

THE NEW YORKER

A REPORTER AT LARGE

DUPED

Can brain scans uncover lies?

by Margaret Talbot

JULY 2, 2007


The most egregious liar I ever knew was someone I never suspected until the day that, suddenly and irrevocably, I did. Twelve years ago, a young man named Stephen Glass began writing for *The New Republic*, where I was an editor. He quickly established himself as someone who was always onto an amusingly outlandish story—like the time he met some Young Republican types at a convention, gathered them around a hotel-room minibar, then, with guileless ferocity, captured their boorishness in print. I liked Steve; most of us who worked with him did. A baby-faced guy from suburban Chicago, he padded around the office in his socks. Before going on an errand, Steve would ask if I wanted a muffin or a sandwich; he always noticed a new scarf or a clever turn of phrase, and asked after a colleague's baby or spouse. When he met with editors to talk about his latest reporting triumph, he was self-effacing and sincere. He'd look us in the eye, wait for us to press him for details, and then, without fidgeting or mumbling, supply them.

One day, the magazine published an article by Steve about a teen-ager so diabolically gifted at hacking into corporate computer networks that C.E.O.s paid him huge sums just to stop messing with them. A reporter for the online edition of *Forbes* was assigned to chase down the story.

You can see how Steve's journalism career unravelled if you watch the movie "Shattered Glass": *Forbes* challenged the story's veracity, and Steve—after denying the charges, concocting a fake Web site, and enlisting his brother to pose as a victimized C.E.O.—finally confessed that he'd made up the whole thing. Editors and reporters at the magazine investigated, and found that Steve had been inventing stories for at least a year. The magazine disavowed twenty-seven articles.

After Steve's unmasking, my colleagues and I felt ashamed of our gullibility. But maybe we shouldn't have. Human beings are terrible lie detectors. In academic studies, subjects asked to distinguish truth from lies answer correctly, on average, fifty-four per cent of the time. They are better at guessing when they are being told the truth than when they are being lied to, accurately classifying only forty-seven per cent of lies, according to a recent meta-analysis of some two hundred deception studies, published by Bella DePaulo, of the University of California at Santa Barbara, and Charles Bond, Jr., of Texas Christian University. Subjects are often led astray by an erroneous sense of how a liar behaves. "People hold a stereotype of the liar—as tormented, anxious, and conscience-stricken," DePaulo and Bond write. (The idea that a liar's anxiety will inevitably become manifest can be found as far back as the ancient Greeks, Demosthenes in particular.) In fact, many liars experience what deception researchers call "duping delight."

Aldert Vrij, a psychologist at the University of Portsmouth, in England, argues that there is no such thing as "typical" deceptive behavior—"nothing as obvious as Pinocchio's growing nose." When people tell complicated lies, they



frequently pause longer and more often, and speak more slowly; but if the lie is simple, or highly polished, they tend to do the opposite. Clumsy deceivers are sometimes visibly agitated, but, over all, liars are less likely to blink, to move their hands and feet, or to make elaborate gestures—perhaps they deliberately inhibit their movements. As DePaulo says, “To be a good liar, you don’t need to know what behaviors really separate liars from truth-tellers, but what behaviors people *think* separate them.”

A liar’s testimony is often more persuasive than a truth-teller’s. Liars are more likely to tell a story in chronological order, whereas honest people often present accounts in an improvised jumble. Similarly, according to DePaulo and Bond, subjects who spontaneously corrected themselves, or said that there were details that they couldn’t recall, were more likely to be truthful than those who did not—though, in the real world, memory lapses arouse suspicion.

People who are afraid of being disbelieved, even when they are telling the truth, may well look more nervous than people who are lying. This is bad news for the falsely accused, especially given that influential manuals of interrogation reinforce the myth of the twitchy liar. “Criminal Interrogation and Confessions” (1986), by Fred Inbau, John Reid, and Joseph Buckley, claims that shifts in posture and nervous “grooming gestures,” such as “straightening hair” and “picking lint from clothing,” often signal lying. David Zulawski and Douglas Wicklander’s “Practical Aspects of Interview and Interrogation” (1992) asserts that a liar’s movements tend to be “jerky and abrupt” and his hands “cold and clammy.” Bunching Kleenex in a sweaty hand is another damning sign—one more reason for a sweaty-palmed, Kleenex-bunching person like me to hope that she’s never interrogated.

Maureen O’Sullivan, a deception researcher at the University of San Francisco, studies why humans are so bad at recognizing lies. Many people, she says, base assessments of truthfulness on irrelevant factors, such as personality or appearance. “Baby-faced, non-weird, and extroverted people are more likely to be judged truthful,” she says. (Maybe this explains my trust in Steve Glass.) People are also blinkered by the “truthfulness bias”: the vast majority of questions we ask of other people—the time, the price of the breakfast special—are answered honestly, and truth is therefore our default expectation. Then, there’s the “learning-curve problem.” We don’t have a refined idea of what a successful lie looks and sounds like, since we almost never receive feedback on the fibs that we’ve been told; the co-worker who, at the corporate retreat, assured you that she loved your presentation doesn’t usually reveal later that she hated it. As O’Sullivan puts it, “By definition, the most convincing lies go undetected.”

Maybe it’s because we’re such poor lie detectors that we have kept alive the dream of a foolproof lie-detecting machine. This February, at a conference on deception research, in Cambridge, Massachusetts, Steven Hyman, a psychiatrist and the provost of Harvard, spoke of “the incredible hunger to have some test that separates truth from deception—in some sense, the science be damned.”

