

We Robot 2015

PERSONAL RESPONSIBILITY IN THE AGE OF USER-CONTROLLED NEUROPROSTHETICS

**Patrick Moore, Timothy Brown, Jeffrey Herron,
Margaret Thompson, Tamara Bonaci, Sara Goering,
and Howard Chizeck**

Possibilities for improving lives through computerized augmentation of the human brain abound. In the next few decades we will likely see advanced prostheses which restore natural mobility to amputees, devices that may eliminate debilitating chronic pain, and devices capable of restoring lost or damaged sensory organs. Indeed, basic versions of these systems already exist. Eventually we may even see technology capable of improving cognition, expanding or reshaping the sensorium, managing emotional states, or enabling data transfer between digital and biological memory.

The common thread in all of these possibilities is the need to develop technology capable of enabling direct, high-fidelity communication between the human brain and a computer. This means that computers will be able to exert control over the mind. But who will be in control of such technology—the computer, or the brain? Programming or conscious choice? How do we determine which is in control? How do we determine who or what is responsible for the user’s actions? To answer these questions, we must open new doors in the realms of law and ethics.

This paper presents an exploration of how law and ethics may interact with and guide the implementation of a present-day technology that represents a significant step forward along the path towards advanced neuroprosthetics. While we are decades away from a society in which individuals routinely undergo elective neurosurgery for the purpose of facilitating deeper interaction with machines and computers, there are already FDA-approved medical devices capable of transmitting computerized signals into the brain.

Though the real-world precursors to this class of technology are still in their infancy, and significant boundaries exist in the form of medical ethics, restrictions on human research, and the sheer complexity of human biology, researchers are making steady progress on merging neurological and digital systems. With an installation base of over

100,000 users, so-called “brain pacemakers,” or deep-brain stimulators, represent the cutting edge of our ability to integrate computer technology directly into the human brain. Deep-brain stimulation (DBS) systems are an accepted and clinically effective form of neuroprosthetic treatment for a variety of common and debilitating neurological movement disorders. Most current implementations of DBS technology, however, lack the capacity to determine whether a user is currently experiencing pathological symptoms. These “open-loop” systems remain active continuously—whether or not user is experiencing symptoms.

Recent research suggests that more advanced implementations, where stimulation is selectively delivered depending on the state of the patient of DBS, are viable. Their research offers proof-of-concept for a closed-loop system capable of detecting the onset of tremor via co-implanted sensors, then delivering stimulation proportional to the symptoms’ duration and severity.¹ A closed-loop device offers obvious advantages: the minimization of adverse side effects, the capability to tune implant response to symptom severity, and extended system lifespan.

It is the basic capacity of human minds to decide. This capacity arises directly from the physical processes that occur within the human brain. It therefore follows that technology which manipulates those processes—like DBS—can alter our faculty for making decisions. Critically, our decisions have consequences. This paper investigates whether giving users *volitional* control over their DBS systems is ethically and legally responsible, despite the fact that putting the user in charge of the device creates the possibility of scenarios in which patient makes a wrong choice, and thus harms others.

We believe that it is not only responsible, it is in fact advantageous when compared to the alternative of making the system’s operation entirely automatic. From an ethical perspective, volitional control maintains (or even extends) the user’s autonomy both by allowing the user to decide when she receives treatment and by making it easier for the user to incorporate DBS into her life. This preservation of autonomy carries into the legal realm, where, through the tort system, giving users control of the system keeps responsibility in human hands.

The first section of this paper serves as a thorough, but by no means exhaustive technical briefing on a proposed volitionally-controlled DBS system. (VDBS). It addresses the history of the technology, the medical

¹ J. Herron and H. J. Chizeck. “Prototype closed-loop deep brain stimulation systems inspired by Norbert Wiener.” *In the Proceedings of the IEEE 2014 Conference on Norbert Wiener in the 21st Century*, Boston, MA, June 2014

conditions it is designed to treat, the criteria for patient selection, and suggests near-future developments and directions for research.

The subsequent sections of the paper tell the story of Janet Parsons. She is 64 years old, lives in the Ballard neighborhood of Seattle with her husband and caretaker Curtis, and is living with an advanced stage of Parkinson's Disease. Her disease has not been responsive to conventional therapy, making her a prime candidate for implantation of a VDBS system.

Janet, however, is not a real person—she lives in Seattle of 2018, at a time when VDBS systems have entered the marketplace, but are not yet widespread. Her story forms a narrative backbone through which we explore the issues and questions that are likely to arise as VDBS systems are implanted in patients. Chief among those questions: Do the benefits of giving users volitional control of the implant outweigh the risks to society and to the users' self-integrity? What types of training will be required of users before they are allowed control of the implant outside of clinical settings? What are the legal consequences of failing to use the implant responsibly? And, finally, under what circumstances might a patient's control of the implant be revoked?

Janet's story begins with her neurologist's recommendation for implantation. It follows her through surgery, through healing and acclimatization to the presence of the implant, and out into the world as a user of the technology. We pause the narrative at points where questions logically arise, and use these opportunities to explore while seeking technological, ethical, or legal solutions to the problems presented. The goal here is not to find concrete answers or argue in favor of strong conclusions. Rather, our goal is to stimulate imagination, debate, and to provide a point of origin for future research and exploration.

