Journal of Scientific Exploration, Vol. 32, No. 2, pp. 433–445, 2018 0892-3310/18

BOOK REVIEW

Into the Grey Zone: A Neuroscientist Explores the Border Between Life and Death by Adrian Owen. Guardian Books, 2017. ix + 304 pp. £16.99 (hardcover). ISBN 978-1-78335-098-8.

DOI: https://doi.org/10.31275/2018.1297

Dramatic modern advances in emergency and resuscitation medicine, starting perhaps with the development of effective mechanical ventilators in the mid-20th century, have created a large class of persons who in earlier times would almost certainly have died, but who can now go on existing, suspended at least temporarily in a state somewhere between death and the conscious life they formerly pursued. A very wide range of brain injuries lead first to *coma*, in which the patient shows no sign of conscious *awareness*, or even of *wakefulness*, in which eye openings and closings indicate the presence of a sleep/wake cycle unconsciously mediated by structures in the brain stem. Following emergence from coma, which may take days to months or more if it happens at all, patients typically show signs of wakefulness without conscious awareness—they are then in a so-called *vegetative state* (VS). For many this state becomes permanent, but some go on to a more recently described condition called the *minimally conscious state* (MCS), in which signs of conscious awareness can be detected by careful neurological examination (see Laureys & Tononi 2009: Chapter 14). A very few such persons ultimately progress to more or less full recovery, but another and particularly horrifying possible outcome is the "locked-in" state, in which a patient is fully conscious but has extremely little or no capacity for voluntary motor action. A famous modern example is that of Jean-Dominique Bauby, author of *The Diving Bell and the Butterfly*, who suffered a stroke to his upper brainstem and awakened 20 days later to find himself fully conscious but capable only of blinking his left eye. A tiny fraction of surgical patients find themselves in similarly terrifying conditions caused by the combination of muscle relaxants with insufficient levels of anesthetic agents (Kelly et al. 2007:387 n. 18).

The distinction between VS and MCS is often difficult to make, clinically, and there are high rates of error in both directions. Skilled neurologists are good at making it, given sufficient information, but they see patients intermittently at best and not for long. Grieving family members and loved ones, on the other hand, may spend far more time with the patient

and thus have much greater opportunity to observe relevant evidence, but as in the infamously sensationalized case of Terri Schiavo's parents they are also far more vulnerable to over-interpreting what they observe as evidence of awareness and potential recovery.

Against this backdrop, British neuropsychologist Adrian Owen describes in this engaging new book the 20-year trajectory of his increasingly successful use of modern functional neuroimaging methods including PET (positron emission tomography), fMRI (functional magnetic resonance imaging), and most recently EEG (electroencephalography) to detect signs of conscious awareness in patients who had been diagnosed as vegetative using standard neurological criteria. He tells this story primarily in terms of gripping accounts of 10 of his own cases, interspersed with related autobiographical material of an often very poignant sort, and with occasional reflections on the weighty ethical and legal ramifications of his research. Let me next sketch the highlights of this remarkable journey.

The story begins with "Kate," in whom encephalomyelitis had produced widespread white-matter destruction. In a hastily organized PET study, Owen and colleagues carried out 12 scans in which they displayed images of Kate's family members and friends on a monitor inside the scanner, and discovered to their surprise that her face-recognition area (fusiform gyrus) responded selectively to those images in comparison with responses to unfocused versions of the same pictures. These preliminary findings were published as a letter in *Lancet* (Menon et al. 1998) and resulted in a substantial media uproar. Owen felt that he had made contact with Kate, but critics argued that the observed activations might have been automatic brain responses, and were insufficient to demonstrate the presence of conscious awareness. She later slowly recovered to a substantial extent, and reported having experienced intense feelings of pain, terror, anger, and thirst during her supposedly vegetative period (pp. 36, 39). In a still later conversation, Owen deliberately declined to ask her whether she remembered her scanning session (p. 41).