This hunger has kept the polygraph, for example, in widespread use. The federal government still performs tens of thousands of polygraph tests a year—even though an exhaustive 2003 National Academy of Sciences report concluded that research on the polygraph’s efficacy was inadequate, and that when it was used to investigate a specific incident after the fact it performed “well above chance, though well below perfection.” Polygraph advocates cite accuracy estimates of ninety per cent—which sounds impressive until you think of the people whose lives might be ruined by a machine that fails one out of ten times. The polygraph was judged thoroughly unreliable as a screening tool; its accuracy in “distinguishing actual or potential security violators from innocent test takers” was deemed “insufficient to justify reliance on its use.” And its success in criminal investigations can be credited, in no small part, to the intimidation factor. People who believe that they are in the presence of an infallible machine sometimes confess, and this is counted as an achievement of the polygraph. (According to law-enforcement lore, the police have used copy machines in much the same way: They tell a suspect to place his hand on a “truth machine”—a copier in which the paper has “LIE” printed on it. When the photocopy emerges, it shows the suspect’s hand with “LIE” stamped on it.)

Over the past two decades, inventors have attempted to supplant the polygraph with new technologies: voice-stress analysis; thermal imaging of the face; and, most recently and spectacularly, brain imaging. Though these methods remain in an embryonic stage of development, they have already been greeted with considerable enthusiasm, especially in America. Private companies are eager to replace traditional modes of ascertaining the truth—such as the jury system—with a machine that can be patented and sold. And law-enforcement agencies yearn to overcome the problem of

suspects who often remain maddeningly opaque, even in the face of sustained interrogation. Although one immediate result of the September 11th attacks was the revival of an older, and even more controversial, form of interrogation—torture—the war on terror has also inflamed the desire for a mind-reading machine.

Not long ago, I met with an entrepreneur named Joel Huizenga, who has started a company, based in San Diego, called No Lie MRI. Most methods of lie detection look at the activity of the sympathetic nervous system. The polygraph, for instance, is essentially an instrument for measuring stress. Heart and respiration rates, blood volume, and galvanic skin response—a proxy for palm sweat—are represented as tracings on graph paper or on a screen, which fluctuate with every heartbeat or breath. The method that Huizenga is marketing, which employs a form of body scanning known as functional magnetic resonance imaging, or fMRI, promises to look inside the brain. “Once you jump behind the skull, there’s no hiding,” Huizenga told me.

Functional MRI technology, invented in the early nineties, has been used primarily as a diagnostic tool for identifying neurological disorders and for mapping the brain. Unlike MRIs, which capture a static image, an fMRI makes a series of scans that show changes in the flow of oxygenated blood preceding neural events. The brain needs oxygen to perform mental tasks, so a rise in the level of oxygenated blood in one part of the brain can indicate cognitive activity there. (Blood has different magnetic properties when it is oxygenated, which is why it is helpful to have a machine that is essentially a big magnet.) Brain-scan lie detection is predicated on the idea that lying requires more cognitive effort, and therefore more oxygenated blood, than truth-telling.

Brain scanning promises to show us directly what the polygraph showed us obliquely. Huizenga expects his company to be a force for justice, exonerating customers who are, as he put it, “good people trying to push back the cruel world that is indicting them unfairly.” Brain scans already have clout in the courtroom; during death-penalty hearings, judges often allow images suggesting neurological impairment to be introduced as mitigating evidence. In theory, an improved method of lie detection could have as profound an impact as DNA evidence, which has freed more than a hundred wrongly accused people since its introduction, in the late eighties. If Huizenga has perfected such a technology, he’s onto something big.

At Huizenga’s suggestion, we met at a restaurant called the Rusty Pelican, on the Pacific Coast Highway, in Newport Beach. A television screen on one wall showed a surfing contest; Huizenga, who is fifty-three, with dirty-blond hair in a boyish cut, is a surfer himself. He has a bachelor’s degree from the University of Colorado, a master’s degree in biology from Stony Brook, and an M.B.A. from the University of Rochester. No Lie is Huizenga’s second startup. The first, ISCHEM Corporation, uses body scanning to look for plaque in people’s arteries. Before that, he worked for Pantox, a company that offers blood tests to gauge a person’s antioxidant levels.

After we sat down, Huizenga recounted the origins of No Lie. A few years ago, he came across an item in the *Times* about some tantalizing research conducted by Daniel Langleben, a psychiatrist and neuroscientist at the University of Pennsylvania. Subjects were placed inside an fMRI machine and told to make some true statements and some false ones. Brain scans taken while the subjects were lying frequently showed a significantly increased level of activity in three discrete areas of the cerebral cortex. Langleben suggested that “intentional deception” could be “anatomically localized” by fMRI scanning. Huizenga immediately saw a business opportunity. “I jumped on it,” he told me. “If I wasn’t here sitting in front of you, somebody else would be.”

The Web site for No Lie claims that its technology, which is based on the Penn protocol, “represents the first and only direct measure of truth verification and lie detection in human history!” No Lie just started offering tests commercially, and has charged about a dozen clients approximately ten thousand dollars apiece for an examination. (No Lie sends customers to an independent imaging center in Tarzana, a suburb of Los Angeles, to insure that “quality testing occurs according to standardized test protocols.”) Some of these initial clients are involved in civil and criminal cases; the first person to use the service, Harvey Nathan, was accused in 2003 of deliberately setting fire to a deli that he owns in South Carolina. A judge dismissed the charges, but Nathan wanted to bring suit against his insurance company, and he thought that documented evidence of his innocence would further his cause. So in December he flew to California and took No Lie’s test. He passed. Nathan said, “If I hadn’t, I would have jumped from the seventeenth floor of the hotel where I was staying. How could I have gone back to South Carolina and said, ‘Oh that machine must not have worked right?’ I believed in it then and I believe in it now.” Nathan’s exam was filmed for the Discovery Channel, which may

soon launch a reality series centering on brain-scanning lie detection.

