NINE NEUROLAW PREDICTIONS

*Morris B. Hoffman**

A sitting trial judge, and member of the MacArthur Foundation's Research Network on Law and Neuroscience, makes short-term, long-term and "never happening" predictions about the impacts neuroscience will have on law.

Keywords: *neurolaw, law and neuroscience, predictions, brains, responsibility, right and wrong, will, intention*

Predictions about the impact neuroscience will have on law seem to be grouped around two extremes: it will have no impact or it will have broad and paradigm-shifting impacts. In my experience, neuroscientists tend toward the latter view, I think in part because they underestimate the difficulty of moving the battleship that is the law. Some also fail to appreciate that the law represents thousands of years of pretty good accumulated folk psychology about the human condition. Some legal scholars, by contrast, tend toward the opposite, overly pessimistic, extreme. I think many of them overestimate both the law's inertia and its pedigree. Some no doubt under-appreciate the extraordinary advances made by neuroscience in the last few decades.

In this essay I will try to stake out some middle ground. In the end, I think neuroscience's legal impacts in the next 50 years will likely be rather

---

*District Judge, State of Colorado, Second Judicial District (Denver), morris.hoffman@judicial.state.co.us; Member, John D. and Catherine T. MacArthur Foundation's Research Network on Law and Neuroscience; Research Fellow, Gruter Institute for Law and Behavioral Science; Adjunct Professor of Law, University of Denver. I want to thank my friend and collaborator Frank Krueger, a social neuroscientist at George Mason University, for his helpful comments on earlier versions of this essay. The views expressed in this essay do not necessary reflect the views of my judicial or academic colleagues.

New Criminal Law Review, Vol. 21, Number 2, pps 212–246. ISSN 1933-4192, electronic ISSN 1933-4206. DOI: https://doi.org/10.1525/nclr.2018.21.2.212.

lumpy, with some significant, but not paradigm-shifting, impacts in a few discrete areas and not much impact anywhere else. I divide my nine predictions into three temporal segments: near term (next 10 years), long term (10 to 50 years in the future), and never happening:

1. In the next 10 years, neuroscience will be able to detect chronic pain accurately and reliably, and may even be able to distinguish chronic pain from malingering. This will have significant impacts on tort and disability law.
2. In the next 10 years, neuroscience will be able to diagnose many legally relevant psychiatric conditions, including several that often bear on criminal competence and insanity (such as severe schizoaffective disorder), as well as mental conditions claimed to have been caused by torts or justifying disability payments (such as PTSD). This development will not have much legal impact because most legally relevant mental conditions will continue to be diagnosed by traditional clinical methods, but it could help when experts disagree on their diagnoses.
3. Ten to 50 years from now, neuroscience will develop accurate and reliable lie detection methods, which could in theory have significant and widespread impacts for both the criminal and civil systems. Most likely, however, this development will not have significant impacts in the courtroom because the law will continue to be resistant to the admissibility of lie detection results. But it will have significant impacts pre-trial.
4. Ten to 50 years from now, neuroscience will show that some drugs seriously affect our ability to assess risk but not our ability to form intentions, and for those drugs the law may reverse its age-old rule that voluntary intoxication is a defense to purposeful crimes but not to knowing or reckless ones.
5. Ten to 50 years from now, neuroscience will demonstrate that there is no distinction between the mental states of "knowing" and "reckless" when it comes to results elements, and the criminal law may abandon the distinction.
6. Ten to 50 years from now, neuroscience will develop accurate and reliable ways of detecting autobiographical memories for faces and places. This development could theoretically have significant and widespread impacts on both civil and criminal law, but likely won't have much courtroom impact because it is a form of lie detection to which the law of evidence will continue to be resistant.

7. Ten to 50 years from now, neuroscience will be able to determine, on an individual basis, how mature a brain is, dispensing with the law's need to draw some age-based lines for things like criminal responsibility or the age of consent.
8. Neuroscience will never completely solve the mystery of addiction. Even though it may solve the puzzle of tolerance, it will not solve the riddle of dependence.
9. Neuroscience will never convince the law to abandon notions of free will or responsibility.

I make these predictions with great trepidation, fully aware not only of the special kind of schadenfreude in which we all seem to delight when our fellow humans get predictions wrong, but also of the sad fact that so many predictions, even by (maybe especially by) well-informed people, end up being spectacularly wrong.[1] I nevertheless make the effort because, as this special themed issue demonstrates, neuroscience is already having impacts on the law. Anything we can do to prepare broadly for future impacts seems to me to be an effort worth making.

My approach is limited in three ways. First, I am considering the neuroscience of law and not the law of neuroscience. By that I mean that my predictions focus on how neuroscience may change the way the law regulates social relations and the processes it uses to do so, not on special legal issues those applications may raise. Thus, for example, I consider how lie detection by neuroimaging may affect trial processes, not whether it will violate the Fifth Amendment to conduct nonverbal lie detection on a criminal suspect.[2] Of course, the extent to which neuroscience will impact the law will necessarily be affected by the extent to which legal principles will permit that impact. But I've chosen this path because my main focus in this

---

[1] Famously wrong scientific predictions include Albert Einstein's 1932 prediction that nuclear energy would never be obtainable because the atom could not be reliably split, and the *New York Times*' 1936 prediction that rockets will never be able to leave the Earth's atmosphere. Wrong predictions are not limited to science. *Variety* magazine predicted in 1955 that rock and roll "will be gone by June." Wrong predictions are also not restricted to the dim past. Readers may remember who most pundits thought would win the 2016 U.S. presidential election.

[2] For example, by showing a suspect a picture of a crime scene and determining, by neuroimaging alone, whether he was at that scene. See the discussion of the detection of autobiographical memories below.

essay is to inform legal readers about neuroscience's legally relevant trajectory, not scientists about law.

Second, I am not covering any issues about neuroprosthetics or the brain/machine interface. These developments may have deep and wide impacts across all of society, and the law will surely not be immune to them. But I limit this essay to neuroscience developments that I think will likely be primarily important to law.

Third, I do not consider the impact neurolaw will have in the legal academy. It is already having some impact. Several law schools offer joint degree programs in law and neuroscience, and many more regularly offer courses in law and neuroscience. Whether law and neuroscience will reach the academic heights that law and economics has reached remains to be seen. But that's not a question I address in this essay. I am more interested in neuroscience's impacts on law than on its impacts on the study of law, though I certainly recognize that the academy can be an important engine of legal change.

There is one prediction readers can take to the bank, so let's get that out of the way now: lawyers, judges, and, well, everyone will continue to misuse neuroscience to reach conclusions that the neuroscience simply does not warrant.

Criminal defense lawyers will continue to conflate cause with excuse, and will continue to drift toward the ideal they seek but will never admit: no one is ever responsible because their brains made them do it.[3] Prosecutors will continue to misuse neuroimaging by claiming that the lack of any positive findings indicates a mentally healthy defendant—a claim, given the current spatial resolutions of PET and fMRI, akin to saying that the world is at peace based on images of earth taken from the moon. Civil and criminal lawyers, on both sides of the fences, will continue to rely on neuroimaging as a kind of poor man's diagnostic: something seems to be wrong with this person's brain, here's even a neuroimage showing something is wrong, and this something wrong should therefore hint at some otherwise insupportable legal conclusion, like this person is or is not disabled and therefore is or is not entitled to disability benefits, or this person is or is not criminally responsible.

Some judges, too, will continue to use neuroscience to reach results they wanted to reach in the first place, whether that's to outlaw capital

---

[3] See the discussion of free will in a later section.

punishment and life without parole for 17-year-old murderers,[4] or to sentence psychopaths less harshly once judges learn that psychopathy seems to be a problem with the brain.[5]

---

[4] Roper v. Simmons, 543 U.S. 551 (2005) (death penalty for 17-year-old murderers violates Eighth Amendment); Graham v. Florida, 560 U.S. 48 (2010) (life without parole for 17-year-olds in non-homicide cases violates Eighth Amendment); Miller v. Alabama, 567 U.S. 460 (2012) (life without parole for 17-year-old murderers violates Eighth Amendment). Although Justice Kennedy's majority opinion in *Roper* did not cite to any neuroscience, several amicus briefs discussed developmental evidence that the prefrontal cortex in the average male is not fully developed until the mid-twenties. *See, e.g.*, Amicus Brief for American Psychological Association and Missouri Psychological Association at 4 (No. 03–633), *available at* http://www.apa.org/about/offices/ogc/amicus/roper.pdf, citing the work of two of my MacArthur colleagues: Laurence Steinberg & Elizabeth S. Scott, *Less guilty by reason of adolescence: Developmental immaturity, diminished responsibility, and the juvenile death penalty*, 58 AM. PSYCHOL. 1009–18 (2003). Justice Kennedy did cite to this science in his majority opinion in *Graham*, 560 U.S. at 69. This developmental fact has since become a common talking point among those pushing to raise either the age for death-eligibility or the age for adult criminal responsibility generally. *See* Steinberg & Scott, *supra*. It has certainly made the rounds of public defender training programs. My public defenders regularly argue that I should sentence young offenders less harshly because their prefrontal cortices are not fully developed. But of course there is a huge base rate gap between the proposition that some defendants don't have fully developed brains and the proposition that those defendants are less responsible. When I teach these cases, I like to ask my students how many of them have been 17 years old, then I ask them to keep their hands up if they burglarized a home, kidnapped the woman there, wrapped her up in duct tape and a carpet, put her in the trunk of her own car, drove her to a bridge, and threw her into the river below while she was still alive. That's what Christopher Simmons did. Of course, these kinds of legal and normative questions become a matter of line drawing—without sufficiently reliable neuroscience on an individual basis, lines have to be drawn that will necessarily preclude the death sentence for some juveniles who are fully mature and responsible as the price of not executing those who are not. But as the science gets better, we may be able to do away with some of this line-drawing. See my Seventh Prediction, discussed below.