There are no great leaps of technological faith. Belief need not be suspended. We have done our best to ground the discussion in terms of what is possible today, or what will inevitably be possible within the next few years.

TECHNICAL BRIEFING

Neurostimulation refers to the delivery of therapeutic stimulation, in the form of electrical current, to the body's nervous system. Systems capable of delivering neurostimulation can be referred to as *neuroprosthetics*. Neuroprosthetics systems can be designed to interact with any portion of the nervous system, including—perhaps most importantly—the human brain.

Cerebral neurostimulation has been the focus of decades of clinical research, and has played a critical role in developing our understanding of the complex relationships between the physical structure of the brain and the functions of the mind and body. Neurostimulation can cause a subject to move,² to perceive somatosensory effects,³ and can even modulate a subject's willingness to engage in high-risk behavior.⁴

The most common form of clinical cerebral neurostimulation is known as *deep brain stimulation* (DBS). Used to treat neurological movement disorders such as Parkinson's Disease (PD) and Essential Tremor (ET), DBS stimulates structures near the very center of the brain, such as the thalamus, with the goal of mitigating the a patient's symptoms. Although current understanding of *why* stimulation helps patients is limited, the systems are effective enough that the devices are widely used, with an installation base of approximately 100,000 patients. In addition to the intracranial leads used to deliver stimulation to the brain, DBS systems require an implanted neuro-stimulator (INS) that houses the system's batteries and control circuitry. This is generally implanted within the chest, and is connected to the stimulation leads with a subcutaneous wire which runs up the back of the patient's neck. This allows for battery replacement without the need for additional neurosurgery.

The current generation of systems are quite simple: they deliver constant stimulation at a fixed intensity, set by a clinician. Although a patient's symptoms can be intermittent, current clinical systems lack the ability to sense their environment. They cannot modulate stimulation in real time in response to the onset or conclusion of a symptomatic episode. Instead, they provide constant stimulation. As such, current systems are referred to as *open-loop systems*.⁵

The clinician is responsible for programming the DBS system. Because the required level of stimulation varies from patient to patient, clinicians must use an exploratory, trial-and-error process of adjustment and observation to hone in on appropriate level of stimulation—one

² Graziano, Michael SA, Charlotte SR Taylor, and Tirin Moore. "Complex movements evoked by microstimulation of precentral cortex." *Neuron* 34.5, 841–851 (2002)

³ Romo, Ranulfo, et al. "Somatosensory discrimination based on cortical microstimulation." *Nature* 392.6674, 387–390 (1998)

⁴ HMM Smeding, AE Goudriaan, EMJ Foncke, PR Schuurman, JD Speelman, and B Schmand. "Pathological gambling after bilateral subthalamic nucleus stimulation in Parkinson disease." *Journal of Neurology, Neurosurgery & Psychiatry* 78 (5), 517–519 (2007)

⁵ Hebb, Adam O., et al. "Creating the feedback loop: Closed-loop neurostimulation." *Neurosurgery Clinics of North America* 25.1, 187–204 (2014)

which maximizes therapeutic benefit while minimizing side effects and battery usage.⁶

The ideal amount of stimulation varies not only from patient-to-patient, but also may vary from episode to episode, as most patients experience a range of symptom severity. Since current-generation systems lack sensors, they must be tuned to account for the most severe symptoms a patient may experience. Because of this, current systems waste power and require more frequent replacement of the INS, on a schedule ranging from a few months to a decade.⁷

Additionally, some patients report side effects caused by the stimulation. These include somatic effects such as tingling or burning,⁸ impaired speech,⁹ and even reduced inhibition or altered judgment.¹⁰ A patient receiving constant open-loop stimulation may experience these side effects at all times, even when they are not experiencing symptoms.

Despite these shortcomings, open-loop DBS systems are approved as an effective treatment for Parkinson's Disease,¹¹ Essential Tremor,¹² and dystonia.¹³ In research settings, DBS has been used to treat Tourette's syndrome, chronic pain, and depression.¹⁴ Studies suggest that DBS may also be an effective therapy for obesity, addiction, and dementia.¹⁵

The obvious solution to the shortcomings of open-loop DBS is the addition of sensors and control logic to the system.¹⁶ Modifying an open-

⁶ Kuncel, Alexis M., et al. "Clinical response to varying the stimulus parameters in deep brain stimulation for essential tremor." *Movement Disorders* 21.11, 920–1928 (2006)

⁷ Medtronic Corp, website, "System Eligibility, Battery Longevity: Neurostimulation systems for deep brain stimulation." <http://tinyurl.com/ltf8jh9> (Last accessed: March 18th, 2014)

⁸ Kuncel (2006)

⁹ Hariz, Marwan I., et al. "Multicenter study on deep brain stimulation in Parkinson's disease: an independent assessment of reported adverse events at 4 years." *Movement Disorders* 23.3, 416–421 (2008)

¹⁰ Smedding (2007)

¹¹ Deuschl, Günther, et al. "A randomized trial of deep-brain stimulation for Parkinson's disease." *New England Journal of Medicine* 355.9, 896–908 (2006)