Several companies have expressed interest in No Lie's services, Huizenga told me. (He would not name them.) He said that he will be able to accommodate corporate clients once he has signed deals with other scanning facilities; he is in talks with imaging centers in a dozen cities, including New York and Chicago. No Lie also plans to open a branch in Switzerland later this year.

Huizenga has been criticized for his company's name, but he said, "It's not about being dignified—it's about being remembered." He believes that the market for fMRI-based lie detection will one day exceed that of the polygraph industry, which brings in hundreds of millions of dollars annually. Investment analysts say that it is too soon to judge if Huizenga's optimism is warranted, but No Lie has attracted some prominent backing. One of its prime investors is Alex Hart, the former C.E.O. of MasterCard International, who is also serving as a management consultant. And it has a "scientific board" consisting of four paid advisers, among them Terrence Sejnowski, the director of the Crick-Jacobs Center for theoretical and computational biology at the Salk Institute. In an e-mail, Sejnowski explained that he offers counsel on "advanced signal processing and machine-learning techniques that can help improve the analysis of the data and the accuracy of the performance." He said of No Lie, "The demand is there, and to succeed as a company the new technology only needs to be better than existing approaches."

Huizenga speaks of his company's goals in blunt terms. "What do people lie about?" he asked me. "Sex, power, and money—probably in that order." (The company's Web site recommends No Lie's services for "risk reduction in dating," "trust issues in interpersonal relationships," and "issues concerning the underlying topics of sex, power, and money.") "Parents say, 'Yes, this is *perfect* for adolescents,' " he went on. "People who are dating say, 'Yes, this is great for dating, because people never tell you the truth.' "

He said that his company receives dozens of inquiries a week: from divorcing men accused of child abuse; from women wanting to prove their fidelity to jealous spouses or boyfriends; from people representing governments in Africa and the former Soviet republics; from "the Chinese police department." He said that he understood why governments were interested in lie-detection technology. "Look at Joe Stalin," he said. "Joe wanted power, he wanted to be on top. Well, it's hard to murder massive numbers of opponents. People in our government, and in others', need more effective ways of weeding out those who aren't their puppets." Some potential foreign clients had explained to him, he said, that in societies that lacked "civilization, there is not trust, and lie detection could help build that trust." (He wasn't sure about that—he was "mulling it over.") Huizenga said that the United States government was "interested" in the kind of technology offered by No Lie; the company has hired Joel S. Lisker, a former F.B.I. agent, to be its "sales liaison for the federal government." (Lisker declined to be interviewed, saying that his government contacts were "confidential.")

The Pentagon has supported research into high-tech lie detection, including the use of fMRI. The major scientific papers in the field were funded, in part, by the Defense Advanced Research Projects Agency, which develops new technologies for military use, and by the Department of Defense Polygraph Institute, which trains lie-detection experts at Fort Jackson, South Carolina. (The Polygraph Institute underwent a name change in January—it's now the Defense Academy for Credibility Assessment—apparently in deference to new technologies such as fMRI.) Last June, the A.C.L.U. filed several Freedom of Information Act requests in an attempt to learn more about the government's involvement with the technology. Chris Calabrese, an A.C.L.U. lawyer, said that the C.I.A. would neither "confirm nor deny" that it is investigating fMRI applications; the Pentagon produced PowerPoint presentations identifying brain scans as a promising new technology for lie detection. Calabrese went on, "We were motivated by the fact that there are companies trying to sell this technology to the government. This Administration has a history of using questionable techniques of truth verification."

Many scholars also think that Huizenga's effort is premature. Steven Hyman, the Harvard professor, told me that No Lie was "foolish." But the history of lie-detection machines suggests that it would be equally foolish to assume that a few scholarly critics can forestall the adoption of such a seductive new technology. "People are drawn to it," Huizenga said, smiling. "It's a magnetic concept."

In comic books of the nineteen-forties, Wonder Woman, the sexy Amazon superhero, wields a golden "lasso of truth." Anybody she captures is rendered temporarily incapable of lying. Like the golden lasso, the polygraph, its inventors

believed, compelled the body to reveal the mind's secrets. But the connection between the lasso and the lie detector is even more direct than that: Wonder Woman's creator, William Moulton Marston, was also a key figure in the development of the polygraph. Marston, like other pioneers of lie detection, believed that the conscious mind could be circumvented, and the truth uncovered, through the measurement of bodily signals.

This was not a new idea. In 1730, Daniel Defoe published "An Effectual Scheme for the Immediate Preventing of Street Robberies and Suppressing All Other Disorders of the Night," in which he proposed an alternative to physical coercion: "Guilt carries fear always about with it, there is a tremor in the blood of a thief, that, if attended to, would effectually discover him; and if charged as a suspicious fellow, on that suspicion only I would feel his pulse."

In the late nineteenth century, the Italian criminologist Cesare Lombroso invented his own version of a lie detector, based on the physiology of emotion. A suspect was told to plunge his hand into a tank filled with water, and the subject's pulse would cause the level of liquid to rise and fall slightly; the greater the fluctuation, the more dishonest the subject was judged to be.