[5] Researchers have shown that judges sentence psychopaths much more harshly than non-psychopaths. L. Aspinwall et al., *The double-edged sword: Does biomechanism increase or decrease judges' sentencing of psychopaths?*, 331(6096) SCIENCE 846–49 (Aug. 17, 2009). No surprise there. Psychopaths are substantially more likely to recidivate than non-psychopaths and are famously immune to any kind of treatment—two traditional sentencing considerations. But these same researchers found that judges reduced these harsh sentences once they were told psychopathy has a neurobiological explanation. *Id.* This is part of what my friend Stephen Morse has so cleverly dubbed "the psycho-legal error"—the conflation of cause with excuse. Stephen J. Morse, *Excusing and the New Excuse Defenses: A Legal and Conceptual Review*, 23 CRIME & JUST. 329, 350–53 (1998).

The media will continue to hyperventilate regularly about how neuroscientists have discovered the brain's "love center," "racism center," or other centers *du jour*. These pronouncements will continue to obscure neuroscience's real modern story—that it is becoming clearer and clearer that the brain is substantially more interconnected than anyone had ever imagined. It is a gross caricature to label any brain area the "center" of any complex behavior.[6]

Finally, modern snake oil salesmen will continue to claim to be able to cure this or that brain problem by treating, or "training," this or that brain center. I was once on a panel presenting some neuroscience to a conference of the Ninth Circuit, and in a later session a non-practicing J.D. presented training exercises about how to get rid of implicit bias by "retraining your amygdala." We were all flabbergasted. After spending an hour trying to educate these federal judges about how cautious we all need to be about the intersection of law and neuroscience, here was this charlatan trying to sell her phony training services. Ah, commerce![7]

Before I discuss my nine predictions, it may be worth pausing to remember that the marriage of law and neuroscience is part of a larger tableau animated by the two Big Questions humans have been asking ourselves since our emergence: Are we fundamentally good or fundamentally bad? and Do we have the free will to choose between the two?[8] I want to take a moment to summarize humankind's Big Guesses about these Big Questions, not because I have anything to add to them, but because much of what has been written about neurolaw touches on, and in my view sometimes gets confused by, those Big Questions and our Guesses about them. I won't dawdle, and less philosophically inclined readers may wish to skip the next two sections and head straight to my summary of neuroscience trends and the predictions flowing from them.

---

[6] To give readers just a small glimpse into the interconnectedness of the human brain, a single cortical neuron is connected to an average of 1,000 to 10,000 other neurons. Not at all like that drawing of a neuron you remember from high school biology, that had a few upstream and downstream connections! With 80 billion neurons in the average human brain, that's a lot of connections and a lot of complexity.

[7] There are many other examples, from companies peddling brain scans they claim can tell what's wrong with your teenager's brain, to ones selling "foolproof" neuro lie detection.

[8] There is of course a third Big Question: What, exactly, is right and wrong?

I. SOCRATES, THRASYMACHUS, AND ALL THOSE OTHER GUYS

Grand theories about human nature can and do affect the lives we live and the institutions we construct, including our legal institutions. More than 2,000 years ago Socrates and an Athenian sophist named Thrasymachus got the ball rolling by engaging in a famous debate about human nature and the meaning of justice.[9] We've been having the same debate ever since. Are humans fundamentally good, built to appreciate beauty, truth, and cooperation, and therefore in need only of modest intervention by the state? Or are we fundamentally bad, built only to maximize our self-interest and therefore in need of heavy-handed restraint by a robust state? Is justice "an expression of the excellence of the soul," as Socrates argued,[10] or nothing more than "the interest of the stronger," as Thrasymachus put it?[11]

The answers we have given to these questions have in large measure defined our political and legal institutions. It is no coincidence that America's founders were steeped in the Enlightenment's decidedly mixed version of this debate. The Constitution's distribution of power between different branches was a reflection of the framers' nuanced views about human nature. We are good enough to be free of an overbearing Leviathan, but not quite good enough to populate our limited government without checks and balances between its parts.

Most modern and certainly post-modern takes on human nature have been distinctly less balanced, skewed heavily toward Thrasymachus' dreary views. The central simplifying assumption of classical economics was that each of us is relentlessly self-interested. Markets are an efficient integration of all those individual greedy unseen hands. Darwin's insights strengthened the belief that we are self-interest machines, and biology's great synthesis of evolution and genetics simply moved the locus of that self-interest from the selfish individual animal down to the animal's selfish genes. Freud did much the same for the psyche, and Marx for the allegedly relentless march of economic and political history framing the wars between self-interested classes.

---

[9] This debate was recorded (or invented) by Plato, and is reported in Book I of his *Republic*. PLATO, THE REPUBLIC 272–97 (Amazon Classics 2017, trans. B. Jowett).

[10] *Id.* at 296.

[11] *Id.* at 274.

But then something delightful happened on this dark road to modern pessimism. Some economists, anthropologists, and biologists had the audacity to look at how humans actually behave instead of how these dreary theories predicted we should behave. And there were many surprising results.

Experimental economists discovered that when we play economic games in the laboratory, we engage in all kinds of cooperative behaviors that cannot be explained by classical economics, including sharing with and trusting other players, even when they are strangers.[12] Behavioral economists added the observation that we are not only not rational self-interest machines, but that our irrationality is predictable in many important decision-making domains.[13]

Anthropology and biology contributed important pieces to this emerging picture of an animal subtly torn between selfishness and cooperation. Anthropologists discovered that human societies have vastly more in common with each other than previously believed, and in fact that many of these commonalities involve social norms that reflect and encourage cooperative behaviors, and that punish antisocial ones.[14] Biologists have long known that this kind of cooperation—between and within species—is not unique to humans, but is widespread throughout the animal kingdom.[15]

In fact, because within-species cooperation is so common—individuals even sacrificing their own lives to save other family members—biologists first dubbed this the "puzzle of altruism." But they now realize it's no puzzle at all. There is an accepted evolutionary explanation for these self-sacrificing and cooperative behaviors, the gist of which is that cooperating sometimes gave us long-term fitness advantages that exceeded its short-term costs.[16] This is especially true with intensely social animals, like humans.

---

[12] These economic games include the One-Shot Ultimatum Game, Prisoner's Dilemma, Stag Hunt Game, and Public Goods Game. For a more detailed explanation of these games and their contribution to this rediscovery of the cooperative side of human nature, see MORRIS B. HOFFMAN, THE PUNISHER'S BRAIN: THE EVOLUTION OF JUDGE AND JURY 18–25 (Cambridge Univ. Press 2014).

[13] *See, e.g.*, R. THALER, MISBEHAVING: THE MAKING OF BEHAVIORAL ECONOMICS (W.W. Norton, 2015); D. Kahneman & A. Twersky, *Prospect Theory: An Analysis of Decision under Risk*, 47(2) ECONOMETRICA 263–91 (Mar. 1979).

[14] DONALD E. BROWN, HUMAN UNIVERSALS (McGraw-Hill 1991).

[15] LEE ALAN DUGATKIN, COOPERATION AMONG ANIMALS: AN EVOLUTIONARY PERSPECTIVE (Oxford Univ. Press 1997).

[16] W.D. Hamilton, *The genetical evolution of social behavior*, 7 J. THEOR. BIOL. 1–52 (1964); R.L. Trivers, *The evolution of reciprocal altruism*, 46 Q. REV. BIOL. 35–57 (1971).

We are still self-interest machines, it's just that the way evolution solved the survival problem for us, and for many other animals, was to take advantage of the fitness benefits that come from working together. For us those benefits were significant—everything from mutual defense, to group hunting and gathering, to divisions of labor. But they were not so overwhelming to turn our predisposition to cooperate into a command, as with the social insects. After all, evolution is about *individual* fitness, not group fitness. So we cooperate with other group members because that gave us a long-term fitness advantage over individuals not part of any group, but we also cheat because that can give us a short-term fitness advantage over other cooperating members.

Evolution built this profound tension into our brains. We are presumptive, but grudging, cooperators. We have built-in prosocial moral intuitions that we tend to follow—don't steal (property or well-being) from one another and don't break promises—but that we are also willing to jettison if the short-term benefits of doing so are tempting enough. This new evolutionary perspective on human nature harkens back to our more nuanced Enlightenment views.