Next comes "Debbie," victim of a head-on car collision resulting in severe anoxia, with damage to the upper brainstem and absence of pupillary reflexes. PET scans were again carried out, but this time the experimental design contrasted responses to meaningful words (two-syllable nouns) versus carefully matched bursts of noise—stimuli crafted by colleagues proficient in psycholinguistics The switch to auditory stimuli was motivated by the fact that Kate's eyes had been closed during 3 of her 12 scans. The main result was that areas of the brain known to be involved in speech processing were selectively activated in Debbie's brain by the words, but not by the noise controls. Possible indications of residual awareness had thus again

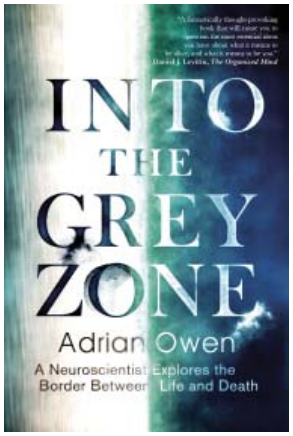

been found, although it was again not possible to conclude with certainty that awareness was really present (Owen et al. 2002).

"Kevin," the victim of a massive stroke involving the brainstem and thalamus, was studied next, again using PET in conjunction with auditory stimuli provided by Owen's psycholinguistics colleagues. The stimuli this time consisted of whole sentences, each of which was varied systematically in intelligibility by addition of differing amounts of background noise. These stimuli had been shown in normal volunteers to evoke activation in specific speech-processing areas, with the level of activation proportional to the difficulty of comprehension. A patient showing similar effects, Owen reasoned, might plausibly be thought to be comprehending the sentences, and hence to some degree be conscious. The original PET study was repeated with Kevin nine months later, together with a related preliminary study featuring two notable advances: First, the scans were carried out this time using newly available fMRI technology, which has much better spatial and temporal resolution and is free of the limitation on PET scanning imposed by considerations of radiation burden. Second, the stimuli were now sentences carefully calibrated in terms of the amount of lexical ambiguity they contained; these had been shown in normal volunteers to produce a more complicated pattern of brain response, including areas of the frontal lobe that became active in proportion to the amount of ambiguity needing to be resolved. All three studies yielded similar results, with Kevin's brain consistently producing responses resembling those of the normal volunteers (Owen et al. 2005). Owen now felt confident he had obtained convincing evidence of awareness in a VS patient and presented these results "excitedly" to colleagues, but to his chagrin they remained unpersuaded (pp. 85–86). It is perhaps also worth noting here that Kevin has remained VS.

Following this disappointing reception of his findings, during a 4-month sabbatical in Australia, Owen realized that what he needed to do was to show that a patient could *voluntarily* perform mental tasks that had predictable and divergent neurophysiological consequences. This would in principle allow persons incapable of voluntary motor acts of the sorts relied upon by neurologists, to provide analogous evidence, through brain activity alone, of their continued conscious presence. He returned to Cambridge, and following careful pre-testing of a variety of possible tasks settled on

two that reliably produced radically divergent patterns of brain activity in healthy volunteers. The first—imagining oneself playing tennis—robustly activated premotor cortex, while the second—imagining oneself walking through the rooms of one's own dwelling—similarly activated a spatially distant area called the parahippocampal gyrus. When these tasks were first presented to VS patient "Carol" in an fMRI scanner, her brain activated in strikingly similar ways. These and other results from Carol were published in a one-page article in *Science* (Owen et al. 2006) which attracted enormous media attention and convinced most readers, myself included, that Carol probably *was* to some degree conscious. A few very determined skeptics suggested that perhaps just hearing the instructions was sufficient to elicit automatically the corresponding patterns of activity, but that possibility was ruled out by further experiments in which healthy volunteers listened to the instructions but deliberately avoided carrying out the tasks, and failed to produce the normal activity patterns.

Over the next few years this 2-task procedure was carried out with 54 additional patients; 23 of these had been diagnosed VS, and 4 succeeded in the tasks (Monti et al. 2010).

Readers will likely anticipate what comes next: If one can reliably produce distinct and easily recognizable patterns of brain activity by performing two different mental tasks, then perhaps one can use those tasks to answer yes/no questions in appropriate ways. After verifying this in himself and 16 other volunteers, Owen tried it first in a long-distance fMRI scanning session with an East European VS patient named "John" who had suffered traumatic brain injury in a motorcycle accident five years previously and was being studied at the facilities of Steven Laureys and his team at Liége, in Belgium. During that single brief session John correctly answered 5 yes/no questions about himself and his family, and Owen began to think about the implications: From a clinical point of view it would be wonderful if we could ask such patients whether they are in pain, and if they respond "yes" give them pain medication and verify that it works. But what if we ask whether they want to live and they say "no"—what then? Increased capacity to communicate with severely brain-damaged persons clearly carries with it momentous potentials for both good and ill, and we will ultimately have to come to terms with this.