Lombroso's student Angelo Mosso, a physiologist, noticed that shifts in emotion were often detectable in fair-skinned people in the flushing or blanching of their faces. Based on this observation, he designed a bed that rested on a fulcrum. If a suspect reclining on it told a lie, Mosso hypothesized, resulting changes in blood flow would alter the distribution of weight on the bed, unbalancing it. The device, known as Mosso's cradle, apparently never made it past the prototype.

William Moulton Marston was born in 1893, in Boston. He attended Harvard, where he worked in the lab of Hugo Münsterberg, a German émigré psychologist, who had been tinkering with an apparatus that registered responses to emotions, such as horror and tenderness, through graphical tracing of pulse rates. One student volunteer was Gertrude Stein. (She later wrote of the experience in the third person: "Strange fancies begin to crowd upon her, she feels that the silent pen is writing on and on forever.")

In 1917, Marston published a paper arguing that systolic blood pressure could be monitored to detect deception. As Ken Alder, a history professor at Northwestern, notes in his recent book, "The Lie Detectors: The History of an American Obsession," Münsterberg and Marston's line of inquiry caught the imagination of police detectives, reporters, and law-enforcement reformers across the country, who saw a lie-detecting machine as an alternative not only to the brutal interrogation known as the third degree but also to the jury system. In 1911, an article in the *Times* predicted a future in which "there will be no jury, no horde of detectives and witnesses, no charges and countercharges, and no attorney for the defense. These impediments of our courts will be unnecessary. The State will merely submit all suspects in a case to the tests of scientific instruments."

John Larson, a police officer in Berkeley, California, who also had a doctorate in physiology, expanded on Marston's work. He built an unwieldy device, the "cardio-pneumo-psychograph," which used a standard cuff to measure blood pressure, and a rubber hose wrapped around the subject's chest to measure his breathing. Subjects were told to answer yes-or-no questions; their physiological responses were recorded by styluses that scratched black recording paper on revolving drums.

In 1921, as Alder writes, Larson had his first big chance to test his device. He was seeking to identify a thief at a residence hall for female students at Berkeley. Larson gave several suspects a six-minute exam, in which he asked various questions: "How much is thirty times forty?" "Will you graduate this year?" "Do you dance?" "Did you steal the money?" The result foretold the way in which a polygraph would often "work": as a goad to confession. A student nurse confessed to the crime—a few days after she'd stormed out during the exam.

In the early twenties, another member of the Berkeley police force, Leonarde Keeler, increased the number of physical signs that the lie detector monitored. His portable machine recorded pulse rate, blood pressure, respiration, and "electrodermal response"—again, palm sweat. Today's lie detector looks much like Keeler's eighty-year-old invention. And it bears the same name: the polygraph.

Polygraphs never caught on in Europe. But here their advent coincided with the Prohibition-era crime wave; with a new fascination with the unconscious (this was also the era of experimentation with so-called truth serums); and with the wave of technological innovation that had brought Americans electricity, radios, telephones, and cars. The lie detector quickly insinuated itself into American law enforcement: at the end of the thirties, a survey of thirteen city police

departments showed that they had given polygraphs to nearly nine thousand suspects.

In 1923, Marston tried without success to get a polygraph test introduced as evidence in the Washington, D.C., murder trial of James Alphonso Frye. In its ruling, the Court of Appeals for the D.C. Circuit declared that a new scientific method had to have won “general acceptance” from experts before judges could give it credence. Since this decision, the polygraph has been kept out of most courtrooms, but there is an important exception: about half the states allow a defendant to take the test, generally on the understanding that the charges will be dropped if he passes and the results may be entered as evidence if he fails.

The polygraph became widely used in government and in business, often with dubious results. In the fifties, the State Department deployed the lie detector to help purge suspected homosexuals. As late as the seventies, a quarter of American corporations used the polygraph on their employees. Although Congress banned most such tests when it passed the Polygraph Protection Act, in 1988, the federal government still uses the polygraph for security screenings—despite high-profile mistakes. The polygraph failed to cast suspicion on Aldrich Ames, the C.I.A. agent who spied for the Soviets, and wrongly implicated Wen Ho Lee, the Department of Energy scientist, as an agent of the Chinese government.

One excellent way to gauge the polygraph’s effectiveness would be to compare it with an equally intimidating fake machine, just as a drug is compared with a placebo. But, strangely, no such experiment has ever been performed. In 1917, the year that Marston published his first paper on lie detection, his research encountered strong skepticism. John F. Shepard, a psychologist at the University of Michigan, wrote a review of Marston’s research. Though the physical changes that the machine measured were “an index of activity,” Shepard wrote, the same results “would be caused by so many different circumstances, anything demanding equal activity (intelligence or emotional).” The same criticism holds true today. All the physiological responses measured by the polygraph have causes other than lying, vary greatly among individuals, and can be affected by conscious effort. Breathing is particularly easy to regulate. Advice on how to beat the lie detector is a cottage industry. “Deception Detection: Winning the Polygraph Game” (1991) warns potential subjects, “Don’t complain about a dry mouth. An examiner will interpret this as fear of being found out and will press you even harder.” (Many people do get dry-mouthed when they’re nervous—which is apparently why, during the Inquisition, a suspect was sometimes made to swallow a piece of bread and cheese: if it stuck in his throat, he was deemed guilty.) Other well-known “countermeasures” include taking a mild sedative; using mental imagery to calm yourself; and biting your tongue to make yourself seem anxious in response to random questions.