Neuroscience is both contributing to and being affected by this rediscovered view of human nature. A raft of neuroscience work is now being done, much of it discussed in the predictions sections below, exploring the way our brains accomplish various legally relevant social tasks, including having moral intuitions (or in the case of psychopaths, not having them), lying, forming intentions, assessing risks, judging the intentions and risk-taking behaviors of others, blaming, punishing, and forgiving. Much of this work is being done in a new subspecialty called social neuroscience. Many neuroscientists are also looking at these issues in a developmental context. They are asking themselves what brain regions, networks, and subnetworks are required to accomplish these tasks, when in our development we acquire this neural hardware, and what legal implications these development horizons might have, especially in juvenile law.

Hamilton showed that traits, both behavioral and physical, could be selected for even though they made an individual less fit, as long as they made the individuals' genes for that trait—shared by relatives—more fit. Trivers showed that this phenomenon could be generalized beyond relatives, and indeed even beyond members of a single species (some fish, for example, refraining from eating so-called "cleaner fish" because the long-tern benefits of being cleaned outweigh the short-term benefits of eating the cleaners).

II. OMG, NOT FREE WILL AGAIN

The second Big Question derives from the first. If our natures are painted both with a broad predisposition to cooperate and with a sharp willingness to cheat, what mechanisms drive those different behaviors when the rubber meets the road of an actual moral decision? Do we freely choose to cooperate or cheat, and are we therefore morally responsible agents? Or is it neurons all the way down, our "decisions" no more freely taken than the "decision" of a ball to respond to gravity? Or is it something in between?

Readers who think this kind of freshman-bull-session inquiry is too abstract to have anything practical to do with either law or neuroscience may be surprised to learn that it seems to have been an important animating factor in the MacArthur Foundation's 2007 foray into law and neuroscience.[17] When the Foundation asked its "genius grant" recipients for big ideas that could have big public impacts, Robert Sapolsky, a primatologist and neuroscientist at Stanford, responded that neuroscience could, and should, do away with the criminal law as we now know it. As Sapolsky has put it, our actions are no more the product of some mysterious force called "will" than a car wills its brakes to fail.[18]

---

[17] The John D. and Catherine T. MacArthur Foundation awarded a $10 million grant for the study of law and neuroscience in 2007. That first effort, dubbed "The Law and Neuroscience Project," was headed by Michael Gazzaniga, generally considered the father of cognitive neuroscience. There were around 50 of us involved in that first iteration, divided roughly equally between neuroscientists and criminal law professors, with a smattering of philosophers and statisticians. Five judges were part of that first effort, and I was lucky to be among them. Retired Justice Sandra Day O'Connor was an honorary member of the governing Board. Jed Rakoff (S.D.N.Y.) was also on the Board and was an active participant. The other judge-participants were Gerard Lynch (2d Circuit) and William Fletcher (9th Circuit). For a history of that first effort, go to the Research Network's current website at www.lawandneuro.org, click on "About Us," and then on the dropdown menu click on "History."

In 2012, the Foundation awarded a second grant, and dubbed its participants "The Research Network on Law and Neuroscience." This was a much smaller effort. There are 13 of us in the Network, again divided roughly equally between neuroscientists and criminal law professors. I was again lucky enough to be one of the two judge-members, along with former judge Andre Davis (6th Circuit).

[18] R. Sapolsky, *The frontal cortex and the criminal justice system*, 339 PHIL. TRANS. B 1787–96 (2004). To be fair, Professor Sapolsky also makes the point that although this "neurobiological framework," as he calls it, would eliminate blame and responsibility, it would not eliminate intervention by the state against criminals for consequentialist reasons.

Many neuroscientists share this kind of extreme determinism and deep skepticism about free will.[19] Philosophers, by contrast generally divide into three camps: a few are hard determinists (no free will, no responsibility), a few are libertarians (free will and responsibility), but most are soft determinists, also called compatibilists (no free will, but responsibility). Most philosophers, especially legal philosophers, also reject the proposition that neuroscience will one day prove, or otherwise justify, hard determinism.[20]

Regardless of one's views on the Big Question of free will, we can probably all agree that this debate can sometimes feel like a time-wasting detour. My friend and MacArthur colleague Owen Jones put it this way:

> The problem with free will is that we keep dwelling on it. Really, this has to stop. Free will is to human behavior what a perfect vacuum is to terrestrial physics—a largely abstract endpoint from which to begin thinking, before immediately moving on to consider and confront the practical frictions of daily existence.[21]

I tell my law and neuroscience students to think of the debate about free will as a giant intellectual black hole—something we should always be aware of as we saddle up to questions like the meaning of pain or the nature of intention, but something we can't get too close to for fear of getting swallowed up. I will try to live by these words of cautious neglect throughout the rest of this essay, though I do admit I'll be getting uncomfortably close to the black hole, and maybe revealing some of my own free will biases, in the "never happening" predictions set forth below.

As he put it: "To understand is not to forgive or to do nothing; whereas you do not ponder whether to forgive a car that, because of problems with its brakes, has injured someone, you nevertheless protect society from it." *Id.* at 1794.

[19] Sapolsky, *supra* note 18; Jerry A. Coyne, *You Don't Have Free Will*, CHRONICLE REVIEW, Mar. 23, 2012, at B6; J. Greene & J. Cohen, *For the Law, Neuroscience Changes Nothing and Everything*, 359 PHIL. TRANS B. 1775–85 (2004) (arguing that "determinism really does threaten free will and responsibility as we intuitively understand them [but] most of us, including most philosophers and legal theorists, have yet to appreciate it."). A more traditionally compatibilist argument has been made by another renowned neuroscientist, Michael Gazzaniga, *Free Will is an Illusion, but You're Still Responsible for Your Actions*, CHRONICLE REVIEW, Mar. 23, 2012, at B7.

[20] *See, e.g.*, Hilary Bok, *Want to Understand Free Will? Don't Look to Neuroscience*, CHRONICLE REVIEW, Mar. 23, 2012, at B8; Stephen J. Morse, *Criminal Law and Common Sense: An Essay on the Perils and Promise of Neuroscience*, 99 MARQ. L. REV. 39 (2015).

[21] Owen D. Jones, *The End of (Discussing) Free Will*, CHRONICLE REVIEW, Mar. 23, 2012, at B9.

III. ADVANCING NEUROSCIENCE

Most of what's been learned in the last 30 years about how the brain works comes directly from advances in neuroimaging.[22] Before the emergence of positron emission tomography (PET) in the 1980s and functional magnetic resonance imaging (fMRI) in the 1990s, there was no reliable way to measure the brain activity of live experimental subjects systematically. X-rays and CT scans were increasingly good at detecting brain structures, but what those structures did and how they worked together remained shrouded in mystery, penetrated only by occasional insights from disease or trauma.

Animal studies—in which electrodes are inserted directly into animal brains to measure neuronal activity—contributed and continues to contribute greatly to scientists' knowledge about reflex and motor pathways.[23] But of course these methods cannot be used on healthy human subjects. And animals presumably don't have sophisticated beliefs about legally relevant propositions, like whether to punish or forgive, and even if they did, those beliefs cannot be examined through language-directed experiments.

PET and fMRI allowed neuroscientists, for the first time, to examine what human brains are actually *doing* when they engage in complex behaviors. Almost entirely because of these two technologies, we've learned more about the human brain in the last 30 years than we learned in all of previous human history. That's a rate we need to keep in mind when we hazard guesses about the future impacts of neuroscience. And it's a rate that seems not only sustainable but likely accelerating.

Two main engines are driving this acceleration: (1) new imaging technologies, and (2) increasingly sophisticated methods of collecting and analyzing data from those new technologies.

---

[22] The "imaging" part of "neuroimaging" is a bit misleading. None of the technologies used to measure brain activity delivers an "image" of the brain in the sense of taking the experimenter inside the brain to record what could be seen there visually. Instead, these technologies typically measure brain metabolism by measuring indirect physical effects of that metabolism, such as blood oxygenation or changes in glucose.

[23] Technological advances are also happening in animal neuroscience. Late last year scientists announced they had developed a new kind of mobile animal brain probe that has 384 separate channels on its tiny surface, and that can therefore measure 384 separate neurons simultaneously (up from the previous maximum of a few dozen neurons per probe). J. Jun et al., *Fully integrated silicon probes for high density recording of neural activity*, 551 NATURE 232–36 (Nov. 9, 2017).

There is a veritable alphabet soup of new technologies that neuroscientists have added to their experimental repertoires since the invention of PET and fMRI, and it seems new ones get added every year.[24] These technologies have acronyms like MEG (magnetoencephalography), QUEEG (quantitative electroencephalography), fNIRS (functional near-infrared spectroscopy), SPECT (single positron emission tomography), TMS (transcranial magnetic stimulation), DOI (diffuse optical imaging), EROS (event-related optical signal), DBS (deep brain stimulation), and DTI (diffusion tensor imaging). Each of these technologies has its own limitations,[25] and neuroscientists often combine different technologies in a single experiment to minimize these limitations.

