (belief) previously acquired or possessed by a person.[2] The idea of an *engram* (the "trace" of past experience or learning in the brain that is a cornerstone of the contemporary neuroscience of memory) is nonsense.

How is memory possible? Thomas Reid, who objected to the engram on similar grounds as Pardo and Patterson, candidly admitted, "I think it appears that memory is an original faculty, given us by the Author of our being, of which we can give no account, but that we are so made" ([4]: 197). Pardo and Patterson can do better.

Start with their idea that knowledge (or belief) is an ability. The ability in question, we presume, involves behaving in specific ways in the diverse circumstances that require the knowledge or belief in question. If we retain the knowledge that Columbus is the capital of Ohio, we will say so when asked, we will board a bus for Columbus if we plan to protest on Ohio's Statehouse steps, we will answer C if asked for the first letter of the name of the capital of Ohio, and so on. Let us suppose that the ability in question could be characterized in terms of a complex set of conditionals relating various stimuli to the appropriate responses. A person properly deserving this attribution of knowledge must be able to satisfy this complex input-output relationship tolerably well. If this is roughly correct, then to explain the retention of knowledge we need to explain the retention of abilities.

Of course, brain mechanisms lie between the stimulus and the response in each case. The ability-defining inputs presumably enter via the senses. The behavioral outputs are produced by coordinated nerve impulses to muscle fibers. How is the ability to effect these transformations of inputs onto outputs maintained over time? An economical possibility is that there is an internal mechanism that has been tuned by experience to do just that. And it is the existence of such mechanisms that marks the difference between remembering a previously performed action or piece of information and encountering such things for the first time (as well as the difference between remembering and relearning). The fact that we are capable of recalling information reliably argues that there must be some such mechanism or set of mechanisms. At the very least, it appears to be an inference to the best explanation that there is some underlying, subpersonal change that makes remembering possible. This hypothesis is surely more plausible than Reid's.

The crucial thing to note is that one need not insist that memories *are* states in the brain in order to think that the brain might contain indicators of the knowledge a person has or retains. Again, Pardo and Patterson are confusing the question of whether memories simply are brain states with the question of whether brain states can tell us anything about memory. Brain parts are components in causal mechanisms necessary for the maintenance of knowledge (abilities) over time. There must be some feature

[2] It is well known that this epistemic conception of memory fails to distinguish cases of memory proper from cases of relearning, such as when a person relearns something that they had forgotten about their past by reading their biography or, perhaps more interestingly, their own diary. Martin and Deutscher [13] develop a causal theory of memory to honor this distinction (see [14]).

of those brain parts that explains (causally) why we engage in the specific behaviors that we do in the sensory circumstances in which we engage in them. If we can detect the mechanisms responsible for those input-output dispositions, and if we can distinguish them from one another, then we can detect memories. That is, we would be able to find reliable indicators of the abilities (beliefs, memories) that constitute a person's retention of knowledge. If this is nonsense, it is the same nonsense that has fueled centuries of progress in biology, engineering, and medicine. It is the nonsense idea that mechanisms underlie manifest abilities and that there can be no difference in those manifest abilities without some difference in the underlying mechanism. Let's call that mechanistic difference-maker in these cases the engram.

It will not do to object, as Pardo and Patterson do, that engrams would be useless because people would have to know how to interpret them (2010). Engrams are not interpreted on this picture. They are sub-personal (causal) facts about the brain mechanisms responsible for maintained abilities. Nor will it do to object (as Pardo and Patterson do) that people cannot see their brains, and so could not read their engrams. One would have to be very confused about the explanatory project of contemporary neuroscience to think that models of memory require a homunculus or require people to look at their brains. Engrams are not (or should not be) posited as part of the person-level explanation of remembering. Rather, engrams are features of the causal mechanism that makes remembering possible.

We have thus far argued that the idea of an engram makes sense given basic principles of mechanistic explanation. There is also compelling evidence that neuroscientists have already discovered indicators of memory in the brain (e.g., [15]). They have also begun to discover how such indicators can be modified electrophysiologically and pharmacologically so as to change an organism's manifest abilities. For example, one can predict a rat's path through a maze by recording from the neurally encoded "spatial map" in its hippocampus [16]. The rat's hippocampus engages the same firing pattern repeatedly after the rat runs the maze, and the hippocampus generates similar maps when the rat is returned to the same maze again [17]. One can also intervene to degrade or erase memories. Rats treated with ZIP, a peptide that inhibits production of the protein PKM, have damage to their long-term memory for taste aversions, even when ZIP is given a month after the aversion was acquired [18]. These and other experimental paradigms help neuroscientists to understand how the brain makes learning and memory possible, even if learning and memory are not states of the brain, but retained person(or animal)-level abilities.