Why, then, is the polygraph still used? Perhaps the most vexing thing about the device is that, for all its flaws, it’s not pure hokum: a racing pulse and an increased heart rate can indicate guilt. Every liar has felt an involuntary flutter, at least once. Yet there are enough exceptions to insure that the polygraph will identify some innocent people as guilty and some guilty people as innocent.

At the Cambridge conference, Jed S. Rakoff, a United States district judge in New York, told a story about a polygraph and a false confession. Days after September 11th, an Egyptian graduate student named Abdallah Higazy came to the attention of the F.B.I. Higazy had been staying at the Millennium Hotel near Ground Zero on the day of the attacks. A hotel security guard claimed that he had found a pilot’s radio in Higazy’s room. Higazy said that it wasn’t his, and when he appeared before Rakoff he asked to be given a polygraph. As Rakoff recalled, “Higazy very much believed in them and thought it would exonerate him.” During a four-hour interrogation by an F.B.I. polygrapher, Higazy first repeated that he knew nothing about the radio, and then said that maybe it was his. He was charged with lying to the F.B.I. and went to prison. Within a month, a pilot stopped by the hotel to ask about a radio that he had accidentally left there. The security guard who found the radio admitted that it hadn’t been in Higazy’s room; he was prosecuted and pled guilty. Higazy was exonerated, and a subsequent investigation revealed that he had felt dizzy and ill during the examination, probably out of nervousness. But when Higazy asked the polygrapher if anyone had ever become ill during a polygraph test he was told that “it had not happened to anyone who told the truth.”

To date, there have been only a dozen or so peer-reviewed studies that attempt to catch lies with fMRI technology, and most of them involved fewer than twenty people. Nevertheless, the idea has inspired a torrent of media attention, because scientific studies involving brain scans dazzle people, and because mind reading by machine is a beloved

science-fiction trope, revived most recently in movies like “Minority Report” and “Eternal Sunshine of the Spotless Mind.” Many journalistic accounts of the new technology—accompanied by colorful bitmapped images of the brain in action—resemble science fiction themselves. In January, the *Financial Times* proclaimed, “For the first time in history, it is becoming possible to read someone else’s mind with the power of science.” A CNBC report, accompanied by the Eurythmics song “Would I Lie to You?,” showed its reporter entering an fMRI machine, described as a “sure-fire way to identify a liar.” In March, a cover story in the *Times Magazine* predicted transformations of the legal system in response to brain imaging; its author, Jeffrey Rosen, suggested that there was a widespread “fear” among legal scholars that “the use of brain-scanning technology as a kind of super mind-reading device will threaten our privacy and mental freedom.” *Philadelphia* has declared “the end of the lie,” and a *Wired* article, titled “Don’t Even Think About Lying,” proclaimed that fMRI is “poised to transform the security industry, the judicial system, and our fundamental notions of privacy.” Such talk has made brain-scan lie detection sound as solid as DNA evidence—which it most definitely is not.

Paul Bloom, a cognitive psychologist at Yale, believes that brain imaging has a beguiling appeal beyond its actual power to explain mental and emotional states. “Psychologists can be heard grousing that the only way to publish in *Science* or *Nature* is with pretty color pictures of the brain,” he wrote in an essay for the magazine *Seed*. “Critical funding decisions, precious column inches, tenure posts, science credibility, and the popular imagination have all been influenced by fMRI’s seductive but deceptive grasp on our attentions.” Indeed, in the past decade, *Nature* alone has published nearly a hundred articles involving fMRI scans. The technology is a remarkable tool for exploring the brain, and may one day help scientists understand much more about cognition and emotion. But enthusiasm for brain scans leads people to overestimate the accuracy with which they can pinpoint the sources of complex things like love or altruism, let alone explain them.

Brain scans enthrall us, in part, because they seem more like “real” science than those elaborate deductive experiments that so many psychologists perform. In the same way that an X-ray confirms a bone fissure, a brain scan seems to offer an objective measure of mental activity. And, as Bloom writes, fMRI research “has all the trappings of work with great lab-cred: big, expensive, and potentially dangerous machines, hospitals and medical centers, and a lot of people in white coats.”

Deena Skolnick Weisberg, a graduate student at Yale, has conducted a clever study, to be published in the *Journal of Cognitive Neuroscience*, which points to the outsized glamour of brain-scan research. She and her colleagues provided three groups—neuroscientists, neuroscience students, and ordinary adults—with explanations for common psychological phenomena (such as the tendency to assume that other people know the same things we do). Some of these explanations were crafted to be bad. Weisberg found that all three groups were adept at identifying the bad explanations, except when she inserted the words “Brain scans indicate.” Then the students and the regular adults became notably less discerning. Weisberg and her colleagues conclude, “People seem all too ready to accept explanations that allude to neuroscience.”

Some bioethicists have been particularly credulous, assuming that MRI mind reading is virtually a done deal, and arguing that there is a need for a whole new field: “neuroethics.” Judy Illes and Eric Racine, bioethicists at Stanford, write that fMRI, by laying bare the brain’s secrets, may “fundamentally alter the dynamics between personal identity, responsibility, and free will.” A recent article in *The American Journal of Bioethics* asserts that brain-scan lie detection may “force a reexamination of the very idea of privacy, which up until now could not reliably penetrate the individual’s cranium.”

Legal scholars, for their part, have started debating the constitutionality of using brain-imaging evidence in court. At a recent meeting of a National Academy of Sciences committee on lie detection, in Washington, D.C., Hank Greely, a Stanford law professor, said, “When we make speculative leaps like these . . . it increases, sometimes in detrimental ways, the belief that the technology works.” In the rush of companies like No Lie to market brain scanning, and in the rush of scholars to judge the propriety of using the technology, relatively few people have asked whether fMRIs can actually do what they either hope or fear they can do.