Along with PET and fMRI, these different technologies are all contributing to an exploding understanding of how brains work. There's no apparent reason to expect this pace of technological innovation to slow. In the next 50 years scientists may even invent entirely new ways to measure brain activity that will be just as revolutionary as PET and fMRI.

These advances are proceeding along several different lines, four of which I'd like to highlight: connectivity, resting state studies, causal technologies, and mobile imaging.

Connectivity technologies allow experimenters to measure how neurons are connected to one another. Neurons are not in fact physically

---

[24] There are two basic types of approaches in neuroscience research, correlational and causal. All imaging technologies are correlational. They take physical measurements of things scientists believe are correlated with brain activity (like glucose in the case of PET, oxygenated hemoglobin in the case of fMRI, or the electromagnetism being released by firing neurons in the cases of EEG and QEEG). Causal technologies (like transcranial magnetic stimulation) actually affect parts of the brain, allowing scientists to observe how these effects impact brain function.

[25] The two most significant limitations are spatial resolution and temporal resolution. The best spatial resolutions, achieved by fMRI, are currently around 3 mm^3, which means that the measurement of brain activity represented in a single voxel of an fMRI neuroimage is actually an average coming from 3 mm^3 of brain, a volume roughly equal to a small pea. If that volume is in the cortex, it contains around 50,000 neurons, so the signal coming from that section is actually an average of the signals coming from 50,000 neurons. Other imaging technologies have much lower spatial resolutions. On the other hand, fMRI suffers from relatively low temporal resolutions—meaning that there is a significant delay (around 1/2 second) between brain activity and the metabolic trace of that activity read by fMRI. That delay, called the hemodynamic delay, comes from the fact that once the brain calls for blood, it takes the blood some amount of time to reach its target. Other imaging technologies have much better temporal resolutions.

connected; there is a tiny space between them called the synapse. (So scream out if anyone ever says our brains are "hard-wired.") Structural connectivity methods like DTI inferentially measure which neurons share synapses. Functional connectivity inferentially measures which neurons are actually firing across those synapses when a particular task is being done in the scanner.

These connectivity methods allow neuroscientists not only to explore which brain areas are involved in certain behaviors, but to begin to map out how networks and subnetworks in those areas actually work with one another. They are headed toward addressing one of neuroscience's biggest challenges: learning about the brain's organization at the intermediate level. Much is known about large-scale brain regions and networks, and even more is known about how individual neurons work; it is at the intermediate level of organization that the brain remains largely a black box. Connectivity technologies promise to begin to open this black box.

A second important technological trend is that neuroscientists are starting to look at so-called "resting states." Traditional functional imaging experiments involve subjects performing tasks inside the scanner, then experimenters comparing the brain activities of subjects doing the tasks and subjects not doing the tasks. Subtracting these two sets of measurements should leave measurements of the areas associated only with the targeted task.

Neuroscientists have long realized that this traditional "subtraction" method involves a huge untested assumption—that the two sets of brains are sufficiently stable to allow the inference that the only differences between them must be attributable to the targeted task. Until recently, though, very little research attention has been devoted to this co-called resting state problem. That has started to change, and in a big way. The literature is now rife with resting state studies, or what are sometimes called "task-free" studies—neuroimaging of brains that are not doing any assigned tasks in the scanner. These studies are not only strengthening the inferences from traditional subtraction-type research, they are becoming their own significant research tool. Several of the predictions discussed below—including the idea that neuroscience may eventually be able to diagnose many psychiatric conditions—rely on the assumption that these resting state studies will continue to improve our knowledge about the neurological differences correlated to different psychological states.

The third technological trend involves causal technologies—methods that temporarily alter the brain. These technologies are used both clinically and experimentally. Some of them have become an elegant research tool to test hypotheses about whether particular brain areas are necessary (though of course not sufficient) for a given task. For example, transcranial magnetic stimulation (TMS) is a method by which neuroscientists can direct a magnetic beam at an area on or near the surface of the cortex to temporarily stimulate or deactivate that area. If subjects with that particular area deactivated cannot perform an experimental task, neuroscientists can be pretty confident that the deactivated area was necessary for the task.[26] Most of these technologies currently work by affecting fairly large areas very near the surface of the cortex. As these technologies improve, neuroscientists should be able to learn a lot more about the role smaller and deeper brain regions play in human decision-making.[27]

The last technological trend I want to discuss is that many of these other technologies are becoming mobile. There are already mobile EEG devices that allow experimenters to collect vast amounts of EEG data from subjects around the world, rather than from the handful of subjects who are able to come into a particular lab over a particular window of time. The same is true of fNIRS (functional near-infrared spectroscopy). If this trend continues, other neuroimaging technologies—either existing ones[28] or completely new ones—may open up neuroscience to the internet in much the same way that the behavioral sciences have already been opened up to the internet. It is now the rule, not the exception, that behavioral experiments, both in psychology and in economics, are done online rather than in the

---

[26] These causal experiments can be conducted "on-line" or "off-line," a reference not to the internet but to whether the task is performed while the brain-affecting technology is still being administered ("on-line") or during the temporary period after which the machinery is turned off but the subject's brain remains affected ("off-line").

[27] Drugs, as opposed to machines, can also be used to alter the brain in experimental contexts. For example, subjects given oxytocin, a neuromodulator associated with childbearing and lactation, become more trusting in economic games, and their brain regions associated with trust more active. M. Kosfield et al., *Oxytocin increases trust in humans*, 435 NATURE 673–76 (2005). Conversely, men with high testosterone levels are on average less trusting (and less trust-worthy) than men with lower testosterone levels. T. Burnham, *High testosterone men reject low ultimatum offers*, 274 PROC. ROYAL SOC. 2327–30 (2007).

[28] But probably not PET or fMRI, at least in their current forms. Fixed principles of physics require the large disruption and detection magnets that these technologies currently use.

laboratory (and, in fact, almost all of the behavioral studies about subjects' assessments of mental states discussed below were online studies). This has given these disciplines an extraordinarily cheap way to conduct experiments with a very large number of subjects, and there's no reason to believe that the same won't be happening in one form or another with some aspects of neuroscience.

The second driver of neuroscience innovation is not new data-gathering technology but rather new ways to measure the data collected. Beginning several years ago, neuroscientists started using the same kinds of "big data" software to examine brain data as retailers use to predict consumer preferences based on their buying, and even just browsing, histories. Instead of looking for particular brain regions or areas that seem to have increased activity associated with a particular behavior, and then constructing models (assisted by connectivity studies) of how each of those regions contributes to the behavior in question, many neuroscientists are now looking at whole brains. They are using deep learning software not to construct models of *how* brains reach decisions, but rather to discover patterns that distinguish brains that made one kind of decision or were exposed to one kind of stimulus from brains that made a different decision or were exposed to a different stimulus. Several of the mental state imagining studies discussed below use these big data methods, called collectively "multivariate pattern analysis" (MVPA).[29] Better and better data-crunching software is a virtual certainty in the future, and in fact one of the studies discussed below uses a new and improved kind of MVPA.

Because they look at patterns in whole brains without worrying about the details of causation, MVPA methods will likely be especially useful in the diagnosis of some mental conditions, as discussed below.

Now, to the predictions.

---

[29] You might see slightly different phrases for this data-crunching, depending on the discipline. Statisticians call it multivariate regression analysis whereas some neuroscientists call it multi-voxel regression analysis or multi-voxel pattern analysis. But the idea is the same: the causal relationship between an outcome (in neuroscience, say, an observed experimental task) and the many variables that not only contribute to that outcome but that can also affect each other (the signals of metabolism coming from each voxel being measured by some neuroimaging technique) can be mathematically modeled and therefore predicted. The more data, the better the prediction.

IV. NEAR TERM (NEXT 10 YEARS)

In this section I predict that in the next 10 years there will be two neuroscience developments that will have important legal impacts, with the first one having significant impacts: (1) the detection of chronic pain, and (2) the diagnosis of some legally relevant mental conditions.

1. Chronic Pain

A huge problem in law is science's inability to measure pain. If you've ever complained to your doctor about pain, you were probably asked to fill out a pain diagram. In addition to asking patients to show on the diagram where the pain is, the diagrams also typically ask patients to rate the intensity of the pain from 1 to 10, 10 being the worst pain you can imagine.[30] This is of course not only a completely subjective inquiry, it is as much a measure of patients' imaginations as a measure of their pain.

Because pain is so utterly and irreducibly subjective, it is vulnerable both to exaggeration by strategic actors and to misjudgment by fact finders. And yet the problem of pain is ubiquitous in the law—from social security disability determinations to crime ("bodily injury" is defined in most jurisdictions to include pain) to, most prominently, torts. My guess is that there is nothing on the immediate neuroscience horizon that will have more significant and widespread impact on the law than the discovery of a reliable method to detect pain, especially chronic pain.