¹² Sydow, Olof, et al. "Multicentre European study of thalamic stimulation in essential tremor: a six year follow up." *Journal of Neurology, Neurosurgery & Psychiatry* 74.10, 1387–1391 (2003)

¹³ Vidailhet, Marie, et al. "Bilateral deep-brain stimulation of the globus pallidus in primary generalized dystonia." *New England Journal of Medicine* 352.5, 459–467 (2005)

¹⁴ Perlmutter, Joel S., and Jonathan W. Mink. "Deep brain stimulation." *Annual Review of Neuroscience* 29, 229–257. (2006)

¹⁵ Miocinovic, Svjetlana, et al. "History, applications, and mechanisms of deep brain stimulation." *JAMA Neurology* 70.2, 163–171 (2013)

¹⁶ Hebb, Adam O., et al. "Creating the feedback loop: Closed-loop neurostimulation." *Neurosurgery Clinics of North America* 25.1, 87–204 (2014)

loop system by adding sensors and control logic is known as *closing the loop*. In a closed-loop system, each component of the system now interacts with the others, with the goal of responding to changes in the system's environment.

In the context of DBS, a *closed-loop system* (CLDBS) would deliver fine-tuned stimulation, responsive to the duration and severity of the patient's symptoms. It would be triggered automatically by the onset of symptoms, modulate stimulation depending on symptom severity, and withdraw stimulation when symptoms cease. This would directly improve a patient's quality of life by reducing the impact and severity of side effects and the frequency of INS replacement surgeries.

One such system has been prototyped by a team of researchers at the University of Nihon, led by Dr. Takamitsu Yamamoto. In their work,¹⁷ patients suffering from Essential Tremor wore accelerometers on their arms; software then interpreted the raw accelerometer data to estimate—in real time—the degree of tremor experienced by the patient, and adjusted DBS stimulation accordingly. This process occurred transparently, without input from the patient.

Another, more advanced, subset of potential closed loop neurostimulation systems would rely not on external signals such as accelerometer data. Instead, the system would directly measure the changes in neural activity that signal the onset of symptoms to determine what amount of stimulation is required at any given moment. Such systems can be referred to as *bidirectional neuroprosthetics*.

This form of direct measurement has already been shown to be useful in providing adaptive deep-brain stimulation to Parkinson's Disease patients. A group lead by Simon Little, at the University of Oxford, has used sensors embedded within the leads of a DBS system to determine when a patient may need stimulation.¹⁸ Parkinson's Disease patients who suffer from severe bradykinesia¹⁹ generate high amplitude beta-waves at one stimulation site for DBS. When Dr. Little's group delivered stimulation only when these beta waves were above the average level, the patient's movement scores were improved over the open-loop case. This indicates that by applying stimulation selectively, not only was the system saving power and reducing side effects, but was also more therapeutic.

¹⁷ Yamamoto, Takamitsu, et al. "On-demand control system for deep brain stimulation for treatment of intention tremor." *Neuromodulation: Technology at the Neural Interface* 16.3, 230–235 (2013)

¹⁸ Little, Simon, et al. "Adaptive deep brain stimulation in advanced Parkinson disease." *Annals of Neurology* 74.3, 449–457 (2013)

¹⁹ Bradykinesia is a form of hypokinesia (reduction in motor function) characterized by slowness of movement. Rather than being a slowness in initiation of movement (akinesia), bradykinesia describes a slowness in the *execution* of movement.

It is important to note that the closed-loop systems described so far rely on involuntary or *nonvolitional* tremor signatures—the user has no input in the system’s stimulation state. Patients may wish, however, to make deliberate tradeoffs between symptom reduction and side effects. For example, a patient may be willing to tolerate tremor in some situations to enable clearer speech. Patients could therefore consciously choose adjust stimulation levels in anticipation of movements that typically cause tremor, or in order to reduce uncomfortable side effects. As such, a *volitionally-controlled* DBS (VDBS) system may provide benefits over a (nonvolitional) CLDBS system.

It is well known fact that individuals who make long-term use of brain-computer interfaces can learn to control computer systems by thought alone.²⁰ This ability can reasonably be expected to extend to DBS systems as well. To perform these adjustments, the patients would “think” to the DBS device that they wished adjust the level of stimulation. This conscious change in brain activity would in turn be picked up by the device’s sensors, and the system would respond appropriately. It is this type of implant—a DBS system equipped with sensors and software capable of responding to conscious thought—which we will consider during our discussion below.

THE CASE FOR VOLITION

Janet Parsons was fifty-two years old when the tremors started. Like many things in life, the gradual onset of her Parkinsonism went unnoticed at first. Sometimes, when reaching out to answer the phone, she’d notice her arm trembling. A moment’s focus halted the tremor, and all was well. An active woman, she attributed her occasional tremors to fatigue from yesterday’s workout, or to the extra cup of coffee she’d had with lunch. Or, maybe, this was just one of the subtle ways her body was reminding her that she was getting older.

Janet was familiar with the symptoms of Parkinson’s Disease, but like many people facing a serious illness, she attempted to rationalize her worsening symptoms away in an attempt to avoid the reality of the fact that a life with Parkinson’s would be very different from a life without it. Her career, her hobbies, and her independence were all threatened by a shadow of a degenerative disease with no known cure.

