3-8-2018

Lying, Deception, and fMRI: A Critical Update

Michael S. Pardo
University of Alabama - School of Law, mpardo@law.ua.edu

Lying, Deception, and fMRI: A Critical Update

Michael S. Pardo

NEUROLAW AND RESPONSIBILITY FOR ACTION
(Bebhinn Donnelly-Lazarov ed., Cambridge University Press,
forthcoming 2018)

Lying, Deception, and fMRI: A Critical Update

Michael S. Pardo*

Among the topics at the intersection of law and neuroscience—or 'neurolaw'—the use of fMRI as a lie detector has been one of the most prominent for scientists, academic commentators, and popular media. In one sense, this is understandable. Determining whether someone is speaking sincerely or is lying is a tremendously important legal topic, and it is one that is notoriously difficult to ascertain in the context of litigation. The ability to bypass the usual complex inferential process of assessing testimony with 'direct' evidence of whether someone is lying is thus understandably attractive.[1] Moreover, fMRI lie detection does not appear to raise the same sorts of thorny philosophical and empirical issues that affect other proposed uses of neuroscience in law—such as those involving free will,[2] ascriptions of criminal responsibility,[3] or even the related issue of mind-reading.[4] Rather, on the surface at least, fMRI

* Henry Upson Sims Professor of Law, University of Alabama School of Law. My thanks to Bebhinn Donnelly-Lazarov and Dennis Patterson for helpful comments and to Dean Mark Brandon and the Alabama Law School Foundation for generous research support.

[1] The idea that fMRI offers 'direct' access to lies in the brain (an idea advanced by some scholars and proponents of the technique) is based on a mistaken conceptual assumption about the evidence. As with more traditional lie-detection techniques such as polygraph, fMRI lie detection also relies on supposed correlations between lying and physiological changes in the body. Physiological changes in the brain no more constitute "lies" than do an increased heart rate or sweating. For a discussion of this conceptual issue in the context of fMRI lie detection, see Michael S. Pardo and Dennis Patterson, *Minds, Brains, and Law: The Conceptual Foundations of Law and Neuroscience* (New York: Oxford University Press, 2013), pp. 105-06. Moreover, even when a witness's sincerity is certain, this does not eliminate the inferential work necessary to evaluate the witness's testimony. Sincerity is not the same as veracity—it is also necessary to evaluate whether the witness is sincere but mistaken.

[2] Adam J. Kolber, 'Free Will as a Matter of Law', in Dennis Patterson and Michael S. Pardo (eds.), *Philosophical Foundations of Law and Neuroscience* (Oxford: Oxford University Press, 2016), p. 9.

[3] Stephen J. Morse, 'The Inevitable Mind in the Age of Neuroscience', in Dennis Patterson and Michael S. Pardo (eds.), *Philosophical Foundations of Law and Neuroscience* (Oxford: Oxford University Press, 2016), p. 29.

lie detection appears to involve a straightforward application of probative evidence being used to prove a relevant and important factual issue in law.[5] Thus, the understandable attention devoted to the topic.[6]

In *Minds, Brains, and Law*, Dennis Patterson and I examined the scholarly literature and judicial opinions on the topic.[7] We discussed several empirical and conceptual issues that affect the use of such evidence in law. Although the empirical limitations with the evidence have been explored in detail, the conceptual issues with the evidence have received considerably less attention. Consistent with the general theme our book[8], we focused in detail on the conceptual issues and the challenges that they raise for law.

In this chapter, I will focus on two conceptual issues and examine several studies that have been published since the publication of our book. The conceptual issues concern: (1) the distinction between deception and lying, and (2) the concept of lying itself (or the criteria for

---

[4] Gideon Yaffe, 'Mind-Reading by Brain-Reading and Criminal Responsibility', in Dennis Patterson and Michael S. Pardo (eds.), *Philosophical Foundations of Law and Neuroscience* (Oxford: Oxford University Press, 2016), p. 137.

[5] Despite this appearance, several difficult conceptual issues underlie the application of this evidence in law. Moreover, the fact that the conceptual issues are typically presupposed (and harder to detect) leads to some of the confusion about them. For discussion, see Pardo and Patterson, supra note 1.

[6] In another sense, however, the focus on lie detection is somewhat strange given the historical reluctance courts have shown toward lie-detection evidence in general. This reluctance, moreover, often concerns issues beyond perceptions about the reliability of such evidence. For excellent discussions, see D. Michael Risinger, 'Navigating Expert Reliability: Are Criminal Standards of Certainty Being Left on the Dock?' (2000) *Albany Law Review* 99 at 129-30; Frederick Schauer, 'Lie-Detection, Neuroscience, and the Law of Evidence', in Dennis Patterson and Michael S. Pardo (eds.), *Philosophical Foundations of Law and Neuroscience* (Oxford: Oxford University Press, 2016), p. 85.