By this time Owen's work had attracted an enormous amount of mostly positive attention, and he moved to a lavishly funded new position at the University of Western Ontario in London, Ontario. There he employed the 2-task procedure with a patient named "Scott," who had been nominally VS for 12 years after being T-boned at the wheel of his car by a policeman traveling at high speed on the way to a crime scene. Scott

began, dramatically enough, by answering "no" when asked whether he was in pain during a session filmed with Owen's approval by the BBC. He went on during the months that followed to answer numerous other yes/no questions, demonstrating that he knew who and where he was, remembered his personal history prior to the accident, knew the names of his caregivers, and remembered other persons and events from the period of his VS. This was communication on an unprecedented scale with a person repeatedly diagnosed as VS by a highly competent neurologist, and it validated the confidence that Scott's parents had expressed all along that their son was still there. Sadly, Scott died about a year later.

The next patient, Abraham, had suffered an intraventricular hemorrhage due to an aneurysm in his anterior communicating artery. Owen uses this case primarily as a vehicle for discussion of the gut-wrenching life-and-death issues involved in such cases, touching also upon various others including those of Terri Schiavo, Karen Quinlan, and Nancy Cruzan that had attracted large amounts of media attention. No scanning was performed, and Abraham died in the hospital.

Having by now discovered that highly specific probes such as the tennis/house tasks sometimes failed in patients who could be shown in other ways to have some degree of residual awareness, Owen and his expanding research group next began to explore "naturalistic" paradigms that could potentially canvas brain function in a more general and searching fashion and hence detect awareness more reliably. A high point of this effort is their work with a short but intense Alfred Hitchcock film called *Bang, You're Dead*, about a small child who discovers a revolver and some bullets and begins to play with his discoveries. The basic idea is that the film places a variety of strong demands on the viewer's attentive, executive, and affective capacities, demands which result in a surprisingly consistent dynamic pattern of fMRI responses in normal volunteers. Owen describes vividly the initial use of this paradigm with a patient named Jeff Tremblay, who had been nominally VS for 15 years following cardiac arrest produced by a kick to his chest, but who nonetheless displayed fMRI response patterns strikingly similar to those of normal volunteers while watching the movie—including activation of frontal and parietal brain areas thought to be required for understanding its dynamically unfolding plot (Naci et al. 2014, Naci, Sinai, & Owen 2017). As Owen himself somewhat oddly puts it,

> We had shown for the first time that the brain activity produced by similar conscious experiences in different individuals could be used to infer conscious awareness in physically nonresponsive patients without any need for self-report. (p. 204)

The next case, Juan, is the most interesting of all. Upon arrival at the hospital following a choking incident, he had a score of 3 out of 15 on the Glasgow Coma Scale—the worst possible score—and a CT scan showed widespread diffuse white matter damage apparently caused by the resulting anoxia. He was declared vegetative after remaining totally unresponsive for two months, but his parents subsequently transported him from the hospital to visit Owen's group for 4 days in hopes that their neuroimaging procedures might suggest some possibility of recovery. Neither the tennis/house task nor the Hitchcock movie, however, revealed any clear evidence of residual awareness, even after repeated testing. At this point Juan's future looked bleak at best. A routine followup call to his parents seven months later, however, revealed that against all expectations he was well on the way to recovery—eating, walking, and talking. Subsequently, through further interactions with Owen and his team, Juan demonstrated that he remembered not only people and events from the neuroimaging sessions but even earlier events that had occurred in the hospital shortly after his admission. Owen acknowledges never having seen "*anything* remotely like Juan's recovery" (p. 209, italics his) and being mystified by many aspects of the case (p. 215), but he never comes directly and fully to grips with its real challenge. I will return to this shortly, after completing my summary of the book's main contents.