Functional MRI is not the first digital-age breakthrough that was supposed to supersede the polygraph. First, there was “brain fingerprinting,” which is based on the idea that the brain releases a recognizable electric signal when processing a memory. The technique used EEG sensors to try to determine whether a suspect retained memories related to a crime—an image of, say, a murder weapon. In 2001, *Time* named Lawrence Farwell, the developer of brain fingerprinting,

one of a hundred innovators who “may be the Picassos or the Einsteins of the 21st century.” But researchers have since noted a big drawback: it’s impossible to distinguish between brain signals produced by actual memories and those produced by imagined memories—as in a made-up alibi.

After September 11th, another technology was widely touted: thermal imaging, an approach based on the finding that the area around the eyes can heat up when people lie. The developers of this method—Ioannis Pavlidis, James Levine, and Norman Eberhardt—published journal articles that had titles like “Seeing Through the Face of Deception” and were accompanied by dramatic thermal images. But the increased blood flow that raises the temperature around the eyes is just another mark of stress. Any law-enforcement agency that used the technique to spot potential terrorists would also pick up a lot of jangly, harmless travellers.

Daniel Langleben, the Penn psychiatrist whose research underpins *No Lie*, began exploring this potential new use for MRIs in the late nineties. Langleben, who is forty-five, has spent most of his career studying the brains of heroin addicts and hyperactive boys. He developed a side interest in lying partly because his research agenda made him think about impulse control, and partly because his patients often lied to him. Five years ago, Langleben and a group of Penn colleagues published the study on brain scanning and lie detection that attracted Huizenga’s attention. In the experiment, which was written up in *Neuroimage*, each of twenty-three subjects was offered an envelope containing a twenty-dollar bill and a playing card—the five of clubs. They were told that they could keep the money if they could conceal the card’s identity when they were asked about it inside an MRI machine. The subjects pushed a button to indicate yes or no as images of playing cards flashed on a screen in front of them. After Langleben assembled the data, he concluded that lying seemed to involve more cognitive effort than truth-telling, and that three areas of the brain generally became more active during acts of deception: the anterior cingulate cortex, which is associated with heightened attention and error monitoring; the dorsal lateral prefrontal cortex, which is involved in behavioral control; and the parietal cortex, which helps process sensory input. Three years later, Langleben and his colleagues published another study, again involving concealed playing cards, which suggested that lying could be differentiated from truth-telling in individuals as well as in groups. The fMRI’s accuracy rate for distinguishing truth from lies was seventy-seven per cent.

Andrew Kozel and Mark George, then at the Medical University of South Carolina, were doing similar work at the time; in 2005, they published a study of fMRI lie detection in which thirty people were instructed to enter a room and take either a watch or a ring that had been placed there. Then, inside a scanner, they were asked to lie about which object they had taken but to answer truthfully to neutral questions, such as “Do you like chocolate?” The researchers distinguished truthful from deceptive responses in ninety per cent of the cases. (Curiously, Kozel’s team found that liars had heightened activity in different areas of the brain than Langleben did.)

Langleben and Kozel weren’t capturing a single, crisp image of the brain processing a lie; an fMRI’s record of a split-second event is considered unreliable. Instead, they asked a subject to repeat his answer dozens of times while the researchers took brain scans every couple of seconds. A computer then counted the number of “voxels” (the 3-D version of pixels) in the brain image that reflected a relatively high level of oxygenated blood, and used algorithms to determine whether this elevated activity mapped onto specific regions of the brain.

One problem of fMRI lie detection is that the machines, which cost about three million dollars each, are notoriously finicky. Technicians say that the scanners often have “bad days,” in which they can produce garbage data. And a subject who squirms too much in the scanner can invalidate the results. (Even moving your tongue in your mouth can cause a problem.) The results for four of the twenty-three subjects in Langleben’s first study had to be thrown out because the subjects had fidgeted.

The Langleben studies also had a major flaw in their design: the concealed playing card came up only occasionally on the screen, so the increased brain activity that the scans showed could have been a result not of deception but of heightened attention to the salient card. Imagine that you’re the research subject: You’re lying on your back, trying to hold still, probably bored, maybe half asleep, looking at hundreds of cards that don’t concern you. Then, at last, up pops the five of clubs—and your brain sparks with recognition.

Nearly all the volunteers for Langleben’s studies were Penn students or members of the academic community. There were no sociopaths or psychopaths; no one on antidepressants or other psychiatric medication; no one addicted to alcohol or drugs; no one with a criminal record; no one mentally retarded. These allegedly seminal studies look exclusively at

unproblematic, intelligent people who were instructed to lie about trivial matters in which they had little stake. An incentive of twenty dollars can hardly be compared with, say, your freedom, reputation, children, or marriage—any or all of which might be at risk in an actual lie-detection scenario.

The word “lie” is so broad that it’s hard to imagine that any test, even one that probes the brain, could detect all forms of deceit: small, polite lies; big, brazen, self-aggrandizing lies; lies to protect or enchant our children; lies that we don’t really acknowledge to ourselves as lies; complicated alibis that we spend days rehearsing. Certainly, it’s hard to imagine that all these lies will bear the identical neural signature. In their degrees of sophistication and detail, their moral weight, their emotional valence, lies are as varied as the people who tell them. As Montaigne wrote, “The reverse side of the truth has a hundred thousand shapes and no defined limits.”