[7] Pardo and Patterson, supra note 1 at pp. 79-120.

[8] Id. at p. xxi. ('The conceptual issues focus on the scope and contours of the concepts being employ in claims involving law and neuroscience.')

what constitutes a lie). As with the array of studies that we examined in our book, the more-recent studies also face serious limitations because of these conceptual issues. Part I summarizes the empirical and conceptual issues concerning fMRI lie detection. Part II illustrates these issues by discussing a concrete example: the attempted use of such evidence in the case of *United States v. Semrau*.[9] Finally, Part III summarizes recent studies and explains the conceptual problems at the heart of the experiments.

I. Functional MRI Lie Detection: Empirical and Conceptual Issues

The basic idea behind fMRI lie detection is that identifiable areas or patterns in the brain will correlate more with lying than with truthful behavior. Based on the strength of these correlations, so it is claimed, reliable inferences may be drawn about whether a person is lying or telling the truth on a particular occasion (or with regard to a particular subject matter).

The research supporting such claims depends on experiments such as the following.[10] Subjects are placed in an fMRI scanner, and they are presented (for example, on a computer screen or through headphones) with a series of binary questions (for example, yes-no, true-false, heads-tails, and so on). Subjects are instructed to answer by pressing one of two buttons, each corresponding to one of the possible answers. Subjects are also instructed either to 'lie' or to answer truthfully about an array of different topics such as biographical information, daily

[9] United States v. Semrau, 693 F.3d 510 (6th Cir. 2012); United States v. Semrau, 2010 WL 6845092 (W.D. Tenn., June 1, 2010).

[10] For examples of the several studies employing this paradigm, see Pardo and Patterson, supra note 1 at pp. 82-85.

events, numbers, playing cards, past events, or under which object money was hidden. While subjects answer, the scanner measures brain activity indirectly by measuring blood flow to areas of the brain.[11] Scanning data is then processed and translated into a brain 'image,' which is a visual representation of the underlying statistical measurements of blood flow.[12]

Based on these studies, researchers reported several positive findings on fMRI lie detection. Initial studies examined group results to explore whether activity in any brain areas in general correlated with lying.[13] These studies identified several brain areas, including the prefrontal and the anterior cingulate cortices.[14] Subsequent studies then focused on individual-level determinations, sometimes reporting impressive results. For example, one study described a 90 percent success rate in identifying which object was stolen in a mock-crime experiment.[15] A second study reported a 100 percent success rate in identifying the participants in a mock

---

[11] The technique depends on a number of important scientific principles—most importantly the BOLD ('blood oxygen dependent level') signal to infer conclusions about brain activity. The relationship between the signal and brain activity depends on two ideas: when hemoglobin in blood delivers oxygen to the brain it becomes 'paramagnetic' and will disrupt a magnetic field (such as one created by an MRI scanner), and when activity increases in a particular brain region blood flow will increase in order to supply more oxygen. For an excellent overview, see Adina L. Roskies, 'Brain Imaging Techniques', in Stephen J. Morse and Adina L. Roskies (eds.), *A Primer on Criminal Law and Neuroscience* (New York: Oxford University Press, 2013), pp. 37-71.

[12] Id.; Teneille Brown and Emily Murphy, 'Through a Scanner Darkly: Functional Neuroimaging as Evidence of a Criminal Defendant's Past Mental States' (2012) 62 *Stanford Law Review* 1119.

[13] Pardo and Patterson, supra note 1 at pp. 85-86.

[14] Id. at 85.

[15] F. Andrew Kozel et al., 'Detecting Deception Using Functional Magnetic Resonance Imagining' (2005) 58 *Biological Psychiatry* 605.

crime (9 out of 9), along with a 33 percent rate in identifying 'no crime' participants (5 out of 15).[16]

Several empirical issues, however, limit the utility of such research for providing a foundation for evidence in legal settings. These issues include: (1) the small sample sizes of the subjects tested; (2) the lack of diversity among those tested; (3) the time differences between experimental and litigation settings; (4) differences among types of lies; (5) the low-stakes conditions for experimental subjects; (6) the fact that the identified brain regions are also correlated with behaviors and activities other than lying; (7) the effectiveness of countermeasures; and (8) the procedures used for constructing brain 'images.'[17] The fundamental empirical question in a litigation setting is how likely a particular answer given by an expert reporting the results from fMRI technology is likely to be accurate when applied to a particular person, answer, or subject matter. Each of the issues listed above affect the answer to this question and thus the inferences that may be drawn from such evidence.