By this point Owen had recognized that his fMRI-based procedures, despite their demonstrated virtues, were not going to be the full answer to what was now emerging as a major medical need for fast, cheap, and portable bedside testing procedures. Good MR scanners are still extremely expensive, after all, and patients have to be conveyed to the machines. To this end he began to investigate the possible use of EEG procedures, in which he had shown surprisingly little interest up to this point in the book. His initial foray in this direction was reported by Cruse et al. (2011), who found 3 of 16 VS patients able to produce fairly distinctive EEG patterns when asked to imagine performing divergent motor tasks (squeezing all fingers of the right hand into a fist *vs*. wiggling the toes of both feet). Only ¾ of the healthy controls succeeded with this protocol, however, suggesting that there remained plenty of room for improvement. The state of this ongoing development as of mid-2015, including its embodiment in a dedicated "EEJeep", is described by Owen in conjunction with the case of patient #10, Leonard, who in 2010 had suffered cardiac arrest while asleep. His wife had detected the problem and called for EMTs, but it was on the order of 15 minutes before they arrived and were able to restart his heart. Subsequently diagnosed as VS, Leonard had shown no signs of awareness in prior fMRI scanning by Owen's team with the tennis/house procedure,

but his wife invited retesting with the new EEG routines, which at this point relied upon possible differences in the neuroelectric responses evoked by speech sounds *vs*. meaningless noise, or by pairs of words that were related versus unrelated in meaning. Once again, no signs of awareness were detected.

In his final chapter—"Reading Minds"—Owen turns to the future of grey-zone science. EEG paradigms now represent for him the growing edge of this development, and he has recognized the deep affinity between his own ongoing work and a burgeoning area of research on "brain–computer interfaces" (BCI), much of it devoted to development of prostheses for fully conscious war veterans who have lost motor organs. He imagines a bright future in which sophisticated artificial intelligence techniques combined with enhanced EEG recording capabilities (possibly including arrays of electrodes implanted in the brain) will enable millions of brain-damaged and disabled persons to communicate effectively with their caregivers and loved ones, and to take back control of their own lives. He even imagines a scenario in which an individual who had been brain-damaged in a criminal attack provides investigators with information enabling them to identify and capture the attacker. All of this, he repeatedly suggests, flows from recognition and acceptance of the *fact* that we are nothing more than our brains (see for example pp. 27, 68, 72, 225, 255).

The book ends with a brief but poignant Epilogue, Acknowledgments, extensive chapter-by-chapter Notes including references to key papers, and an Index.

Turning now to evaluation of the book, the first thing that must be said is that Owen and his colleagues deserve great credit for systematically and doggedly opening up this new window into states of impaired consciousness in brain-damaged persons. I have no doubt that some nominally VS patients are in fact consciously aware to some degree, and the ability to detect that fact using neuroimaging methods certainly carries enormous implications for improved diagnosis and treatment going forward. The techniques they have introduced will continue to improve, and I think it essentially certain that standardized diagnostic procedures will soon begin to incorporate them.

That said, however, I must next focus on several aspects of his presentation that I find less than satisfactory. I'll begin where he ends, with the prospects for "mind-reading" by machines. This is potentially a big subject in itself, but in brief I think Owen is overly impressed with glitzy technologies in general, and with this one in particular. A particularly startling specimen occurs on p. 252, where he exclaims "Emerging technologies will undoubtedly one day allow us to *read* the minds of others. Not in the rudimentary sense that we do already—decoding yes and no responses

based on changes of fMRI activity—but in the sense of interpreting and understanding exactly what another person is thinking based solely on some sort of readout from his or her brain" (italics his).

As in the AI world generally, enthusiastic researchers have been quick to suggest that the rather modest successes achieved so far represent first steps leading inevitably to eventual complete triumph. In this case that is anything but certain, however, and in fact there are many reasons for doubting that anything on this scale will ever be possible. For starters, all existing neuroimaging technologies have significant limitations in terms of spatial and temporal resolution and the manner of coupling between the measured signals and the underlying brain activity, and it is not self-evident that we could ever measure patterns of brain activity at the requisite level of detail even if the postulated highly specific brain–mind correlations existed. There must also be vast numbers if not an infinite number of mental states, and specific phenomenal contents with potentially recognizable physiological correlates often occur together, or as components of some larger organization. These are already massive practical obstacles, but they pale in comparison with still deeper issues of a more abstract and philosophical sort: In particular, anyone who imagines that the correspondence between mental states and brain states is clear and consistent enough to support full-scale mind-reading by machines should also consult the weighty arguments against this possibility by philosopher Stephen Braude (2002:123–140 and Part II). For example, brain-states corresponding to a mental image of an old rabbi are sure to differ endlessly in detail, like the images themselves, both within and across individuals, and even explicit knowledge that such an image is present as a component of some mental state would ultimately reveal very little about the total intentional content of that state in the specific context in which it occurs. On the physical side, meanwhile, despite their gross overall similarity in appearance, human brains vary widely in the details of their structural and functional organization, implying that specific mental-state "types"—even if such things existed, which appears highly doubtful—would likely take on widely varying appearances across subjects in terms of the accompanying brain activity. Much of the work to date on "mind-reading" by machines has in fact focused on artificially simple tasks such as discrimination of brain responses to small numbers of distinct sensory stimuli, usually in terms of activation patterns evoked by those stimuli in the corresponding primary sensory areas of the brain. Even in that restricted setting, brain damage could well disturb whatever correlations antecedently existed, and things will certainly get much tougher as we move toward more central parts of the mind.