Then came the day when she picked up the phone and couldn’t hold onto the receiver. The tremor had been getting worse for a few months, and with it, the amount of focus required to control her hands had increased. She had, for the most part, been able to manage her symptoms without

²⁰ Eleanor A Curran and Maria J Stokes. “Learning to control brain activity: A review of the production and control of EEG components for driving brain-computer interface (BCI) systems.” *Brain and Cognition* 51, 326–336 (2003)

significant disruption to her life. But as she stared at the telephone and listened to the tinny voice asking “Hello? Are you there?” she knew that she had no choice but to visit her doctor. Her diagnosis came quickly. The shadow she feared was swept aside and replaced with a tangible, inescapable thing: Janet had Parkinson’s Disease, and she was going to have to adapt.

Her doctor prescribed a variety of drugs—the blue Levocar and orange Tasmar tablets became a fixture of her daily life. Her symptoms abated, and in the first few years following her diagnosis, little changed in her private and professional life. But Parkinson’s is a progressive disease, and after a time the drugs and the physical therapy could no longer keep her symptoms under control. Eight years after her first visit to the doctor, she entered unofficial retirement, giving control of her share of the financial advising business she had founded to her partners.

As her condition worsened, Janet did her best to stay informed of new treatment options and therapies for Parkinson’s. She learned about DBS during one of her many forays into the medical literature surrounding her disease, and discussed the possibility of implantation with her neurologist. Together, they decided that DBS was not the right option for Janet. In particular, Janet worried that DBS would make little difference in restoring her independence—while her tremors might be reduced, she would likely pay for that benefit with impaired speech. The thought of trading one set of symptoms for another set of “side effects” didn’t seem worth it, especially given the need for invasive surgery.

At her last doctor’s visit, however, Janet’s neurologist told her about a new type of DBS capable of minimizing side effects while maximizing therapeutic benefit. She would be able to control the DBS system with her thoughts and might, in time, learn to control the implants almost instinctively. Together with ongoing drug therapy, this VDBS system offered Janet the best of both worlds—control of her symptoms when it mattered most, with minimal side effects at other times. Janet was hopeful that, for the first time in years, she might feel as if she had some control over her disease. Most enticingly, she might be able to return to work.

The range of side effects that DBS systems can generate represent alterations of basic features of conscious experience—motor control, somatic effect and executive function. While the scope and magnitude of “side effects” caused by DBS are limited, this technology indicates that is possible for a computer to interact with the brain at a very fundamental level.

Setting aside for the moment the deeper philosophical question of whether or not the very act of giving someone the ability to directly manipulate their brain *at all* is ethically responsible, let us instead focus

on whether it is ethically responsible to give users this control in light of the fact that their manipulations may cause harm.

To begin with, why is it desirable for Janet to have volitional control of her implant in the first place? If we assume that she was given a choice of either a VDBS or CLDBS system, why would she opt for the former over the latter? Put another way, which is better for Janet: voluntary or involuntary control?

If we recall, a closed-loop DBS system for Parkinson's can either detect the onset of symptoms, and then apply stimulation when the user needs it, or it can apply stimulation when the user gives a volitional neural command. While a nonvolitional system would manage Janet's tremor and reduce the amount of time she experiences unwanted, uncomfortable side-effects, she will still experience *some* side-effects. And, because the device determines when she needs treatment, those side-effects could occur at inopportune times.

For example, Janet may not be able to speak well when her DBS is active.²¹ If Janet was to give an important presentation at work, and her DBS begins to apply stimulation halfway through her presentation. In this case, Janet would likely rather deal with tremors than speech impairment—and it would be convenient if she could just turn the device off. The nonvolitional device, however, leaves Janet out of the decision loop. With the volitionally controlled device, however, Janet would always be in the loop. The device—in the best case scenario—would use a nonvolitional control scheme until the user gives a command to turn off stimulation. This would give Janet the power to decide when she wants to deal with tremors and when she wants to deal with the side-effects of neurostimulation, while minimizing her workload under most day-to-day circumstances.

In other words, a volitionally controlled device would likely give Janet a means of self-determination that would not be available otherwise. Ryan and Deci (2001) defend Self-determination Theory as an explanation of why it is important to respect and bolster self-determination in clinical contexts: they argue that perceived autonomy, competence, and relatedness to others are fundamental human needs. That is, if persons are not able to perceive themselves as autonomous agents that complete tasks competently, they may have trouble finding the (intrinsic) motivation they need to live a productive life. A plethora of studies demonstrate positive results when clinicians use techniques that both

²¹ The prevalence of speech problems in DBS patients has been reported to vary between 1% after 6 months, and up to 70% at three-year followup, with an average of 9.3%. Sabine Skodda, "Effect of deep brain stimulation on speech performance in Parkinson's Disease," *Parkinson's Disease*, vol. 2012

respect and support self-determination, Examples include encouraging people with diabetes to adhere to treatment (Williams 2009), and teaching people with Parkinson's-related postural instability to perform motor tasks.