In addition to the empirical issues, a number of *conceptual* issues also affect the quality of fMRI evidence in this context. Although the conceptual issues are less noticeable and have received considerably less attention in the literature, they are both significant and problematic. The array of empirical issues noted above takes place within a conceptual framework that

---

[16] F. Andrew Kozel et al., 'Functional MRI Detection of Deception After Committing a Mock Sabotage Crime' (2009) 54 *Journal of Forensic Science* 220.

[17] See Pardo and Patterson, supra note 1 at pp. 97-99 (discussing these issues). Three recent articles provide excellent overviews of the empirical limitations regarding such evidence. Martha J. Farah et al., 'Functional MRI-Based Lie Detection: Scientific and Societal Challenges' (2014) 15 *Nature Reviews Neuroscience* 123; Joshua W. Buckholtz and David L. Faigman, 'Promises, Promises for Neuroscience and Law' (2014) 24 *Current Biology* R861; Henry T. (Hank) Greely, 'Mind Reading, Neuroscience, and the Law', in Stephen J. Morse and Adina L. Roskies (eds.), *A Primer on Criminal Law and Neuroscience* (New York: Oxford University Press, 2013), p. 120.

includes a number of presuppositions about what exactly is being measured in the experiments.[18] These issues are conceptual rather than empirical because they do not concern the empirical relationship between brain activity and behaviors; rather, they concern the concepts that are used to identity and pick out the target behaviors in the first place. For this reason, the conceptual issues are in an important sense prior to, or foundational for, the empirical issues. Confusing or erroneous conceptual presuppositions will render empirical claims built on top of them problematic. Experiments looking to confirm the brain activity that correlates with any variable *X*, in other words, need correct conceptual assumptions about what exactly constitutes *X*. In the lie-detection context, these conceptual issues include, for example, what constitutes lying and related issues.

In our book, we discussed several conceptual issues underlying fMRI lie detection. In this chapter, I will focus on two of them: the distinction between lying and deception, and the definition of lying employed or presupposed in the studies.[19]

The first conceptual issue concerns the distinction between lying and deception. In order to discover the brain correlates of lying and truthful behavior, the studies focus in particular on deception or an 'intent to deceive' on the part of subjects.[20] In some cases, 'lies' are defined in terms of 'an intent to deceive.' In one sense, this focus on deception makes sense because there is likely to be substantial overlap between lying and intending to deceive, and evidence that

[18] Gabriel Abend, 'What are Neural Correlates Neural Correlates of?' (2016) *BioSocieties* 1-24.

[19] Other conceptual issues with brain-based lie detection are discussed in Pardo and Patterson, supra note 1 at pp. 99-115.

[20] For examples, see Pardo and Patterson, supra note 1 at p. 108.

someone is acting deceptively makes it less likely that they are speaking sincerely. But the equation of lying with deception is also conceptually problematic because the two are not identical.[21] An intent to deceive is neither necessary nor sufficient to lie. It is possible to deceive without lying, and it is possible to lie without deceiving or even having an intent to deceive. An example of the latter includes a witness who has been threatened into giving false alibi testimony for a witness.[22] The witness may give such testimony not intending to deceive but hoping the audience (the police, a judge, or a jury) will see through the false statements and refuse to rely on them. The witness will have undoubtedly lied, and the testimony will also fit legal definitions for perjury and false-statement crimes.[23] Importantly, however, the witness will not have had an intent to deceive. Thus, fMRI technology that *could* reliably measure deceptive intentions would nevertheless be *under*-inclusive in measuring whether someone is lying or speaking sincerely.

---

[21] Id.; Don Fallis, 'Lying as A Violation of Grice's First Maxim of Quality' (2012) 66 *Dialectica* 563 ('several philosophers . . . have pointed out that there are lies that are not intended to deceive'); pp. Seana Valentine Shiffrin, *Speech Matters: On Lying, Morality, and the Law* (Princeton: Princeton University Press, 2014), pp. 19-26 (distinguishing lying and deception). Shiffrin also observes that '[t]he failure to distinguish carefully between lies and deception represents a defect in the empirical literature about lying.' Id. at p. 20 n.28.

[22] Lies without an intent to deceive include the category of so-called 'bald-faced lies' in which 'a speaker "goes on record" with something even though everybody knows it is false.' Adam J. Arico and Don Fallis, 'Lies, Damned Lies, and Statistics: An Empirical Investigation of the Concept of Lying' (2013) 26 *Philosophical Psychology* 790. See also Roy Sorensen, '*Bald-Faced Lies*! Lying Without the Intent to Deceive' (2007) 88 *Pacific Philosophical Quarterly* 251; Thomas L. Carson, 'The Definition of Lying' (2006) 40 *Nous* 284. An example (from Carson, supra) is a student whom everyone knows cheated and who brags openly about doing so. Suppose the Dean will punish the student only if he confesses (and the student knows this). Although everyone (including the Dean) knows the student cheated (and the student knows the Dean knows), the student denies cheating when questioned by the Dean. This is an example of lying, but the student is not trying to deceive anyone.