Clinicians traditionally distinguish acute pain from chronic pain temporally. Pain that lasts twelve weeks or more, when the source of the pain has healed, is typically classified as chronic. Neuroscience has known for quite a while that there is a fundamental neurological difference between acute pain and chronic pain. Acute pain comes from the injury site in a process called "nociception," which is a fancy way of describing the relays of neurons that tell the brain that the injury site has been damaged. In fact, the acute pain message can be carried by nociceptors that have different transmission speeds, which probably explains the difference between the sharp pain we feel at the time of the injury and the more throbbing pain we feel later.

---

[30] They sometimes also ask more qualitative, but nevertheless still mostly subjective, questions such as whether the pain is "numbness," "pins and needles," "burning," "aching," or "stabbing."

But chronic pain is not the product of any nociception; that is, no pain messages are coming from the site of the injury or indeed from anywhere in the peripheral nervous system. The chronic pain message originates in the brain itself. Of course, that is not to say that it is not real pain, or is necessarily hysterical or malingering in origin. In fact even acute pain is, in the end, the brain's interpretation of the nociceptive signals it is getting from the peripheral nervous system, and therefore even acute pain is in large part a mystery of the central nervous system. That's one reason pain—both acute and chronic—is so wildly variable between individuals. It's also why the context of an injury can affect the amounts of pain we perceive.

Because acute pain starts with nociception, methods to detect nociception—like nerve conduction studies—can detect and rule out acute pain, and can even sometimes pinpoint its sources.

In the case of chronic pain, by contrast, the acute pain signal has stopped, there is no nociception going on, and yet the brain is still signaling pain to its user. For no other reason than that legal proceedings typically begin, and certainly end, long after the acute pain period is over, it is chronic pain that the law is most commonly called upon to evaluate. Today, there simply is no reliable way to detect, let alone quantify, chronic pain, and therefore no way to distinguish it from bald-faced lying or, more subtly, from exaggeration or even self-deception. Instead, fact finders have to look at a combination of the sufferer's testimony (do they believe it?—see lie detection section below) and other indirect evidence, like the sufferer's behavior before and after the claimed onset of chronic pain.

But neuroscientists are getting close to being able to reliably detect chronic pain—and perhaps even to distinguish it from malingering. Using MVPA techniques, they can now detect some kinds of chronic pain by neuroimaging alone, at accuracy levels that are already approaching 80 percent.[31] These techniques and their accuracies will almost certainly improve, and it therefore seems inevitable that in the next 10 years, and probably much closer to 5 than 10, jurors and administrative law judges will be able to tell whether the plaintiffs and claimants before them are actually suffering chronic pain or are faking.

---

[31] *See, e.g.*, H. Ung et al., *Multivariate classification of structural MRI data detects chronic low back pain*, 24 CEREBRAL CORTEX 1037–44 (2012), reporting 76% accuracy in distinguishing between control subjects and subjects with chronic lower back pain.

Neuroscience may even be able to tell the difference between real chronic pain and exaggerated pain. My optimism about this prospect comes from a 2013 paper in which experimenters, again using MVPA, were able to distinguish between brains that were suffering experimentally induced acute pain from brains that were remembering such pain previously induced, at 96 percent accuracy.[32] If these methods can distinguish acute pain from the memory of acute pain, it seems plausible that similar methods might be able to distinguish real chronic pain from pain that has been consciously or even unconsciously exaggerated.

One cannot overestimate how such a technology would change both personal injury law and the administrative processes that determine whether claimants are disabled. Pain and suffering is a huge component of virtually every personal injury claim, often dwarfing every other category of damages. Disability claims are also often grounded not on the assertion of physical limitations but on limitations alleged to come from chronic pain. Accurate and reliable methods of separating real chronic pain from exaggerated or faked chronic pain would revolutionize both of these areas of law. Not only would false claims be detected early on, but eventually fakers would not even bring them, leaving more of the compensatory pie for the truly injured. I wager that this development would have a more beneficial impact on the tort system than all past tort reforms put together.

2. Legally Relevant Mental Conditions

Whole-brain MVPA techniques should also allow neuroscientists, some day quite soon, to be able to diagnose a host of psychiatric conditions, including several legally relevant ones, by imaging alone. Indeed, although at the moment no DSM conditions are diagnosable by neuroimaging alone,[33] neuroscientists are getting very close to being able to diagnose several, including some kinds of depression and antisocial personality disorder. There is every reason to believe that as these pattern classifiers are unleashed on more and more subjects, including in resting state studies, we

---

[32] T. Wager et al., *An fMRI-based neurologic signal of physical pain*, 368 NEW ENG. J. MED. 1388–97 (2013).

[33] Morse, *supra* note 20, at 64 and n.55.

will soon have a battery of neuroimaging tests that can accurately and reliably diagnose dozens of psychiatric conditions.[34]

These neurodiagnostic methods probably won't replace more traditional clinical methods, at least initially. I expect psychiatrists to use brain imaging as an additional tool to confirm their clinical judgments, though as time goes on this order of prominence could switch. Neuroimaging might be especially valuable in distinguishing between differential diagnoses.

The same can be said for any forensic uses. I doubt lawyers will rely only on brain imaging, but will use it in conjunction with other kinds of evidence, including neuropsychological testing. The two legal areas in which such diagnosis-by-neuroimaging will probably have its biggest impacts are competency and insanity. There are dozens of psychiatric conditions that can at their extreme render parties incompetent or insane. Neurodiagnosis of such conditions would be valuable to jurors and judges faced with deciding whether a particular legal actor is suffering, say, from paranoid schizophrenia with auditory hallucinations.

But I don't want to overstate the impact of neurodiagnosis. In competency proceedings—which arise much more frequently than the insanity defense—the outcome often turns not so much on the clinical diagnoses of the competing experts but rather on their judgments about the extent to which the diagnosed condition is in fact interfering with a defendant's ability to understand the charges against him and/or assist his counsel in defending against them. These are inquiries that are essentially legal in nature, not medical. Neuroimaging data that confirms a clinical diagnosis might be useful in the rare occasions when the experts disagree about a diagnosis. But it is unlikely to answer the beguiling question of how, once a defendant's mental condition is established, that condition affects his competency.

Neurodiagnosis will certainly be prevalent in the penalty phases of death penalty cases, where neuroimaging is already common because of the lure of the so-called Christmas tree effect.[35] In anything-goes death penalty

---

[34] That is not at all to say that the resulting increase in diagnostic power will necessarily lead to better treatment. In fact, because these diagnostics will likely be based on whole brain pattern analyses, it is unlikely they will lead in the near term to any insights about the neurobiological *causes* of any of these conditions. But they will be an important step in the long and so far unsatisfying journey toward grounding psychiatry in biology.

[35] The phrase "Christmas tree effect" is used to describe the claimed phenomenon that subjects (and therefore jurors and judges) are more likely to reach legal conclusions

litigation, where lawyers on both sides are desperate for anything to sway jurors, the visualization of a mental condition that neurodiagnosis will provide will be irresistible.

Diagnosing mental conditions will probably have its biggest impact in tort cases. Plaintiffs who claim, for example, that a tortious incident has left them with a controversial clinical diagnoses, like PTSD, will be able to be tested for that condition by neuroimaging. Ditto in any area of the law where a party's mental condition is at issue, including will contests and claims of lack of contractual capacity. Again, the neurotesting might not only confirm (or of course eliminate) a given clinical diagnosis, but it could also help resolve any clinical battle of the experts in this regard. As neurodiagnosis gets more common, baseline data could even assist finders of fact in deciding whether a particular mental condition was caused or exacerbated by a tort, or was fully pre-existing.

Finally, much as with the neurodiagnosis of chronic pain, a reliable neuroimaging method of diagnosing a host of mental conditions allegedly interfering with one's ability to perform the ordinary tasks of daily living will undoubtedly find its way into the administrative area, in social security disability cases and other disability-based eligibility determinations.

V. LONG TERM (10–50 YEARS)

There are five areas where I think neuroscience may have significant legal impacts in the long term: lie detection, intoxication as a defense, simplifying the law's traditional mental states, detecting autobiographical memories, and measuring individual brain maturity.

3. Lie Detection by Neuroimaging

Polygraphy has gotten a bit of a bad legal rap, owing mostly to the coincidence that it emerged in its modern form in the early 1900s, at about the

---

unsupported by neuroscience if the neuroscience is accompanied by a neuroimage, that is, if they see the brain "lighting up like a Christmas tree." The literature is actually mixed about whether the Christmas tree effect is real, and how strong it is. Compare D. McCabe & A. Castel, *Seeing is believing: The effect of brain images on judgment of scientific reasoning*, 107 COGNITION 343–52 (2008) (detecting an effect), with N. Schweitzer & M. Saks, *Neuroimage evidence and the insanity defense*, 29 BEHAV. SCI. & L. 592–607 (2011) (detecting no effect).

same time as the famous *Frye* case,[36] which required a general scientific consensus before novel scientific evidence could be admitted. *Frye* has since been displaced by generally more forgiving evidentiary standards,[37] and yet courts across the country continue their almost uniform refusal to admit polygraph evidence.