Langleben acknowledges that his research is not quite the breakthrough that the media hype has suggested. “There are many questions that need to be looked into before we know whether this will work as lie detection,” he told me. “Can you do this with somebody who has an I.Q. of ninety-five? Can you do it with somebody who’s fifty or older? Somebody who’s brain-injured? What kinds of real crimes could you ask about? What about countermeasures? What about people with delusions?”

Nevertheless, the University of Pennsylvania licensed the pending patents on his research to No Lie in 2003, in exchange for an equity position in the company. Langleben didn’t protest. As he explained to me, “It’s good for your résumé. We’re encouraged to have, as part of our portfolio, industry collaborations.” He went on, “I was trying to be a good boy. I had an idea. I went to the Center of Technology Transfer and asked them, ‘Do you like this?’ They said, ‘Yeah, we like that.’ ”

Steven Laken is the C.E.O. of Cephos, a Boston-based company that is developing a lie-detection product based on Kozel’s watch-and-ring study. (It has an exclusive licensing agreement for pending patents that the Medical University of South Carolina applied for in 2002.) Cephos is proceeding more cautiously than No Lie. Laken’s company is still conducting studies with Kozel, the latest of which involve more than a hundred people. (The sample pool is again young, healthy, and free of criminal records and psychological problems.) Cephos won’t be offering fMRIs commercially until the results of those studies are in; Laken predicts that this will happen within a year. At the National Academy of Sciences committee meeting, he said, “I can say we’re not at ninety-per-cent accuracy. And I have said, if we were not going to get to ninety per cent, we’re not going to sell this product.” (Nobody involved in fMRI lie detection seems troubled by a ten-per-cent error rate.)

In March, I went to a suburb of Boston to meet Laken. He is thirty-five years old and has a Ph.D. in cellular and molecular medicine from Johns Hopkins. Nine years ago, he identified a genetic mutation that can lead to colorectal cancer. He has a more conservative temperament than Joel Huizenga does, and he told me he thinks that spousal-fidelity cases are “sleazy.” But he sees a huge potential market for what he calls a “truth verifier”—a service for people looking to exonerate themselves. “There are some thirty-five million criminal and civil cases filed in the U.S. every year,” Laken said. “About twenty million are criminal cases. So let’s just say that you never even do a criminal case—well, that still leaves roughly fifteen million for us to go after. Some you exclude, but you end up with several million cases that are high stakes: two people arguing about things that are important.” Laken also thinks that fMRI lie detection could help the government elicit information, and confessions, from terrorist suspects, without physical coercion.

He calmly dismissed the suggestion that the application of fMRI lie detection is premature. “I’ve heard it said, ‘This technology can’t work because it hasn’t been tested on psychopaths, and it hasn’t been tested on children, and it certainly hasn’t been tested on psychopathic children,’ ” he said. “If that were the standard, there’d never be any medicine.”

Laken and I spoke while driving to Framingham, Massachusetts, to visit an MRI testing center run by Shields, a company that operates twenty-two such facilities in the state. Laken was working on a deal with Shields to use their scanners. For Shields, it would be a smart move, Laken said, because customers would pay up front for the scan—there would be no insurance companies to contend with. (Cephos and Shields have since made an official arrangement.) Laken believes that Cephos will prosper primarily through referrals: lawyers will function as middlemen, ordering an fMRI for a client, much as a doctor orders an MRI for a patient.

We pulled into the parking lot, where a sign identifies Shields as “the MRI provider for the 3-X World Champion

New England Patriots.” Inside, John Cannillo, an imaging specialist at Shields, led us into a room to observe a woman undergoing an MRI exam. She lay on a platform that slid into a white tubular scanner, which hummed like a giant tuning fork.

During a brain scan, the patient wears a copper head coil, in order to enhance the magnetic field around the skull. The magnet is so powerful that you have to remove any metal objects, or you will feel a tugging sensation. If a person has metal in his body—for instance, shrapnel, or the gold grillwork that some hip-hop fans have bonded to their teeth—it can pose a danger or invalidate the results. At the N.A.S. meeting in Washington, one scientist wryly commented, “It could become a whole new industry—criminals having implants put in to avoid scanning.”

A Shields technician instructed the woman in the scanner from the other side of a glass divide. “Take a breath in, and hold it, hold it,” he said. Such exercises help minimize a patient’s movements.

As we watched, Laken admitted that “the kinks” haven’t been worked out of fMRI lie detection. “We make mistakes,” he said of his company. “We don’t know why we make mistakes. We may *never* know why. We hope we can get better.” Some bioethicists and journalists may worry about the far-off threat to “cognitive freedom,” but the real threat is simpler and more immediate: the commercial introduction of high-tech “truth verifiers” that may work no better than polygraphs but seem more impressive and scientific. Polygraphs, after all, are not administered by licensed medical professionals.

Nancy Kanwisher, a cognitive scientist at M.I.T., relies a great deal on MRI technology. In 1997, she identified an area near the bottom of the brain that is specifically involved in perceiving faces. She has become a pointed critic of the rush to commercialize brain imaging for lie detection, and believes that it’s an exaggeration even to say that research into the subject is “preliminary.” The tests that have been done, she argues, don’t really look at lying. “Making a false response when instructed to do so is not a *lie*,” she says. The ninety-per-cent “accuracy” ascribed to fMRI lie detection refers to a scenario so artificial that it is nearly meaningless. To know whether the technology works, she believes, “you’d have to test it on people whose guilt or innocence hasn’t yet been determined, who believe the scan will reveal their guilt or innocence, and whose guilt or innocence can be established by other means afterward.” In other words, you’d have to run a legal version of a clinical trial, using real suspects instead of volunteers.