[23] For example, a witness commits perjury under federal law in the United States if, while testify under oath, he 'willfully' asserts something 'he does not believe to be true.' 18 U.S.C. § 1621(1). 'Willfully' in this context means making the statement with the intent to give false testimony. Unites States v. Dunnigan, 507 U.S. 87, 94 (1993). Similarly, a potential witness commits a federal crime in the United States when he 'knowingly' and 'willingly' makes a 'false statement' to a government agent on a matter within the jurisdiction of the United States. 18 U.S.C. § 1001(a). Either of these crimes may be committed without having an intent to deceive.

This under-inclusiveness is perhaps a relatively minor point, suggesting that the proposed technology may fail to pick up one category of lying behavior. Much more important, however, is the second conceptual issue—which reveals that the studies are significantly and problematically *over*-inclusive in their conceptions of lying. This conceptual issue concerns what constitutes a lie and whether the experiments are measuring lies.

Although prototypical cases of lying may by easy to identify, the concept of lying has been notoriously difficult to articulate in terms of necessary and sufficient conditions. A great deal of recent philosophical analysis, however, has helped to clarify its scope and contours.[24] Most importantly, lying involves more than simply stating something that is false or that a speaker believes to be false. There is also a social, normative component to lying that distinguishes it from other activities such as acting in a play, role-playing, playing a game, telling a joke, or using sarcasm, in which actors may say things they believe to be false. In our book, we relied in particular on recent work by philosopher Don Fallis, who has emphasized the 'normative component of assertion that is necessary for lying.'[25] As Fallis explains, the difference that makes saying 'I am the Prince of Denmark' a lie when told at a dinner party but not when on stage during a play are the distinct norms of conversation that are in effect in the former. In the former, but not the latter, there is an operative norm against communicating something believed to be false. Prototypical cases of lying involve communicating something

[24] See, for example, the sources cited in note 22.

[25] Don Fallis, 'What is Lying?' (2009) 106 *Journal of Philosophy* 29, 35.

you believe to be false *while also* violating (or believing yourself to be violating) a norm of communication against doing so.[26]

This normative component is critically important in legal settings. Lying in such settings not only violates an operative communicative norm—it also illegal.[27] This normative component, however, appears to be absent in the fMRI experiments. Subjects who give false answers in the scanner are not violating an operative norm against communicating something believed to be false; indeed, they were instructed to do just that.[28] Thus, the subjects do not appear to be lying. Therefore, the brain activity measured in such experiments does not appear to be brain activity accompanying lying.[29] Rather, it appears to resemble other types of non-lying activities noted above such as acting, joking, or playing a game. This is a conceptual problem that renders the empirical results from fMRI studies—and the inferences that may be drawn in legal settings—problematic. And it does so for reasons that are distinct from the litany of empirical issues noted above.

---

[26] Id. ('You believe that you make this statement in a context where the following norm of conversation is in effect: do not make statements that you believe to be false'). In a more-recent article, Fallis explain that this definition of lying may not cover some cases of lying by children or others who may be unaware of the norm against saying false things. See Fallis, supra note 21.

[27] See supra note 23.

[28] Note that this is distinct from whether the statements are 'instructed lies.' Someone can be instructed to lie and still lie—for example, an employee instructed by an employer to lie in a deposition. In doing so, the employee is still violating a norm against communicating false beliefs. When this norm is not present (such as in the experiments), then the instructed false statements are not examples of lying.

[29] Among the fMRI literature, we mentioned as a notable exception one study whose experimental design did appear to create the possibility for lying by participants. Joshua D. Greene and Joseph M. Paxton, 'Patterns of Neural Activity Associated with Honest and Dishonest Moral Decisions' (2009) 106 *Proceedings of the National Academy of Sciences* 12506; Pardo and Patterson, supra note 1 at pp. 110-11.

The next Part illustrates the practical, real-world problems generated by this conceptual issue.

II. *United States v. Semrau*: An Application

The well-known and much-discussed case of *United States v. Semrau* provides a concrete example for illustrating the significance of the conceptual issues at the heart of fMRI lie detection.[30] The case is noteworthy because it is the first detailed attempt to introduce the results of fMRI lie detection at trial in a federal criminal case in the United States.[31] The case is also noteworthy because of the detailed and careful analysis provided by the trial[32] and appellate[33] courts in assessing expert testimony on the evidence and the underlying scientific studies. The trial court excluded the evidence and the appellate court upheld the ruling, after each court discussed several empirical issues related to the evidence. Although the courts' reasons were sufficient to exclude the evidence, the conceptual issues (not discussed at either level) provide additional, foundational problems with such evidence. The conceptual issues are important to highlight because, in excluding the evidence, the courts in *Semrau* each acknowledged the possible admissibility of such evidence in the future. This Part first discusses the details of the

[30] United States v. Semrau, 693 F.3d 510 (6th Cir. 2012); United States v. Semrau, 2010 WL 6845092 (W.D. Tenn., June 1, 2010).