This despite the fact that modern polygraphy has gotten substantially better since *Frye* was decided in 1920. A 2003 National Academy of Sciences report found that the best polygraph studies had accuracy rates considerably better than chance, ranging from 70 to 80 percent.[38] The continuing problem with polygraphy is not so much its accuracy levels as its reliability—that is, the extent to which its accuracies vary across polygraph examiners, across subjects, and even within the same examiner and same subject but over time. It was this problem of reliability, not accuracy, that caused the NAS to conclude in its 2003 report that polygraphy was still not ready for courtroom prime time.

Lie detection by neuroscience—primarily by fMRI—is already on its way to improving the accuracy of traditional polygraphs. A 2015 meta-study done by some of my MacArthur colleagues found that fMRI lie detection techniques have somewhat increased accuracies compared to polygraph, but that reliability problems remain.[39] That study also expressed concerns about countermeasures, to which fMRI lie detection seems particularly vulnerable.

I don't doubt for a moment that fMRI lie detection will get more accurate and more reliable in the next several years—so accurate and so reliable that courts and legislatures may have to face up to the fact that such methods are more accurate and reliable than juries guessing about who they believe. As with the diagnosis of pain and some mental conditions, this improvement seems an ineluctable result of MVPA classifiers getting better

---

[36] Frye v. United States, 293 F. 1013 (D.C. Cir. 1923). *Frye* actually involved polygraphy by blood pressure.

[37] The *Frye* approach was rejected by the United States Supreme Court in the famous trilogy of scientific evidence cases: Kuhmo Tire Co. v. Carmichael, 526 U.S. 937 (1999); General Elec. Co. v. Joiner, 522 U.S. 136 (1997); and Daubert v. Merrill Dow Pharma., Inc., 509 U.S. 579 (1993).

[38] National Research Council, The Polygraph and Lie Detection (National Academies Press 2003), *available at* https://www.nap.edu/read/10420/chapter/1

[39] M.J. Farah et al., *Functional MRI-based lie detection: Scientific and societal challenges*, 15 Nat. Neurosci. 123–31 (2014).

and better as they examine more and more subjects. But I've put lie detection by neuroimaging in the "long term" category instead of the "near term" category for two reasons.

First, it is not at all clear to me that fMRI lie detection will soon be able to solve the problem of countermeasures. The electromagnetic measurements made by fMRI are very sensitive to subjects' movements, not just head movements that will nullify the results entirely—and which themselves might be a rudimentary kind of lie detection—but also by tiny, hard-to-detect body movements, like the wiggling of a finger. Neuroimaging-based lie detectors might likewise be easily fooled by a whole host of mental versions of wiggling one's fingers—thinking about something else, for example, while being scanned. I don't see any reason why, in theory, enough pattern classifiers unleashed on enough subjects trying to fool them will not eventually win, or at least be able to stay ahead of, this countermeasures arms race. But this will take time.

The second, and I think more profound, reason lie detection by neuroimaging will probably not impact the law for at least another 10 years is the problem of the instructed lie. All lie detection research is hobbled by the fact that, in one way or another, experimenters tell subjects to lie, or at the very least give them permission to lie. In the typical experiment, researchers do something in front of the subjects, such as hiding an object somewhere in the room. They then tell half the subjects to lie and half to tell the truth when later asked about where the experimenter hid the object. In the best experiments, subjects are free to decide on their own whether to tell the truth or lie. But in either case, they are given permission to lie.

These kinds of instructed or invited lies presumably have very different psychologies behind them, and therefore very different neural patterns, than lies told about whether someone robbed a liquor store or committed a rape. Until a few years ago, it seemed like the problem of the invited lie was insoluble, because there didn't seem to be any reliable way to incentivize subjects to voluntarily lie in the laboratory that was anywhere near the ecological equivalent of the lying criminal suspect.

But in 2009, Joshua Greene, now the director of Harvard's Moral Cognition Lab, came up with an experimental paradigm that may solve the problem of the invited lie. He told subjects he was investigating extrasensory perception, and asked the control group to predict whether a computerized coin would land on heads or tails. Not unexpectedly, they all averaged around 50 percent. But he asked a second group to record their predictions *after* the coin toss, thus incentivizing them to lie (both groups got paid for each successful guess). It turns out that about 1/3 of the subjects never lied, 1/3 occasionally lied (an average of 20% of the time), and 1/3 almost always lied (an average of 80%). These tasks were done in the scanner, and Greene and his colleagues could detect with 100 percent accuracy, by looking at subjects' brain images alone, which people never lied and which people were regularly lying.[40]

This work leads me to believe that the problem of the invited lie will eventually be overcome by lie detection researchers, and that we will someday have accurate and reliable lie detection techniques suitable for forensic contexts. I put this prediction in the 10–50 year category because, despite Greene's work, it does not appear much if any lie detection research is being done using methods like Greene's. Perhaps that's because, although he solves the invited lie problem, Greene's method cannot tell us whether subjects were lying on any given coin flip—only whether they were, in the aggregate, liars or truth-tellers.

It is also not clear to me that the law will ever get over, or should get over, its reluctance to embrace lie detection.[41] If it does, lie detection will have a major impact in virtually all areas of the law. Imagine criminal defendants or personal injury plaintiffs entering settlement discussions armed with favorable or unfavorable lie detection results. Imagine civil litigants fighting over whether there was an oral modification of a contract being able to test that allegation reliably, pre-litigation. Plea bargains and settlements would probably soar, and the already disappearing jury would get even more endangered. Indeed, this fear of taking away the truth-telling function of juries is no small part of the law's resistance to lie detection generally.

4. Mental States and Drugs

Although it is difficult to induce legally relevant mental states in experimental subjects, baby steps in this direction are starting. A group of us has recently published the first neuroimaging study ever done that attempts to

---

[40] J. Greene & J. Paxton, *Patterns of neural activity associated with honest and dishonest moral decisions*, 106(30) Proc. Nat'l Acad. Sci. 12506–11 (2009).

[41] For a survey of some of the legal and policy issues that reliable lie detection technology might raise, see Farah et al., *supra* note 39; Owen D. Jones & Morris B. Hoffman, *Lies, Brains and Courtrooms*, 85 U.S.L.W. 94 (2017).

examine the mental states of knowing a fact versus being reckless about it.[42] We induced these two different states by using a double-risk gambling paradigm. Subjects played a game where they were asked whether they would risk taking a suitcase that may or may not contain contraband over a border. If they were caught with contraband, they lost. If they successfully carried contraband, they won. We varied the chances that the suitcase contained contraband from 100 percent (knowing state) down to 50 percent, 33 percent, and 25 percent (these latter chances representing different levels of recklessness). We did that by presenting subjects with varying numbers of suitcases, at least one of which contained contraband, and having them randomly pick which suitcase to carry. Thus, subjects presented with just one suitcase were 100 percent sure it contained contraband. At the other extreme, subjects presented with five suitcases had only a 20 percent chance that the suitcase they picked contained contraband. We also varied their chances of being caught, and informed them of that variation. The subjects played the game in the scanner, having fMRI images recorded as they made their decisions about whether to try to cross the border or not.

We used a new-fangled MVPA technique called "elastic net regression" to crunch the data. We were able, looking at whole brain images alone, to distinguish when a subject knew there was contraband from when they were merely risking that fact, at an accuracy rate of about 80 percent. Our methods could even distinguish between the different lower-levels of risk.

Of course, this one experiment is just the beginning. Much work needs to be done, both from the whole brain level down and from traditional subtraction and connectivity methods up, to begin to develop a neural model of what it means to be reckless. But the payoff could be huge, primarily in the area of intoxication.

In virtually every jurisdiction, voluntary intoxication is not a defense to knowing or reckless crimes but is a defense to purposeful ones. There are perfectly legitimate consequentialist justifications for this rule, and even some morally grounded ones. If you voluntarily choose to degrade your own risk-assessment abilities in exchange for temporarily feeling better, then maybe you should be held accountable for that risk degradation.

---

[42] I. Villares et al., *Predicting the knowledge-recklessness boundary in the human brain*, 114(12) Proc. Nat'l Acad. Sci. 3222–27 (2017).

We will only consider voluntary intoxication when the prosecution has to prove the most culpable of all mental states—purposeful.

But the law may have this exactly wrong as a neurological matter. It may be that some drugs substantially interfere with our abilities to assess risks—including of course our own ability to judge whether those abilities are impaired (this is the drunk driver problem)—but have no impact on our ability to form intentions. As we learn more about how brains assess risks, and how various substances interfere with that task, I think it is quite possible that in the next 50 years this general rule about intoxication may be reversed, at least for some substances that are found to seriously degrade those risk-assessment systems but not to materially interfere with intent formation.