Further, we can say devices like these, by extension, can be designed in a way that both respects and supports Janet's autonomy. There is, however, a great deal of controversy over what autonomy is and what it means to respect it in clinical contexts. Proponents of respect for decisional autonomy²² argue that a clinician respects a patient's autonomy when they present the patient with options, provide the patient with enough information to make an informed decision, and support the patient's decision after the fact. That is, a clinician respects a patient's autonomy when she is allowed to make un-coerced, informed decisions about her own healthcare.

Proponents of relational autonomy views²³ argue that we have focused too much on autonomous decision-making: they argue we must also "attend to the implementation of health care choices with significant self-management implications such as health-oriented lifestyle changes."²⁴ In the health-oriented lifestyle-changes approach, changes happen within the context of "cultural norms and social structures and practices affect the lives and identities we regard as valuable and possible for us" (Ibid). That is, it not only matters what circumstances bring a person to make the health care choices they do, it also matters that they are able to incorporate those choices into the rest of their lives. And so if we take on a relational view of autonomy, we must pay attention not only to Janet's decision to have a DBS system implanted, but also to how a DBS system changes her available options.

It follows that VDBS device would be designed in the service of Janet's relational autonomy and her decisional autonomy. Without volitional control of her DBS system, Janet can only take note of how stimulation affects her, then cope as best as she can until the next opportunity to consult with her clinician. With volitional control, not only does she have the option of allowing her clinician to calibrate the device, she can adapt the stimulation to meet her moment-to-moment needs. That is, when Janet has the power to decide when her system applies

²² Beauchamp, Tom L., and James F. Childress *Principles of Biomedical Ethics*. New York: Oxford University Press 2009

²³ Mackenzie, Catriona, and Mary Walker. "Neurotechnologies, Personal Identity, and the Ethics of Authenticity." *Handbook of Neuroethics*, 373–92, Dordrecht: Springer Netherlands, 2014.

²⁴ Entwistle, Vikki A, Stacy M Carter, Alan Cribb, and Kirsten McCaffery. "Supporting Patient Autonomy: the Importance of Clinician-Patient Relationships." *Journal of General Internal Medicine* 25.7, 741–45 (2010)

stimulation, she has the power to decide how to live with both Parkinson's and DBS (and its potential side-effects) in her particular life: in her workplace, with her family, or even by herself. Volitional control gives her the power to take ownership of her treatment in ways she would not have been able to otherwise.

On the other hand, it is a foundational premise of law that individuals bear responsibility for their actions. With regards to one's responsibility to others, people are held accountable through the frameworks of *tort law* and *contract law*. While DBS's neuropsychological symptoms may influence an individual's capacity to form contracts—for example, if DBS causes personality changes that impede the user's decision-making capacities—contract law governs the enforcement of explicit, voluntarily assumed responsibilities. We are more concerned here with a person's general responsibility to other members of society, and thus focus on the implications of DBS for tort law.

A volitionally controlled system serves to keep the user accountable. Assuming the implant is not malfunctioning in such a way as to deprive her of the ability to control it, she always has the ability to choose to turn the implant off (or on, depending on the circumstances), and therefore will always be responsible for her actions and the consequences that may flow from them. Even in cases where the implant is the “but-for” cause of harm, the law does not generally assign liability to *objects*, it assigns liability to *people*. Janet, as the operator of the implant (an object) bears responsibility for its use in the same way that one bears responsibility for the use of an automobile or a firearm. In the nonvolitional case, however, Janet would not be an operator. Her implant may decide to change state (and thus the possibility of risk of harm) at any time, and so she would either have to assign blame to the implant or restrict her behavior in such a way that a nonvolitional change in stimulation level could never present a risk of harm. Our intuition is that under no circumstances should Janet be able to blame her implant for her actions, and so Janet's autonomy would be restricted by the fact that she would be responsible for inadvertent changes in stimulation level. This may actually result in Janet having less freedom than she would be afforded with an open-loop system. The open-loop patient can always count on the presence of stimulation, and can take action based on that expectation. While some actions may be unreasonable for an open-loop patient, an open-loop patient will never be in a situation where what was once reasonable suddenly becomes unreasonable. A closed-loop, nonvolitional patient would potentially face that type of situation regularly. Janet may then have to refrain from any situation where a sudden change in stimulation could cause harm—even though she does not make a conscious choice

about the stimulation, she does make a conscious choice to put herself in a situation where harm could occur. Volitional control therefore allows her to operate in situations where nonvolitional control would be too risky.

There are technical advantages to volitional control as well. Currently, the physiological mechanisms that DBS therapy relies upon are poorly understood. While we are able to observe a clear relationship between stimulation and symptom abatement, we do not know *why* the one causes the other. Similarly, we do not understand why stimulation that controls motor symptoms also causes non-motor side effects such as changes in judgment, emotional state, or somatic effects. Thus it is simply not yet possible to design a perfect nonvolitional system. Hypothetically, a CLDBS system of ideal design would effectively become part of the patient's autonomic nervous system—using it would be akin to breathing. While individuals can control their breathing consciously, the brainstem is better equipped to manage the task autonomously in response to signals that operate below the level of conscious awareness. There may be a point at which our understanding of the brain allows a “perfect” DBS system to be developed, where including the user in the loop becomes suboptimal. At the present time, however, we are incapable of designing such an advanced system due to the limitations in our understanding of the brain. Thus, Janet will be in a better position to make judgments about the amount of stimulation she needs than will the system's control unit.