[31] In two other cases at the state-court level, parties have attempted unsuccessfully to introduce fMRI lie-detection evidence. Wilson v. Corestaff Services L.P., 900 N.Y.S.2d 639 (May 14, 2010); Maryland v. Smith, Case No. 106589C (Maryland Cir. Ct. 201).

[32] The district court adopted a detailed report and recommendation by Magistrate Judge Tu M. Pham. United States v. Semrau, 2010 WL 6845092 (W.D. Tenn., June 1, 2010).

[33] United States v. Semrau, 693 F.3d 510 (6th Cir. 2012).

evidence and the courts' analyses and then discusses conceptual problems with the foundational studies underlying the evidence.

The defendant, Lorne Semrau, was charged with multiple counts of health-care fraud and money laundering.[34] Semrau, a licensed psychologist, owned two companies that provided mental-health services to patients in nursing homes. The allegations against Semrau concerned a scheme to defraud Medicare, Medicaid, and other government programs by submitting false and fraudulent claims for payments, totaling approximately three million dollars. In order to convict Semrau at trial, the prosecution had to prove that Semrau (1) 'knowingly devised a scheme . . . to defraud,' (2) 'executed or attempted to execute this scheme,' and (3) 'acted with intent to defraud.'[35]

In his defense, Semrau claimed that he did not intend to defraud with his billing practices, and that the mistakes in his billing resulted from confusion caused by unclear billing codes and instructions from insurance providers. To support his defense, Semrau sought to introduce fMRI results purporting to indicate that he was not lying when he claimed the billing practices at issue resulted from confusion rather than an intent to defraud.

Semrau attempted to introduce the evidence through the expert testimony of Dr. Steven Laken, the CEO of the company (Cephos) that produced the fMRI lie-detection results for Semrau.[36] According to Laken, Semrau was scanned according to a protocol developed by

---

[34] Semrau, 2010 WL 6845092, slip op. at pp. 1-2.

[35] Id.

[36] Id.

Cephos. The testing consisted of scanning Semrau while asking him three types of questions: neutral, control, and 'specific incident questions' [SIQs]. The neutral questions (e.g., 'Are you over age 18?') were used to establish a baseline for Semrau; the control questions (e.g., 'Do you gossip?' were used to fill empty space and were not considered part of the analysis; the SIQs concerned the allegations regarding the fraudulent billing practices. In the scanner, Semrau was instructed to 'lie' or 'tell the truth' when presented with SIQs and to tell the truth when presented with control and neutral questions.[37] According to Laken, the fMRI results suggest that Semrau was being truthful in his claims about the billing practices.[38]

The trial court concluded that Laken's testimony and the fMRI evidence should be excluded under Federal Rule of Evidence 702.[39] The court analyzed the evidence according to the factors articulated in *Daubert*: whether the evidence was (1) tested, (2) subject to peer review and publication, (3) had known error rates and standards of control, and (4) was based on principles and methods generally accepted in the relevant scientific community.[40] The court noted, with regard to the first two factors, that the fMRI technique had been tested in a number of studies (including some co-authored by Laken) that were published in peer-reviewed scientific journals.

[37] Id. at 6 ('In each fMRI scan, Dr. Semrau was visually instructed to "Lie" or to tell the "Truth" in response to each SIQ. He was told to respond truthfully to the neutral and control questions.')

[38] Id. at 6-8. In fact, Laken initially concluded that Semrau was being deceptive with regard to one set of question. He re-tested Semrau after concluding that these results may have been caused by fatigue. After re-testing, he concluded that Semrau was being truthful. Laken also explained that his conclusion applies only to the general subject matter of the questions and not to any individual questions.

[39] Id. at 10-14.

[40] Daubert v. Merrell Dow Pharm., Inc. 509 U.S. 579 (1993).

The court, however, discussed several empirical problems with regard to the third factor: error rates and standards of control.[41] First, that court noted the error rates in the laboratory studies may not translate to real-world settings, for which there are no known error rates. Second, the court explained that there are no industry standards for such testing and that Laken appeared to deviate from his own protocol in testing Semrau. Third, Semrau differed in important respects from participants in the studies. Fourth, the studies themselves were limited in several respects, including (1) a small number of nondiverse subjects; (2) inconsistent results; and (3) the unknown effects of countermeasures. In addition, with regard to the fourth factor, the court concluded that the principles and methods of fMRI lie detection were not generally accepted among neuroscientists. For all of these reasons, the court held that the fMRI results were not sufficiently reliable to be admitted at Semrau's criminal trial.