I do not mean to discount the value of BCI research, because for anyone who has been locked-in and unable to do anything at all, the ability to answer a simple yes/no question, even if it takes 5 minutes, already represents essentially infinite progress. But although some practically useful advances in this direction have already occurred, and more are sure to follow, talk like Owen's as quoted above seems to me nothing short of unbridled science fiction—fiction, moreover, which can be pernicious to the extent it engenders unrealistic hopes in persons struggling to care for their brain-damaged or locked-in loved ones.

My second main concern is that Owen overstates considerably what he has definitely discovered about the mental lives of his patients. He seems to see little middle ground between being truly vegetative and being fully intact and conscious. On p. 3 of the Prologue he announces dramatically that ". . . we have discovered that 15 to 20 per cent of people in the vegetative state who are assumed to have no more awareness than a head of broccoli are fully conscious, although they never respond to any form of external stimulation." For all intents and purposes, that is, he is declaring such persons to be locked-in, although he hesitates to apply that term (p. 123). But the ensuing narrative repeatedly exaggerates what he actually knows about his patients' mental condition: Thus for example, "Carol was hopelessly disadvantaged by her useless body but was nevertheless still in there—her personality, attitudes, beliefs, moral compass, memories, hopes and fears, dreams and emotions" (p. 113). Similarly, Owen declares of Scott, the yes/no test subject, that "On that day, and on many occasions in the months that followed, we *conversed* with Scott in the scanner" (p. 162, italics mine), even though he himself acknowledges a short time later that "We had never had a real conversation" (p. 167). Jeff Tremblay, shown the Hitchcock film, "was conscious and experiencing the movie just as you or I would" (p. 201). Again, "Juan remembered everything about his first visit, down to the tiniest detail" (p. 214). Additional examples can be found scattered through the text.

Owen's narrative of mental intactness is reinforced, moreover, by parallel exaggerations about similarities in patterns of brain activation between patients and normals: Thus for example Debbie's brain "responded to speech and noise bursts just like yours or mine" (p. 57), and Kevin's temporal lobe "lit up in exactly the same way it had in the healthy volunteers" (p. 85). When asked to imagine playing tennis, Carol "would activate her premotor cortex just like healthy volunteers" (p. 109), and when asked to imagine walking through her house, "her pattern of brain activity was identical to that of healthy volunteers" (p. 109). Owen acknowledges that when excited he is "famously prone to hyperbole" (p. 209), and there is

certainly plenty of excitement here, but this tendency goes beyond what seems proper even in a book intended for a popular audience. The 2006 *Science* paper on Carol, moreover, states that "Her neural responses were *indistinguishable* from those observed in healthy volunteers performing the same tasks" (italics mine), when in fact they were only *statistically*, not visually, indistinguishable, and occurred in brains that were probably in very different overall functional conditions; vegetative patients for example typically have resting rates of cerebral blood flow and metabolism that are far below normal levels, approaching those observed in deep general anesthesia (Laureys & Tononi 2009: Chapter 13).