TRAINING

Time has passed since Janet's surgery. Eight weeks after being released, she visits Dr. James Bishop, a clinical neurologist at UW Medical Center. Dr. Bishop will be responsible for activating the implant, tuning it, and teaching Janet how to use the system to control her tremors. Her first visit is quite straightforward: the system is activated, and Dr. Bishop walks Janet through a series of basic motor tasks as he manually adjusts the stimulation level. Janet goes home that night with her VDBS system running in open-loop mode, applying limited stimulation. This acclimates Janet to the sensation of implant working, and allows her to see the system's benefits while learning to cope with limited side effects.

At her next visit, Janet is eager to be given the chance to begin learning how to control the system. Her hands are less shaky than they have been in years, and that was at less-than-peak stimulation! Dr. Bishop begins by adjusting several of the device's parameter settings in open-loop mode. He then switches the implant from open-loop to closed-loop mode, and walks her through a set of controlled tasks. Janet notices a difference in her symptoms as system increases stimulation in response to her tremor.

Eventually, Dr. Bishop leads Janet to a computer terminal, places a transcutaneous communications unit over her INS module, and opens a program on the desktop. A small cube appears on the screen and begins to change colors and pulsate as it rotates in front of her. Dr. Bishop walks Janet through the previous set of exercises one more time, and Janet watches while the cube's color shifts as the sensors in her brain detect the onset of tremor, and the cube grows as the level of stimulation increases. The system is then put into volitional-control mode, and Janet practices growing and shrinking the cube—first without stimulation, and then with it.

Over the next several weeks, Janet becomes an expert at the cube game, and several others besides. She learns what she can and cannot do with the implant, how to recognize the onset of symptoms and quickly control them, and how to make decisions about whether stimulation is appropriate at a given time. Eventually, Dr. Bishop stops putting the implant in open-loop mode when she goes home, and Janet is free to use the implant as she sees fit.

Volitional control poses new challenges for patients because the system will require a basic level of skill to operate. Janet will need to learn to modulate her brain activity in the same way that people learn new motor or cognitive skills such as driving a car or doing math problems.

There is a significant body of literature, discussed above, indicating that the brain can learn how to manipulate its own electrical signals to control artificial devices using brain-computer interface (BCI) technology. The co-implanted sensors of a VDBS system form just such an interface, and ought to allow Janet to control her stimulation merely by thinking about it.

There are two major approaches to teaching patients to operate the neural interface: *task-based training* and *operant training*. In task-based training, patients perform specific cognitive or imagined-motor tasks. Operant training does not give users specific cues for generating brain activity; users simply try to perform a task, such as moving a cursor, and learn through trial and error how to better control the device. In either training method, multiple training sessions are often required to become proficient at operating the interface. Though the exact training methods which will be used in teaching VDBS patients to control their implants are still in the conceptual stage, current thinking is to use a series of training games to develop a patient's skills using a flexible combination of task-based and operant training that is responsive to an individual patient's strengths and weaknesses.

We suggested earlier that a VDBS system can respect and support both her decisional and relational autonomy; but that only means Janet, her clinicians, her loved-ones, and her colleagues share responsibilities

they would not share other-wise. Prior to being taught how to use the VDBS system, Janet lacks the capacity to use the device responsibly. This does not mean, however, that she bears no responsibility for its use; arguably any use of the system would be irresponsible, and she ought to be held accountable—both morally and legally—for any resulting harm. From a legal perspective, that responsibility would be framed as a duty to not use the VDBS system until she was able to control it properly. This would follow the framework of *negligence*: Janet knows (or should know)²⁵ that, without training, making use of a device that can cause harm to others violates a person's general duty of care to act reasonably. An aggrieved party might sue her under such a theory, and she could be found liable for damages. At the same time, however, if Janet causes harm due to her lack of training, the clinician—as the party responsible for giving Janet that training—ought to take their fair of any responsibility for said harm. Any plaintiff could thus be expected to name the institutional parties (manufacturer, clinician, etc.) as well.

We do not need to invoke some notion of collective responsibility,²⁶ or even shared agency,²⁷ in order to describe how responsibility is shared between Janet, her clinicians, and her contemporaries. It is quite possible that, upon further consideration of future cases like Janet's, we find difficult-to-parse relationships between patient populations, neuroengineers, device manufacturers, clinicians, regulative bodies, legal professionals, and the public at large. These complicated relationships are, however, outside of the scope of this paper; and, for the moment, it seems sufficient to identify how each individual fails in some particular obligation they have because of Janet's implant.

We should note that at this point in time Janet's responsibility to others *still exists*, whether or not she is in a position to actually breach that responsibility. Thus, this responsibility can be expected to inform the design of the training regimen. The most responsible course of action would be to allow her to control the system only while in isolation (most likely at her clinician's office); the clinician would revert the system to either an open-loop setting or disable the system entirely while she is at home or in public. Only when Janet has demonstrated she can use the system competently, she would be allowed to operate it in everyday life.