Although the court excluded the evidence, the court explained that similar evidence may be admissible in the future, even if there remain unknown error rates for real-world settings:

> Should fMRI-based lie detection undergo further testing, development, and peer review, improve upon standards controlling the technique's operation, and gain acceptance by the scientific community for use in the real world, this methodology may be found to be admissible even if the error rate is not able to be quantified in a real world setting.[42]

---

[41] Semrau, slip op. at 13.

[42] Id. at 12 n.18.

Semrau was convicted at trial of three courts of health-care fraud and appealed his conviction.[43] One issue that he raised on appeal was the exclusion of his fMRI evidence. The appellate court upheld the court's ruling:

> After carefully reviewing the scientific and factual evidence, we conclude that the district court did not abuse its discretion in excluding the evidence under Federal Rule of Evidence 702 because the technology had not been fully examined in 'real world' settings and the testing administered to Dr. Semrau was not consistent with tests done in research studies.[44]

The court also explained that Federal Rule of Evidence 403 provided an independent basis for excluding the evidence because

> the prosecution did not know about the test before it was conducted, constitutional concerns caution against admitting lie detection tests to bolster witness credibility, and the test results do not purport to indicate whether Dr. Semrau was truthful about any single statement.[45]

Similar to the district court, the appellate court also noted the possible admissibility of such evidence in the future.[46]

---

[43] United States v. Semrau, 693 F.3d 510 (6th Cir. 2012).

[44] Id. at 516.

[45] Id.

[46] Id. at n.12 ('Though we are not at that point today, we recognize that as science moves forward the balancing of Rule 403 may well lean toward finding the probative evidence for some advancing technology is sufficient.').

The courts' reasons were sufficient for excluding the evidence and they discussed several important empirical limitations with the evidence. Absent from the discussion at either level, however, were the *conceptual* issues raised by the underlying studies. This is problematic because the conceptual issues also affect inferences that may be drawn from the evidence. And they do so in a way that is even more fundamental than the empirical issues noted by the courts.

The single most important conceptual issue is whether the underlying studies are measuring lies.[47] And, for the reasons discussed above, they are not. This can be shown by examining the primary study submitted by Semrau and discussed by the trial and appellate courts. This study, *Functional MRI Detection of Deception After Committing a Mock Sabotage Crime*, of which Laken (Semrau's expert) is a co-author, reports impressive results in using fMRI in order to determine which test subjects participated in a mock crime.[48] Although the study describes 'lies' and 'lying' behavior by participants in each group, the participants were not in fact lying and the studies were not measuring brain activity during lying behavior. We need to look at the details of experiment to see why.

Participants were randomly divided into two groups: 'mock crime' and 'no crime.'[49] Those in the mock-crime group were given instructions in how to carry out 'a mock crime.' These instructions told participants to pick up an envelope, 'destroy evidence' by breaking a CD,

---

[47] The studies also raise the other conceptual issue discussed above—namely, the mistaken equation of lies with deception. Laken testified, for example, that 'A lie is the intentional act of deceit.' Semrau, Hearing Transcript at 159. For the reasons discussed above, this conception is under-inclusive with regard to lying. See supra note 23.

[48] Kozel, supra note 16. As discussed above, see supra note 16 and accompanying text, the study reports a 100 percent success rate in identifying the participants in a mock crime (but a 33 percent rate in identifying the 'no crime' participants).

[49] Id. at 221-22.

and 'steal' a second CD and bring it to the researcher giving the instructions. For those who successfully completed this portion of the study, '[t]hroughout the remainder of the study, the participants were to report that they went to pick up the envelope but did nothing else (i.e., they were to lie about the mock-crime task).'[50] As the authors explain: 'any deviation from that would result in the subject being excluded from the study.'[51] Those in the 'no crime' group were also instructed to 'lie.' They were told that they would need an alibi for the time a CD was being destroyed and that they should answer that they were retrieving an envelope in another building at the time. For those in the no-crime group, 'they were to answer as if they picked up the envelope but did not damage the CD or visit the crime room.'[52] As the authors explain: '[t]he purpose of having the participants lie about picking up the envelope was to have questions during the fMRI evaluation in which participants in this group were lying.'[53]

What, exactly, is the problem with this experimental set-up? The problem is that lying requires more than simply saying something believed to be false. It also has a normative component. Lying takes place when the speaker communicates something believed to be false in a context where there is a norm against doing so. Violating this norm is what makes prototypical cases of lying constitute 'lies' and not other performative acts (such as jokes, game playing, acting, and so on). Importantly, there is no such norm in play during the experiments. Indeed, the participants were instructed to say false things. If anything, the operative norm during the

---

[50] Id. at 222.

[51] Id.

[52] Id.