What Owen was really talking about in his Prologue, of course, was people who have been *diagnosed* as vegetative, and question number one is whether those diagnoses were correct. I think the answer is often "no," and this brings me to a more fundamental issue. For someone who repeatedly describes himself as being interested in basic science, Owen makes distressingly little effort to connect his work with the massive contemporary literature on the neuroscience of consciousness. Let me explain: Over the last several decades, an overwhelming consensus has developed among neuroscientists that normal human conscious experience occurs only in conjunction with a brain that is capable of generating large-scale oscillatory neuroelectric activity cooperatively linking widespread cortical and subcortical territories (there remains plenty of disagreement as to exactly how the two are related, but that has no bearing on the following discussion). Owen's colleague Steven Laureys has done as much as anybody to situate grey-zone science within that framework, which entails as a general expectation that brain-damaged persons can only recover consciousness to the extent that their brains begin to function as they did before the injury, by virtue of re-establishment of metabolic activity and connectivity within large-scale thalamocortical networks (e.g., Laureys et al. 2005, Laureys & Tononi 2009, Laureys & Schiff 2012). Most of the literature on disorders of consciousness appears consistent with this picture: For example, strong somatosensory (wrist-shock) stimuli that are experienced as noxious and painful by healthy volunteers activate primary somatosensory cortex in vegetative patients in the normal way (because the sensory pathway itself remains intact), but that first-stage activation fails to spread as in normal brains to other regions of the cortical pain network (Laureys et al. 2002). Similarly, it has long been known that severely brain-damaged persons can sometimes exhibit behavioral signs associated with cortical "islands" of relatively preserved metabolism and neuroelectric activity that remain functionally isolated from the rest of the brain (Schiff et al. 1999). By contrast, genuinely locked-in but fully conscious patients

such as Jean-Dominique Bauby typically have normal cerebral blood flow and metabolism and normal (or only mildly abnormal) EEGs (Laureys & Tononi 2009: Chapter 15, Patterson & Grabois 1986).

Owen's own patients illustrate the fact that brain injuries themselves can take an indefinite variety of forms, and in general it is extremely difficult to know with much precision what happened in any given case. That problem is compounded, moreover, by the typically widespread and hard-to-track dynamic changes that follow over time in response to the original injury. As Owen himself acknowledges, "every brain is different, and every brain *injury* is different" (p. 225, italics his). So were his patients really all vegetative? For several of his most important cases I think there are good reasons to doubt it. Kate, for example, could apparently fixate on the displayed images, suggesting MCS. She also produced long-latency evoked-potential components, and of course she eventually recovered to a considerable extent. The Supplementary material for the paper on Carol shows that she displayed background EEG activity including alpha and beta frequencies. Jeff Tremblay appeared capable of visual tracking, and Owen himself suggests on that basis that he was probably MCS rather than VS at the time of his scanning.

I have no doubt, in sum, that neuroimaging methods can detect signs of residual awareness in some severely brain-damaged patients who are behaviorally unresponsive, and this surely represents an important advance in diagnosis and treatment. From a theoretical standpoint, however, most of Owen's patients do not seem to me to pose very clear-cut challenges to the conventional picture sketched above, and I see little justification for his central claim that large numbers of genuinely vegetative patients are in effect locked-in, totally intact and lucid conscious minds trapped inside shattered brains.

The main possible exception here, of course, is Juan, and this brings me to my final and most important concern. Owen never seems quite to appreciate or at least articulate the fact that if his central claim were true it would effectively falsify the contemporary mainstream neuroscientific consensus on consciousness and the brain as described above. Among all the patients described in this book, Juan certainly comes closest to reaching that threshold, and we have already seen how disturbing this was to Owen himself. But even in Juan's case, in my opinion, our knowledge concerning the functional condition of his brain at the time of the events he later recalled is not sufficient—at least not yet—to make that argument stick. But there are numerous other cases already in the literature that make essentially the same argument in a much stronger fashion, and it galls me that Owen makes no contact whatsoever with this large body of material. I

am referring, of course, to the hundreds of cases of near-death experiences (NDEs) occurring under extreme physiological conditions such as adequate deep general anesthesia and/or cardiac arrest. In circumstances such as these, the specific neurophysiological conditions thought by virtually all contemporary neuroscientists to be necessary for conscious experience are definitely abolished, and yet many persons subjected to those circumstances later report having had not just any old conscious experiences but the most extraordinary and transformative experiences of their entire lives. In many such cases, moreover, the experiences can be anchored to the period of brain impairment by the patients' ability to correctly report events occurring during that time (e.g., Holden, Greyson, & James 2009, Kelly et al. 2007: Chapter 6, van Lommel et al. 2001).

The inside front portion of the jacket material of Owen's book intriguingly poses the following question: "We have known for a long time that a body does not define a person—but what if a brain does not define a mind? What does it mean if a mind can exist unharmed within a deeply damaged brain?" That is indeed the theoretically most fundamental question, but Owen himself unfortunately fails in the end to come fully and directly to grips with it. The work described in this book is surely important, but its significance up to this point is almost entirely clinical, and not theoretical. One can only hope that at least some of the many able persons now working on grey-zone science will take the central message of this review to heart, and devote serious attention to cases of this most profoundly challenging sort.