²⁵ Presumably, Janet would be informed of the risks involved in using the VDBS system as part of the informed consent process prior to surgery.

²⁶ Allen, Pauline. "Accountability for clinical governance: developing collective responsibility for quality in primary care." *BMJ* 321.7261 608–611 (2000)

²⁷ Gilbert, Margaret. *Sociality and Responsibility*. Lanham, Md.: Rowman & Littlefield (2000)

It may seem unnecessary to disable the entire system while the subject is undergoing training to completely use it. One can argue that a VDBS system operating in open-loop mode is no different than a current-generation open-loop DBS system in terms of patient responsibility. A savvy plaintiff would likely note the novelty of the VDBS system in his or her claim against Janet and the clinician, arguing that the reasonable course of action would have been for the clinician to completely disable the system until she learned to control it. Rather than having to argue that Janet used the system improperly, and then by necessity considering the alternative courses of action Janet may have taken with regards to the relationship between the state of her implant and the circumstances at the time the harm occurred (an exploration we will conduct in a later section), one could argue that *any* use of the system was unreasonable. “Janet didn’t know how to use her implant, and someone got hurt,—it should have been turned off by her doctor” is a much simpler argument than “Janet made a decision about how to use her implant, and someone got hurt—Janet should have decided to use her implant differently.” Thus it may be safer for the institutional parties to deactivate the system after training sessions.

Liability is a major concern of medical device manufacturers and medical professionals, and can be said to have a chilling effect on innovation.²⁸ How then do we minimize institutional liability? If both Janet and the institutional parties share responsibility (and therefore liability) at the outset, the matter of training can be framed in terms of reducing the degree of responsibility borne by the clinician and manufacturer.

Janet’s duty to others is always the same: to prevent harm to others by acting reasonably under the circumstances. Again, our intuition is that at no time should Janet be able to assign responsibility to the implant; we should always be able to say that Janet is ultimately responsible for any harm that comes from her use of the device. We can therefore consider her degree of responsibility is fixed. Rather, it is the range of circumstances in which we can say she acts responsibly that changes. Prior to training, that universe of circumstances is quite small. Without the ability to control her implant, she would be obliged to avoid circumstances where her inability to control her stimulator could cause harm. Training would expand this universe. Janet’s skill at controlling the

²⁸ Garber, Steven “Product Liability, Punitive Damages, Business Decisions and Economic Outcomes” 1998 Wis. L. Rev. 237 (1998)

implant can be expected to be progressive,²⁹ so as she gains skill, she also gains freedom, while remaining responsible at all times.

But is Janet, in reality, always the bearer of ultimate responsibility? Is there ever a point in time at which Janet bears *no* responsibility at all? Perhaps. There are a few circumstances which present the possibility of Janet being absolved. First there is the possibility of a defect arising from faulty manufacture or system design. The manufacturer is responsible for designing and building a device that is free of defects. Were a defective device implanted in a patient, and that defect caused harm, it would make little sense to hold Janet responsible, so long as she was unaware of the defect. This defect may manifest as an unanticipated malfunction causing an excess or deficit in stimulation. Second is the possibility of a programming error by the clinician. The clinician is responsible for setting the device's parameters for responsiveness, and maximum/minimum stimulation levels. Sufficiently wrong parameters (say, resulting from clinician error during periodic post-training adjustments) may cause Janet to be unable to control her implant appropriately. For example, the clinician may accidentally reduce the maximum amount of stimulation Janet's implant can deliver. Later, unaware of this change, she attempts a dangerous motor task—say, pouring a cup of coffee for her husband. As she pours, a tremor begins. She attempts to activate her implant, but it does not respond, and she burns her husband's hand. Janet as done no wrong here; she had no reason to know the implant was not functioning properly. Finally, Janet may be presented with a no-win situation—one in which her implant is simply incapable of responding appropriately. While she could be expected to identify circumstances such as this through training, it is possible that Janet may find herself in a situation that exceeds the ability of the implant to respond. If this situation is entirely unforeseeable, she should not bear responsibility for the consequences (and, arguably, neither should the institutional parties).

While Janet can be said to almost always bear significant responsibility for her use of the implant, the degree of responsibility borne by the institutional actors varies. Just after surgery, before the implant is active, the only party to whom they are responsible is Janet. Once the implant is activated, they become, through their relationship and contact with Janet, accountable to third parties as well.

²⁹ Skill development times range from several hours to get greater than chance accuracy, to 80+ hours of training for fine control of the interface. Eleanor A Curran and Maria J Stokes. Learning to control brain activity: A review of the production and control of EEG components for driving brain-computer interface (BCI) systems." *Brain and Cognition* 51, 326–336 (2003)

Their chief responsibilities at this stage are twofold. First is their responsibility to Janet, to train her to use the implant.³⁰ She cannot make use of the system without training, and to deny her the ability to use the implant properly would violate her expectations of receiving a better quality of life through the device that she requested and they provided. Second, the responsibility to train Janet also extends to third parties, as without training Janet is more dangerous to others than she would be with it. Initially, the institutional actors would bear a significant responsibility to third parties—this is why training occurs in isolation. As Janet becomes more capable of controlling the implant, she also gains the knowledge and experience necessary to make decisions about her stimulation. Eventually Janet will reach a point at which she is competent enough to use the implant in her day-to-day life, outside the clinical setting. A threshold point will be reached where, absent one of the scenarios described above, it will be presumed that the institutional parties no longer bear any significant responsibility for harms caused by Janet’s actions. This point may not be contemporaneous with Janet being allowed day-to-day control—there may be a period of overlap where a plaintiff could claim that her training was inadequate, and, since the training was administered by the institutional parties, they are still responsible.