[53] Id.

experiment was that they should say false things. In doing so, they were following the rules of the experiment and would be excluded from the study if they failed to do so.[54] The behavior being measured in the experiments more closely resembled game playing, role playing, or acting rather than genuine cases of lying.

This mismatch between the studies and actual lying limits the ability to drawn inferences about whether behavior in real-world cases constitutes lying. Because the studies were not measuring lying, the relationship between that data and actual lies is essentially unknown. The different context between the studies and the statements by Semrau—his knowingly false statements about his billing practices would be lies—provides another limitation on using the former to drawn inferences about Semrau. In addition to the several empirical limitations discussed by the courts, the fact that the studies[55] were not even measuring lies further limits their utility for real-world settings. Moreover, this conceptual limitation cuts deep. It suggests that the empirical results of such studies do not provide support for additional research into lying; rather, they are built upon a conceptual mistake that runs to their foundation. Studies that purport to illuminate lying need to either measure actual lying or provide some reason to think that what is being measured does correlate with lying.

III. Recent Studies

[54] Id. Indeed, under this set-up, actually telling the truth when they were instructed to say falsehoods would have been closer to real lying.

[55] The other studies relied on by Laken and his colleagues involved the same experimental conditions. See Kozel, supra note 15; F. Andrew Kozel et al., 'Replication of Functional MRI Detection of Deception' (2009) 2 *Open Forensic Science Journal* 6.

The conceptual problems that we identified in *Minds, Brains, and Law* occur throughout the fMRI studies on lie detection that had been published at that time. Since we surveyed the literature for our book, the conceptual problems have persisted. This Part discusses five examples, focusing on the most important conceptual issue: the criteria for what constitutes a lie and whether the studies are in fact measuring lies.[56]

1. Ofen et al. (2016)[57]—This study purports to measure brain activity correlated with different types of lies.[58] The experiment distinguishes between two different types of false statements: those about past experiences (episodic memories) and those about personal beliefs. The study reports that some brain regions exhibited increased activation for both types of false statements (in the frontal and parietal cortex), but some regions showed differences between the two types of statements (right temporal pole, precuneus, and right amygdala).[59] Based on their experiments, the authors conclude: '[t]he findings here suggest that the brain signature of a lie is influenced by the type of knowledge one is lying about.'[60] But were the participants lying? And therefore were the researchers measuring the 'brain signature of a lie'? It does not appear so. The participants were instructed to either answer truthfully or falsely to questions about experiences (e.g., 'Have you ever ridden a horse?') or opinions (e.g., 'Do you believe the death penalty is justifiable?').[61]

---

56 As with prior studies, the newer studies also continue to conflate lying and deception.

57 Noa Ofen et al, 'Neural Correlates of Deception: Lying About Past Events and Personal Beliefs (2016) Social Cognitive and Affective Neuroscience 1, available at: doi: 10.1093/scan/nsw151

58 The study defines a lie as a false statement with an intent to deceive. Id. at 4.

59 Id. at 1.

60 Id. at 28.

61 Id. at 6. If a participant had intentionally answered contrary to the instruction, then this would better match the normative component of real lying, even if the participant answered truthfully (when instructed to answer falsely).

In this context, they were not violating a norm of communication in speaking falsely. They were instructed to do so; it was a rule of the game that they were playing in the scanner.

2. Jiang et al. (2015)[62]—This study created 'whole-brain functional connectivity networks' for each participant and then used a 'multivariate pattern analysis approach . . . to distinguish lie-telling from truth-telling.'[63] The study reports that 'lie-telling' could be differentiated from 'truth-telling' with an accuracy rate of 82.81 percent (85.94 percent for lies, and 79.69 percent for truthful statements).[64] The experiment, however, was not measuring lying. Participants in the study were instructed to select three pictures from a group of ten and were then asked during scanning whether a presented picture was one they had selected. During the 'truth condition,' participants were 'required to give accurate, honest responses.'[65] During the 'lie condition, the participants were required to devise a strategy to deceive others, with the goal of lying skillfully and avoiding detection.'[66] Once again, the norm against communicating something false does not apply in this setting and participants likely did not believe such a norm applied (they were explicitly told otherwise). Thus, the participants were not lying.

---

[62] Weixiong Jiang et al., 'Decoding the Processing of Lying Using Functional Connectivity MRI' (2015) 11 *Behavioral and Brain Functions* 1.

[63] The study also equates lying with deception. Id. at 1, 6.

[64] Id. at 4.

[65] Id. at 2.

[66] Id.