Edward F. Kelly
Division of Perceptual Studies
Department of Psychiatry and Neurobehavioral Sciences
University of Virginia, Charlottesville, VA
ek8b@virginia.edu

References Cited

Braude, S. (2002). *ESP and Psychokinesis: A Philosophical Examination* (revised edition). Parkland, FL: Brown Walker Press.

Cruse, D., Chennu, S., Chatelle, C., Bekinschtein, T. A., Fernández-Espejo, D., Pickard, J. D., Laureys, S., & Owen, A. M. (2011). Bedside detection of awareness in the vegetative state: A cohort study. *Lancet*, *378*:2088–2094.

Holden, J., Greyson, B., & James, D. (2009). *The Handbook of Near-Death Experiences: Thirty Years of Investigation.* Santa Barbara, CA: ABC-CLIO.

Kelly, E. F., Kelly, E. W., Crabtree, A., Gauld, A., Greyson, C. B., & Grosso, M. (2007). *Irreducible Mind: Toward a Psychology for the 21st Century.* Lanham, MD: Rowman & Littlefield.

Laureys, S., & Schiff, N. D. (2012). Coma and consciousness: Paradigms (re)framed by neuroimaging. *NeuroImage, 61*:478–491.

Laureys, S., & Tononi, G. (Editors) (2009). *The Neurology of Consciousness: Cognitive Neuroscience and Neuropathology*. New York: Academic Press.

Laureys, S., Faymonville, M. E., Peigneux, P., Damas, P., Lambermont, B., Del Fiore, G., Degueldre, C., Aerts, J., Luxen, A., Franck, G., Lamy, M., Moonen, G., & Maquet, P. (2002). Cortical processing of noxious somatosensory stimuli in the persistent vegetative state. *NeuroImage*, *17*:732–741.

Laureys, S., Perrin, F., Schnakers, C., Boly, M., & Majerus, S. (2005). Residual cognitive function in comatose, vegetative and minimally conscious states. *Current Opinion in Neurology*, *18*:726–733.

Menon, D. K., Owen, A. M., Williams, E. J., Minhas, P. S., Allen, C. M. C., Boniface, S. J., Pickard, J. D., & the Wolfson Brain Imaging Centre Team (1998). Cortical processing in persistent vegetative state. *Lancet*, *352*:200.

Monti, M. M., Vanhaudenhuyse, A., Coleman, M. R., Boly, M., Pickard, J. D., Tshibanda, L., Owen, A. M., & Laureys, S. (2010). Willful modulation of brain activity in disorders of consciousness. *The New England Journal of Medicine*, *362*:579–589.

Naci, L., Sinai, L., & Owen, A. M. (2017). Detecting and interpreting conscious experiences in behaviorally non-responsive patients. *NeuroImage*, *145*:304–313.

Naci, L., Cusack, R., Anello, M., & Owen, A. M. (2014). A common neural code for similar conscious experiences in different individuals. *Proceedings of the National Academy of Sciences*, *111*:14277–14282.

Owen, A. M., Menon, D. K., Johnsrude, I. S., Bor, D., Scott, S. K., Manly, T., Williams, E. J., Mummery, C., & Pickard, J. D. (2002). Detecting residual cognitive function in persistent vegetative state. *Neurocase*, *8*:394–403.

Owen, A. M., Coleman, M. R., Menon, D. K., Johnsrude, I. S., Rodd, J. M., Davis, M. H., Taylor, K., & Pickard, J. D. (2005). Residual auditory function in persistent vegetative state: A combined PET and fMRI study. *Neuropsychological Rehabilitation*, *15*:290–306.

Owen, A. M., Coleman, M. R., Boly, M., Davis, M. H., Laureys, S., & Pickard, J. D. (2006). Detecting awareness in the vegetative state. *Science*, *313*:1402.

Patterson, J. R., & Grabois, M. (1986). Locked-in syndrome: A review of 139 cases. *Stroke*, *17*:758–764.

Schiff, N., Ribary, U., Plum, F., & Llinás, R. (1999). Words without mind. *Journal of Cognitive Neuroscience*, *11*: 650–656.

Van Lommel, P., van Wees, R., Meyers, V., & Elfferich, I. (2001). Near-death experiences in survivors of cardiac arrest: A prospective study in the Netherlands. *Lancet*, *358*:2039–2045.