Langleben believes that Kanwisher is too pessimistic. He suggested that researchers could recruit people who had been convicted of a crime in the past and get them to lie retrospectively about it. Or maybe test subjects could steal a “bagel or something” from a convenience store (the researchers could work out an agreement with the store in advance) and then lie about it. But even these studies don’t approximate the real-world scenarios Kanwisher is talking about.

She points out that the various brain regions that appear to be significantly active during lying are “famous for being activated in a wide range of different conditions—for almost any cognitive task that is more difficult than an easier task.” She therefore believes that fMRI lie detection would be vulnerable to countermeasures—performing arithmetic in your head, reciting poetry—that involve concerted cognitive effort. Moreover, the regions that allegedly make up the brain’s “lying module” aren’t that small. Even Laken admitted as much. As he put it, “Saying ‘You have activation in the anterior cingulate’ is like saying ‘You have activation in Massachusetts.’ ”

Kanwisher’s complaint suggests that fMRI technology, when used cavalierly, harks back to two pseudosciences of the eighteenth and nineteenth centuries: physiognomy and phrenology. Physiognomy held that a person’s character was manifest in his facial features; phrenology held that truth lay in the bumps on one’s skull. In 1807, Hegel published a critique of physiognomy and phrenology in “The Phenomenology of Spirit.” In that work, as the philosopher Alasdair MacIntyre writes, Hegel observes that “the rules that we use in everyday life in interpreting facial expression are highly fallible.” (A friend who frowns throughout your piano recital might explain that he was actually fuming over an argument with his wife.) Much of what Hegel had to say about physiognomy applies to modern attempts at mind reading. Hegel quotes the scientist Georg Christoph Lichtenberg, who, in characterizing physiognomy, remarked, “If anyone said, ‘You act, certainly, like an honest man, but I can see from your face you are forcing yourself to do so, and are a rogue at heart,’ without a doubt every brave fellow to the end of time when accosted in that fashion will retort with a box on the ear.” This response is correct, Hegel argues, because it “refutes the fundamental assumption of such a ‘science’ of conjecture—that the reality of a man is his face, etc. The true being of man is, on the contrary, his act; individuality is real in the deed.” In a similar vein, one might question the core presumption of fMRI—that the reality of man is his brain.

Elizabeth Phelps, a prominent cognitive neuroscientist at N.Y.U., who studies emotion and the brain, questions another basic assumption behind all lie-detection schemes—that telling a falsehood creates conflict within the liar. With the polygraph, the assumption is that the conflict is emotional: the liar feels guilty or anxious, and these feelings produce a measurable physiological response. With brain imaging, the assumption is that the conflict is cognitive: the liar has to work a little harder to make up a story, or even to stop himself from telling the truth. Neither is necessarily right. “Sociopaths don’t feel the same conflict when they lie,” Phelps says. “The regions of the brain that might be involved if you have to inhibit a response may not be the same when you’re a sociopath, or autistic, or maybe just strange. Whether it’s an emotional or a cognitive conflict you’re supposed to be exhibiting, there’s no reason to assume that your response wouldn’t vary depending on what your personal tendencies are—on who *you* are.”

When I talked to Huizenga, the No Lie C.E.O., a few months after I had met him in California, he was unperturbed about the skepticism that he was encountering from psychologists. “In science, when you go out a little further than other people, it can be hard,” he said. “The top people understand, but the middle layer don’t know what you’re talking about.”

Huizenga told me that he was trying to get fMRI evidence admitted into a California court for a capital case that he was working on. (He would not go into the case’s details.) Given courts’ skepticism toward the polygraph, Huizenga’s success is far from certain. Then again we are in a technology-besotted age that rivals the twenties, when Marston popularized lie detection. And we live in a time when there is an understandable hunger for effective ways to expose evildoers, and when concerns about privacy have been nudged aside by our desire for security and certainty. “Brain scans indicate”: what a powerful phrase. One can easily imagine judges being impressed by these pixellated images, which appear so often in scientific journals and in the newspaper. Indeed, if fMRI lie detection is successfully marketed as a service that lawyers steer their clients to, then a refusal even to take such a test could one day be cause for suspicion.

Steven Hyman, the Harvard psychiatrist, is surprised that companies like No Lie have eluded government oversight. “Think of a medical test,” he said. “Before it would be approved for wide use, it would have to be shown to have acceptable accuracy among the populations in whom it would be deployed. The published data on the use of fMRI for lie detection uses highly artificial tests, which are not even convincing models of lying, in very structured laboratory settings. There are no convincing data that they could be used accurately to screen a person in the real world.” But, in the end, that might not matter. “Pseudo-colored pictures of a person’s brain lighting up are undoubtedly more persuasive than a pattern of squiggles produced by a polygraph,” he said. “That could be a big problem if the goal is to get to the truth.”

Laken, meanwhile, thinks that people who find themselves in a jam, and who are desperate to exonerate themselves, simply have to educate themselves as consumers. “People have said that fMRI tests are unethical and immoral,” he said. “And the question is, Why is it unethical and immoral if somebody wants to spend their money on a test, as long as they understand what it is they’re getting into? We’ve never said the test was *perfect*. We’ve never said we can guarantee that this is admissible in court and that’s it—you’re scot-free.” Later that day, I looked again at the Cephos Web site. It contained a bolder proclamation. “The objective measure of truth and deception that Cephos offers,” it said, “will help protect the innocent and convict the guilty.” ♦