5. Simplifying Some Mental States

Do people in fact judge criminal responsibility the way the Model Penal Code (MPC) assumes we do: holding purposeful wrongdoers more responsible than knowing ones, knowing ones more responsible than reckless ones, and reckless ones more responsible than negligent ones? And are we even capable of distinguishing these four different mental states? A group of us began behavioral and neuroimaging experiments several years ago to test these questions. Our results were both good news and bad news for proponents of the MPC. With one exception, subjects were quite good at distinguishing the four mental states and punishing them in the order predicted by the MPC. The one exception: the boundary between knowing and reckless. Subjects could not distinguish these two mental states, and even when given the definitions of them, they simply refused to treat them differently in terms of punishment.[43]

These experiments raise as many questions as they answer. But if their results hold, we may see in the next 50 years a move toward collapsing the

---

[43] Francis X. Shen et al., *Sorting Guilty Minds*, 86 N.Y.U. L. REV. 1306 (2011); Matthew R. Ginther et al., *The Language of Mens Rea*, 67 VAND. L. REV. 1327 (2014). We performed a version of this experiment in the scanner, and the results suggested that different brain systems seem to be involved in assessing harm and in assessing the actor's mental state, and that yet a third system integrates the two into punishment. M. Ginther et al., *Parsing the behavioral and brain mechanisms of third-party punishment*, 36 J. NEUROSCI. 9420–36 (2016). For a survey of the neuroscience of third-party punishment, see F. Krueger & M. Hoffman, *The emerging neuroscience of third-party punishment*, 39(8) TRENDS NEUROSCI. 499–501 (2016).

knowing and reckless mental states into a single mental state, or the elimination of the knowing mental state entirely, at least for results elements. I confess this may be more an occupational hope than a serious prediction. I shudder every time I preside over a criminal trial in which I instruct the jury both on a knowing crime and a lesser-included reckless one (in my jurisdiction, for example, second degree murder and reckless manslaughter), worrying that my jurors may have no idea what the difference is between the two.

6. Detecting Autobiographical Memories

Researchers have discovered that neuroimaging pattern classifiers are pretty good at what may seem to be an impossible, futurist, task that you might think should be listed in my "never happening" category: detecting whether a subject has in fact recognized a face or a place. This work began with facial recognition studies, with the classifiers eventually being able to accurately detect, at levels approaching 90 percent, whether a given brain had or had not recognized a face previously shown in the experiment.[44] This technique could assist in detecting witnesses who are lying about identifying a suspect, and might even be able to quantify the confidence level of a positive identification. Since erroneous eye witness identification seems to be a big factor in wrongful convictions, this technology could be a boon to the wrongfully accused, and of course a major problem for the guilty.

Researchers expanded this work from the identification of faces to the identification of places,[45] and although their results to date are not quite as robust as the work with faces, it is entirely possible that this technique will eventually get accurate enough for forensic use. Imagine being able to show a witness (or defendant) the picture of a crime scene and being able to confirm or eliminate, by brain imaging alone, whether that witness had been there!

My guess is that this kind of technology, which is really a particular kind of lie detection, will suffer the same resistance by the law as lie detection generally. But, as with lie detection by neuroimaging, detecting autobiographical memories could have big legal impacts in the shadows of the trial.

---

[44] M. Uncapher et al., *Goal-directed modulation of neural memory patterns: Implications for fMRI-based memory detection*, 35 J. Neurosci. 8531–45 (2015).

[45] J. Rissman et al., *Decoding fMRI signatures of real-world autobiographical memory retrieval*, 28(4) J. Cogn. Neurosci. 604–20 (2016).

7. Individualizing Measures of Brain Maturity

The law is forced in many different contexts to draw lines based on the chronological age of legal actors. How old should people have to be, to be held to their promises in contract? How old should they be before they can decide to marry, or have an abortion, or vote? How old do they have to be before they are criminally responsible, or can be executed?

Each of these legal questions assumes that age is a proxy for the real question, which is: How mature must a young individual be before we apply a law intended for a fully mature adult? That is, how mature is that individual's brain?

We all know that we are not born with fully mature brains, and that brain development is highly individualized. There are 30-year-olds who act like teenagers and teenagers who act like adults. There are many psychological and neuropsychological tests that can measure different aspects of brain maturity, but at the moment there is no neuroscience that can reliably do so on an individual basis. One reason the law is unwilling to individualize some of its maturity-based applications may be that it is not convinced that these behavioral measures truly reflect differences in brain development. Is that immature 30-year-old immature because she has an underdeveloped brain or because her parents spoiled her (and are these things really different)? In the next 50 years neuroscience may be able to accurately and reliably measure brain maturity on an individual basis. If that happens, I believe the law may abandon some of its chronological age proxies.

From conception through adolescence, the brain develops generally from inside out, from back to front. The prefrontal cortex (PFC) is critical for higher-order tasks, like decision making and inhibition. Because it is located at the very front and outside of the brain, it is among the last brain regions to develop. The average human male PFC is not fully developed until the mid-twenties, the average female a bit earlier. But these are just averages. The wide variations in maturity that we see among individuals come at least in part from the fact that individual brains vary widely in their developmental paths.

Young brains develop in two different ways. As you might expect, they grow neurons and add synapses. But they also pare away synapses in response to the environment. Many of the neurological challenges faced by the adolescent brain come from the problem that these brains are inexperienced in the world, and therefore have not had sufficient paring away of their neural connections. That's one reason why adolescents actually use much more

PFC energy than adults when asked things like whether it's a good idea to eat broken glass or swim with sharks. Adults know from experience—direct or indirect—that these things are too risky, and it seems these learned experiences are in part a matter of paring away connections that are no longer needed to contemplate such risky possibilities. Immature brains have to reason their way through these questions, and uninformed reason alone is often simply not enough to make the right risk-taking decisions.[46]

Advances in neuroscience, especially in connectivity, may eventually allow neuroscientists to measure brain maturity on an individual basis, accurately and reliably enough to dispense with chronological age as a proxy. In addition to the gross measures of brain maturity currently used—including the percentage of gray matter (a proxy for the number of synapses)—I can imagine technologies that will someday also be able to measure the amount and kinds of paring during development. Indeed, using these kinds of approaches, with MVPA, neuroscientists can already distinguish children's brains from adults' brains by resting state neuroimages alone, with 90 percent accuracy.[47] Putting various kinds of measures together might give scientists a way of pinpointing where on the general developmental trajectory any individual's brain might lie. I can picture these tests accumulating into a single number representing "brain age."

This single measure might not only take into consideration where an individual brain lies in terms of endogenous development,[48] but also detect when brains have been so deprived of outside stimulation that their normal development has been stunted by way of insufficient paring.[49]

---

[46] Not unlike psychopaths, who, because they generally lack moral intuition, must reason their way to moral decisions, using much more energy in the PFC than non-psychopaths. K. Kiehl, *A cognitive neuroscience perspective on psychopathy: Evidence for paralimbic system dysfunction*, 142 PSYCHIA. RES. 107–28 (2006).

[47] N. Dosenbach et al, *Prediction of individual brain maturity using fMRI*, 329(5997) SCIENCE 1358–61 (2010).

[48] Of course, there really is no "endogenous" brain development; brains are experience machines, and their "natural" development trajectory is inexorably bound up with the experiences they have. Despite this highly individualized nature of brain "maturity," there are still common developmental signposts that I imagine researchers will someday be able to measure to create a rough but very detailed average brain development trajectory, against which individual brains could be compared.

[49] I have no illusions that such measures could distinguish between all good and bad experiences, but they very well might be able to detect brain environments that are so bad they have had significant impact on normal brain development.

Of course, this won't end all age-based legal debates. Even a normalized neuroscience-grounded measure of brain maturity will not answer the policy questions of whether a 16-year-old brain should vote, should be able to consent to sex or to an abortion, or should be executed. Maturity is not a one-size-fits-all policy inquiry, but at least the law might be making these judgments based on real brain age and not on its unreliable chronological proxy.

VI. NEVER HAPPENING

There are two related areas of law where I doubt neuroscience will ever have any impact: addiction and moral agency. These are the very two areas where hopeful neuroscientists sometimes predict the most significant and earth-shaking impacts.

8. Addiction

Neuroscience will likely never crack the addiction problem because it will never crack the mind/body problem—the problem of how even non-addicted brains make decisions.

The story of the neuroscience of addiction has been profoundly dissonant, and extraordinarily disappointing. On the one hand, enormous amounts have been learned about what various drugs do to various neurotransmitters and their transmission systems. The pace of these discoveries was so great that in the mid-1990s addiction specialists were predicting pharmacological cures for all kinds of addictions.[50] More than twenty years later, none has come to pass. And that's mainly because addiction is a mystery embedded in the brain's still ill-understood risk/reward system. Compelling neurobiological addiction models have been proposed, but they will all remain theoretical until neuroscience cracks the puzzle of how non-addicted brains assess risk and reward, and then make decisions. It's one thing to know that cocaine is a dopamine and norepinephrine

---

[50] The first director of the National Institute for Drug Abuse (NIDA), Alan Leshner, predicted in 1997 that there would be a pharmacological cure for methamphetamine addiction in five years. S. Satel & S. Lilienfeld, *Addiction and the brain disease fallacy*, 4 FRONT. PSYCHIA. 141 (2014). As late as 2012, one of his successors, Nora Volkow, predicted a pharmacological cure for cocaine addiction. *Id.*

re-uptake inhibitor,[51] and quite another to understand what that has to do with the risk-reward system's twin breakdowns when it comes to addiction: dependence and tolerance.