Consider one example of how the discovery of neural mechanisms might have legal implications. Neuroscientists have recently discovered that EEG records show a characteristic response to familiar stimuli (see [19, 20]). In a laboratory setting, subjects are presented with a set of images to study. Later, they are given a second list comprising both images familiar from the first list and novel images. Familiar images produce a characteristic change in the positive-going portion of the EEG wave 300 ms after presentation (known as the "P300" response). Novel items do not produce it. Some believe this technique could be used to determine whether a person has previously experienced a given scene or situation. Of course, outside the laboratory there are significant questions about the source of that apparent familiarity, about the ability of subjects to deploy countermeasures, and about the appropriate protocols for conducting such tests [19]. But such concerns are not relevant here. The question we are addressing is whether this kind of test makes sense. The fact that the test accurately distinguishes familiar from novel stimuli, even in highly controlled laboratory settings, suggests that it does.

It is a further empirical matter whether this method works well enough to provide reliable, time-sensitive information in various extra-laboratory contexts and whether that information can live up to the standards of evidence in civil and criminal cases. These matters cannot be decided in advance by analyzing our ordinary ways of speaking and thinking about memory; they require empirical work. Perhaps Pardo and Patterson would defend a weaker claim, namely, that behavioral evidence must always trump the findings of neuroscience because behaviors are *criterial* for belief whereas brain activities are merely indicative of belief (see p. 25). But again, we urge that the question of how we determine the meaning of mental state terms be set aside. What is at issue here is the evidence that can be used to determine whether a person has, or is capable of having, a particular mental ability or state. We have argued that facts about subpersonal mechanisms might be reliable, if fallible, indicators of such abilities or states. Person-level evidence is also fallible. People can have beliefs that are not expressed in action (e.g., when they are unable or unwilling to express them). It is precisely when we suspect that a person's behavior does not reflect their abilities that subpersonal evidence might help to settle the matter.

Conclusion: The Hard Problems of Person-level Indication

We have argued that it makes sense to look for subpersonal indicators of person-level activities and states, and we have given a few examples of more or less reliable subpersonal indicators for person-level activities and states. This means there is good reason to be both excited and frightened about the potential applications of this knowledge in practical human affairs, such as the law, law enforcement, and in military campaigns. In our view, a priori Wittgensteinian arguments are irrelevant and distracting. The question of whether neuroscience can inform the law and legal theory cannot properly be addressed from the comfort of a philosopher's (or lawyer's) armchair. Instead, answering this question requires knowledge of how the brain works, of the strengths and weaknesses of neuroscientific techniques, and of the principles of statistical inference.

Here are some of the issues most relevant to assessing the proposed applications of neuroscience to the law.

What do you want to know about the person? Do you need to tell whether a person can see, or do you need to know the precise details of some particular belief or store of knowledge? Do you want to learn what they are doing or what they know now, what they were doing or knew in the past, or what they will do or know in the future? In the context of application, how certain do you have to be that the measure is accurate? What are the risks of being wrong?

Under what conditions do you need to know it? Are you on a battlefield making time-urgent decisions, or are you in a quiet neuroscience laboratory with all the time you care to dedicate to answering the question? In those conditions, can you run properly controlled behavioral experiments? Have the techniques been tested in those conditions? Is the subject a willing participant or an opponent? Are the personnel delivering the test trained in its proper use? What constitutes proper use and how do you know the answer to that question?

What are the relevant subpersonal indicators in brain mechanisms? Are these subpersonal indicators idiosyncratic to particular subjects trained in particular environments, or do they generalize across subjects? Are there reliable generalizations about, for example, how patterns of activity in one region of the brain are related to patterns of activity in other areas of the brain, and are those generalizations at the appropriate grain to track the content that one wishes to assess? Do the subpersonal indicators remain constant over time in the same subject or do they evolve with experience? Are the relevant subpersonal indicators localized, or are they distributed in the brain? To what extent do the relevant activities and states in brain mechanisms change across task situations or the context in which the measures are taken?