3. Sip (2013)[67]—This study explored whether participants' own beliefs about whether their lies could be detected would affect the underlying brain regions involved in producing such lies. Participants were instructed to commit a mock crime (taking earphones or a memory stick from one of two rooms). During the scanning phase, the participants were introduced to an 'interrogator' and also hooked up to a fake 'lie detector' machine.[68] They were told the 'lie detector' would allow the interrogator 'to discriminate between honest and deceptive responses' but that it would be turned on for only half of the session (and they would be told when).[69] This would allow the researchers to examine differences in responses when the participants believed their lies could be detected and when they were not being analyzed. The authors report 'a significant interaction between the effects of deception and belief in the left temporal pole and right hippocampus/ parahippocampal gyrus . . . when participants believed their false claims could be detected, but not when they believed the lie detector was switched-off.'[70] Unlike the above studies, the participants were not 'explicitly instructed to produce false statements'; however, they were instructed that if, after the interrogation, 'the interrogator could not tell whether they had taken an object, then they would get to keep the object.'[71] Despite this experimental wrinkle (i.e., not telling subjects to say false things), the participants nevertheless do not appear to be lying. Telling participants that they their goal is to trick the interrogator and

[67] Kamila E. Sip et. al., 'When Pinocchio's Nose Does not Grow: Belief Regarding Lie-Detectability Modulates Production of Deception' (2013) 7 *Frontiers in Human Neuroscience* 1.

[68] Id. at 3. This study also equates lying and deception.

[69] Id.

[70] Id. at 1.

[71] Id. at 3-4.

the 'lie detector,' plainly conveys to the participants that they are expected to say false things. Or at least that there is nothing normatively problematic about them doing so, unlike in genuine cases of lying. The findings in this study about the participants' *beliefs* about lie detection are certainly relevant and interesting, but the experiment is not measuring lying.

4. Jiang et al. (2013)[72]—This study explored deception among participants with sufficiently high scores on a diagnostic questionnaire for antisocial personality disorder. The study reported increased activation in several brain regions during 'lie-telling.'[73] This experiment, however, employed the same select-three-pictures paradigm discussed above (with instructions to answer truthfully when instructed to do so and to 'lie' when instructed to do so).[74] Thus, the experiment was not measuring lying, even in the 'lie-telling' condition.

5. Cui et al. (2013)[75]—This study instructed subjects to participate in a 'mock murder.' The participation consisted in creating a story with several details about a fictitious murder (for example, identity of victim, weapon used, etc.). Other subjects (the 'innocents') did not participate in constructing a story about a murder. Both groups were scanned and asked details about the 'murder,' with innocents instructed to honestly deny knowledge and the 'mock

---

[72] W. Jiang et al., 'A Functional MRI Study of Deception Among Offenders With Antisocial Personality Disorder' (2013) 244 *Neuroscience* 90. The study also equates lying with detection.

[73] Id. at 96 ('The pattern of brain activity in the frontal-anterior cingulated-parietal regions associated with lying that was observed in our study is largely consistent with the findings of previous fMRI deception studies.')

[74] Id. at 93. See supra notes 62-66 and accompanying text.

[75] Qian Cui et al., 'Detection of Deception Based on fMRI Activation Patterns Underlying the Production of a Deceptive Response and Receiving Feedback about the Success of the Deception After a Mock Murder Crime' (2013) *Social Cognitive and Affective Neuroscience* 1.

murders' instructed to lie.[76] The authors report different activation in brain areas in the two groups. Specifically, 'the right VLPF [ventral lateral prefrontal] contributed the most to diagnosing the participants' real identities'—'Based on the activity of this area, 81.25% of the "murderers" and 81.25% of the "innocents" were correctly diagnosed.'[77] Subjects in this study also appear not to be lying. As with the other studies, participants do not appear to be violating a norm against communicating something believed to be false—they were instructed to do precisely that—and thus their false statements are not lies.

*   *   *

The purpose in critiquing these studies—and with the previous ones discussed in *Minds, Brains, and Law*—is to illustrate the importance of conceptual issues for law and neuroscience. The empirical limitations on using neuroscientific evidence in court are real and have received important treatment from scholars in a variety of fields. The success of empirical investigations, however, also depends on clear conceptual presuppositions underlying the investigations. Philosophical attention to these presuppositions can help to clarify the issues and correct mistaken assumptions. This suggests an important role for philosophy to play in neurolaw, not only for issues involving high-level theoretical issues (such as free will and criminal responsibility) but also at the level of evidence and proof on the ground. The conceptual issues involved with fMRI lie detection are one such example. In critiquing the conceptual underpinnings of some of the scientific literature, the aim has not been to challenge the

[76] Id. at 2-3.

[77] Id. at 7. The study also reports increased accuracy rates when subjects were given feedback about whether their prior answer were judged to be deceptive or not. Id. (93.75 percent of the 'murderers' and 87.5 percent of the 'innocents').

possibility of brain-based lie detection. Rather, the aim has been to ensure that any evidence that is produced for legal purposes is probative regarding something the law cares about—in this instance, lying rather than other types of behavior.