My guess is that the neurobiology of the tolerance phenomenon—addicts generally need higher and higher doses to achieve the same effects—will someday be fully understood, perhaps even in my lifetime. But dependence is another thing entirely, because it touches on the black hole of free will. Indeed, many critics of the disease model of addiction believe that model's central failing is that it ignores the role of will, even in addicts.[52] Whatever one's views of the contention that addiction is a "chronic and recurring brain disease,"[53] it seems unlikely we will ever get at the problem of dependence if we cannot even agree about its essential nature.

The best evidence that the mystery of dependence will never be solved in the foreseeable future might be to look at how much money has been thrown at the neuroscience of addiction and to what effect. Billions have been spent in the last several decades,[54] none of which has materially advanced our understanding of dependence.

9. The Abolition of Responsibility

Robert Sapolsky and others think it may only be a matter of time before the mechanistic nature of the brain is fully understood, and that when that happens the very notion of human responsibility will evaporate, precisely for the same reason we do not hold cars or other machines responsible. Like many others, I think that this kind of extreme materialism, even if it comes to be accepted, will not destroy responsibility. But I also don't think it will come to be accepted.

---

[51] Meaning it inhibits substances in the synapse that soak up excess dopamine and norepinephrine.

[52] *See, e.g.*, Satel & Lilienfeld, *supra* note 50.

[53] This phrase was coined by Alan Leshner as a strategy to talk Congress into appropriating more money for addiction treatment and research. *Id.*; S. Vreko, *Birth of a brain disease: Science, the state and addiction neuropolitics*, 23(4) HIST. HUM. SCI. 52–67 (2010).

[54] Addiction is perennially one of the most generously funded of mental health research categories. The NIH alone budgeted $1.343 billion for addiction research in 2017 (consisting of $375 million for alcohol, $102 million for cannabinoids, and $865 for drug abuse generally). NIH, Estimates of Funding (July 3, 2017), available at https://report.nih.gov/categorical_spending.aspx#legend1.

The brain is not only the most complicated machine we have ever encountered, it may be the kind of machine we can never fully understand. With an average of 80 billion neurons, or the equivalent of 1 gigabyte, the human brain is certainly not a massive storage device. A single-layer Blue-Ray disc can hold 25 gigabytes. It's not the size of the brain network that makes it so complicated; it's the way the network's nodes—the synapses between the neurons—communicate with one another. Whether a signal from one neuron crosses the synapse and makes it to a downstream neuron depends on the makeup of the chemical soup sitting in the synapse. The nature of that soup is not only changed by the neuron firing, and releasing neurotransmitters and neuromodulators itself, it is changed by the firing patterns of other nearby neurons. Quite apart from the soup in the synapse, a neuron's physical characteristics can be altered by its own firing and by the firing of nearby neurons. For all these reasons there is a profoundly self-referential aspect to the brain.

Given its self-referential nature, I have lots of doubts about whether we can ever fully understand the brain, for reasons analogous to the indeterminacy of mathematical systems. Kurt Gödel famously proved that every sufficiently rich logical system is indeterminate—that is, that there are true propositions in it that cannot be proved true using the axioms of the system.[55] I suspect that when we say we "understand" some deep aspect of the brain, we may be stating true things about that aspect that we will never be able to prove.

If this sounds a little too mathematical and esoteric, a similar kind of objection to hard determinism is that the brain is simply too complicated ever to be a predictable machine. Even if we knew every fact about every synapse, and every influence every other neuron had on that synapse and on every other neuron, we simply could not predict how a brain would react to given inputs because the complications themselves create a kind of emergent indeterminism. Many well-understood, and quite simple, physical systems exhibit this same phenomenon. Lost in all the achievements of the space program is the mathematical fact there is no algebraic formula that will accurately predict the position of three or more objects in space,

---

[55] K. Gödel, *On formally undecidable propositions of principia mathematica and I, in* KURT GÖDEL: COLLECTED WORKS (Oxford 1986). For a lively and accessible history of Gödel and his famous theorem, see REBECCA GOLDSTEIN, THE PROOF AND PARADOX OF KURT GÖDEL (Norton 2006).

each exerting some gravitational force on the other.[56] Is it any surprise that the brain's 80 billion objects, each influencing the firing potentials of one another, will also be functionally impossible to predict?

There is even some recent neuroscience that seems to suggest that our difficulty in predicting what brains will decide is not just a problem of indeterminism or complexity, but also a problem that brains are inherently random. NYU neuroscientist Paul Glimcher has spent decades examining the brains of macaques as they decide whether to look left or right, a decision Glimcher incentivizes by giving the monkeys different rewards depending on which way they look. He discovered neural patterns that suggest that what the monkeys are doing is changing probabilities depending on the changed payoffs, but that a random process was being applied to those probabilities.[57]

Under Glimcher's model, when you pull your hand away from a hot flame, you are not really "deciding" to pull it away; your brain is rolling dice that are heavily loaded toward pulling your hand away. They are so heavily loaded that we call the resultant behavior a "reflex." By contrast, in Glimcher's world, when you decide to walk to work instead of driving, a different set of dice is being thrown that are not so heavily weighted toward one choice or another. It feels like we are "deciding," but really our random decision generator is just operating on more evenly matched probabilities.

Of course, the idea of free will is just as dead in Glimcher's world as it is in a deterministic one. We are not "deciding" anything, we're just rolling dice. But precisely because we are just rolling dice, we will never be able to predict an individual brain's decision in any context other than when the probabilities are so skewed as to be akin to reflexes.

---

[56] This is so-called "N-body problem," and in 1890, the French mathematician Henri Poincaré proved that for three bodies or more there is no algebraic solution that will describe the bodies' positions at future times. Of course, with just one object, its future position for all of eternity is completely described by the formula distance = velocity × time. With two objects, there are differential equations that will completely describe their future positions. But when just one more body is added, the problem becomes insoluble, at least algebraically. That doesn't mean their positions cannot be approximated. Indeed, in 1991, a Chinese mathematics student, Quidong Wang, proved that for Ns of three and above, there are power series that converge on solutions, though their convergence is so slow that even the fastest computers don't use them to approximate celestial positions. *See, generally*, F. Diacu, *The solution of the N-body problem*, 18(3) MATH. INTELLIGENCER 66–70 (1996).

[57] *See generally*, PAUL W. GLIMCHER, DECISIONS, UNCERTAINTY AND THE BRAIN (MIT Press 2004).

Which circles us back to another reason why I don't think anything we learn about the brain will ever eviscerate our powerful senses of free will and responsibility: these are powerful beliefs because we evolved them. Hard determinists talk a good talk, but, like all of us, they behave as if they had free will. And that's because notions of free will, responsibility, punishment, and forgiveness are all part of our evolved neural architectures.

I remember being asked once by one of our MacArthur philosophers why I was a "moral realist." I told him I was a "moral realist" in the same way I was a "spinal realist" or an "opposable thumb realist." These are all things we evolved to help us navigate our emergent environment, and things that continue to be useful in our current one. The law widely recognizes notions of responsibility, blame, forgiveness, and punishment because they are deeply embedded in all of us. None of us will give them up just because some neuroscientist tells us we have no free will, any more than any of us would agree not to use our thumbs just because some evolutionary theorist tells us they are really no longer necessary.

CONCLUSION

In the next 5 to 10 years, neuroscience will provide law with accurate and reliable ways of detecting chronic pain, and perhaps distinguishing real chronic pain from exaggerated chronic pain, and both from malingering. This will have significant impacts in tort and disability law. In that same period, neuroscience will develop methods to diagnose several legally relevant mental conditions, though neurodiagnosis will likely have only marginal impact on the law, which will continue to use traditional clinical methods of diagnosing these conditions and presenting them to fact finders.

In the next 10 to 50 years, neuroscience will develop accurate and reliable lie detection methods. These developments may include methods to detect autobiographical memories of faces and places. Although the law will probably continue to be skeptical of both of these forms of lie detection as a matter of trial evidence, neuro lie detection will have big impacts in all crevices of the law that lie in the shadow of trial, including settlement and plea-bargaining.

Toward the end of this 50-year period, neuroscientific insights into the effects of certain drugs on the risk-reward network may cause the law to reverse its general rule that voluntary intoxication is a defense to purposeful crimes but not to knowing or reckless ones. In this same vein, neuroscience may convince the law that there is no difference between knowing and reckless states of mind when it comes to results elements.

Also at the far end of the next 50 years, neuroscience may be able to provide individual measures of brain maturity, allowing the law in some domains to abandon the line-drawing it now uses based on chronological age.

But neuroscience will likely never solve the mystery of addiction—at least the dependence part of that mystery. Nor will it ever convince the law to abandon notions of individual free will and responsibility.