What are the available techniques for measuring subpersonal states and activities? How reliable are they at detecting brain indicators? How frequently do they generate false positives and false negatives? How precisely can they measure these states? Do they measure the brain at a resolution appropriate to the question you are trying to answer? Are there significant confounds? Are the measures taken by competent professionals (who know how to avoid the confounds)?

How reliable is the relationship between subpersonal indicators and person-level states? What is required of a properly controlled task condition for the technique in question? Have the indicators been tested on content of the sort sought in the current application? Have the indicators been tested in the conditions of application, that is, on this type of content and in this type of context? Is the subject using countermeasures, and how effective can the countermeasures be? What other information is available to guide task selection, to generate baseline data, and to assess significant differences?

This is just a brief survey of scientific and technological questions that law and society will have to face with the inevitable advance of neuroscience. There are, in addition, a host of philosophical questions about how the personal and subpersonal levels intermingle with one another and about the implications of this intermingling for our commonsense thoughts about agency, free will, and moral responsibility.

If we are to negotiate our way to reflective equilibrium between the empirical findings of neuroscience and the practical needs of the law, then we must be guided to the extent possible by a detailed understanding of the neuroscience in question, its strengths, and its limits. This can be done only if legal scholars understand how the brain works and, more importantly, understand the reliability and validity of available techniques for measuring the brain states necessary for person-level abilities. Pardo and Patterson have shown just how irrelevant the metaphysics of mind is to the challenges raised by the inevitable application of neuroscience to practical human affairs. No-nonsense neurolaw can proceed full steam ahead at the subpersonal level.

References

1. Wittgenstein, L. 1953. *Philosophical investigations*. Oxford: Basil Blackwell.
2. Malcom, N. 1977. *Memory and mind*. New York: Cornell University Press.
3. Bennett, M., and P.M.S. Hacker. 2003. *Philosophical foundations of neuroscience*. Oxford: Basil Blackwell.
4. Reid, T. 1785/1941. In *Essays on the intellectual powers of man*, ed. A.D. Woozley. London: MacMillan and Co.
5. Dennett, D. (1969). The ascription of content. In *Content and consciousness*, 72–96. New York: Routledge.
6. Hornsby, J. (1997). Introduction: Personal and subpersonal levels. In *Simple mindedness: In defense of naïve naturalism in the philosophy of mind*, 157–167. Cambridge, MA: Harvard University Press.
7. McDowell, J. 1994. The content of perceptual experience. *The Philosophical Quarterly* 44: 190–205.
8. Cowey, A., and V. Walsch. 2000. Magnetically induced phosphenes in sighted, blind and blindsighted observers. *NeruoReport* 11(14): 3269–3273.
9. Barker, A.T., R. Jalinous, and I.L. Freeston. 1985. Non-invasive magnetic stimulation of human motor cortex. *The Lancet* 1: 1106–1007.
10. Penfield, W. 1967. *The excitable cortex in conscious man*. Liverpool: Liverpool University Press.
11. Hume, D. 1772/1993. *An enquiry concerning human understanding.* Indianapolis: Hackett Publishing Company.
12. Monti, M. M., A. Vanhaudenhuyse, M. R. Coleman, M. Boly, J. D. Pickard, L. Tshibanda, A. M. Owen, and S. Laureys. 2010. Willful modulation of brain activity in disorders of consciousness. *New England Journal of Medicine,* February 3, 2010 at www.njem.org.
13. Martin, C.B., and M. Deutscher. 1966. Remembering. *Philosophical Review* 75: 161–196.
14. Bernecker, S. (2010). *Memory: A philosophical study.*
15. Thompson, R.F. 2005. In search of memory traces. *Annual Review of Psychology* 56: 1–23.
16. Wilson, M.A., and B.L. McNaughton. 1993. Dynamics of the hippocampal ensemble code for space. *Science* 216: 1055–1105.
17. Diba, K., and G. Buzsaki. 2007. Forward and reverse hippocampal place-cell sequences during ripples. *Nature Neuroscience* 10: 1241–1242.
18. Shema, R., S. Hazvi, T.C. Sacktor, and Y. Dudai. 2009. Boundary conditions for the maintenance of memory by PKM in the neocortex. *Learning & Memory* 16: 122–128.
19. Polich, J., and A. Kok. 1995. Cognitive and biological determinants of P300: An integrative review. *Biological Psychology* 41: 103–146.
20. Farwell, L.A., and S.S. Smith. 2001. Using brain MERMER testing to detect concealed knowledge despite efforts to conceal. *Journal of Forensic Sciences* 46: 135–143.