Thorough testing and training ought to minimize these risks, and is another factor weighing in favor of conducting training in isolation. Any defects in the device would hopefully become apparent during this period. Additionally, any time a change is made to the device’s programming, Janet should be required to demonstrate competence by completing training exercises. Training should focus on teaching Janet not only *how* to adjust the level of stimulation, but also *when* to do it, based on the circumstances. Training games might be designed to give Janet practice making decisions about stimulation and observing the consequences, allowing her to develop a strong sense of the relationship between stimulation level and physical ability. Finally, careful consideration must be given to the point at which Janet is free to use the implant as she sees fit.

³⁰ In some ways, training Janet is congruent with these actors’ responsibilities to themselves. If training Janet limits their risks, and that limitation is beneficial, then out of pure self interest the responsible course of action is to offer training. This is not just about risk, however. Training offers positive benefits in that it provides new insights into patient behavior and implant performance. Further, it is congruent with the notion of relational autonomy.

NEGLIGENCE

Janet has been living with full control over her VDBS system for almost a year, and has been gradually getting back to her old life. While her PD symptoms are still a significant consideration in her day-to-day life, her implant has allowed her to engage in many activities that, prior to surgery, were impossible. Her increase in fine motor control has allowed her to type for the first time in years, and she has returned to work. Her improved mobility allows her to go hiking again. And, perhaps most importantly, she is able to drive again.

On a Saturday afternoon, Janet is driving her grandchildren, Kevin and David, to her daughter's home in Redmond. Without the implants, Janet would not be able to drive, but she has turned up her stimulation to a level that allows her full control over the vehicle. Having taken several car rides with Janet since her surgery, Kevin and David have learned that she cannot speak well while she is driving—the implant, while allowing her to drive, significantly impairs her speech. Mischievous children that they are, it didn't take them long to figure out that their grandmother's inability to speak means that she cannot scold them, and have decided to take advantage of this car ride to settle a long-running dispute about who can make the more annoying noise.

Janet has had enough. She looks at the road around her, and sees nothing that presents itself as an immediate danger. Chastising the children will only take a moment, and so she reduces her stimulation so that she can speak. Her hands begin to shake, but she loosens her grip on the wheel so as not to transmit her tremor into the car's controls. As she turns towards the back seat, however, she notices a large piece of debris in her lane. She wills her implant to increase stimulation, but the system is not designed to respond instantly—ramping stimulation up or down takes several seconds. Unable to take evasive action, Janet's car hits the debris, and she swerves across several lanes of traffic as she tries to recover, sideswiping another motorist and causing both of them to hit the median divider. Janet and the children are thankfully unhurt, but the driver of the other vehicle suffers a broken leg and several other minor injuries. Janet's insurance agency denies the other driver's claim, arguing that Janet's condition voids her coverage, and that she never should have been driving in the first place. Having no other recourse, the other driver sues Janet for negligence to cover the damages.

We have chosen to examine negligence because no other tort encapsulates the notions of duty and responsibility to the extent that negligence does. While an intentional tort such as battery or conversion certainly has a flavor of responsibility, the given responsibility at hand (e.g. not hitting or stealing from someone) is quite specific. Negligence allows us to discuss responsibility in a broader context, examining the

various choices Janet may make regarding her stimulation in a given circumstance.

Negligence is commonly defined as *the failure to exercise reasonable care under the circumstances*.³¹ This failure may take the form of an act (doing something that a reasonably careful person would not do), or an omission (failure to do something that a reasonable person would do). This requirement is commonly known as the *duty of care*. A person breaches this duty when they engage in conduct that creates a foreseeable risk of harm to others. Should another person come to harm as a result of the breach, the negligent actor is liable for that harm.³² In determining whether or not someone was negligent we must consider whether, where a harm has resulted, a reasonable person would have taken a different course of action. With regards to the matter at hand, it is clear that the user of a VDBS system must exercise ordinary care to use the system so as to avoid foreseeable harms to others. This is not as straightforward as it may seem, however, as Janet's available choices for reasonable action vary with the amount of stimulation she receives. What may be reasonable while she is receiving stimulation may very well be unreasonable when she is receiving no (or less) stimulation.

It may be helpful to begin by considering the duty of an untreated Parkinson's patient, and progress through the various types of DBS system and analyze Janet's options with each type of system. The law recognizes that physical disabilities such as Parkinson's Disease change the scope of what may or may not be reasonable for a person to do.³³ This is not to say that a Parkinson's patient is somehow less culpable for their actions. Rather, the nature of the disease may make things that are reasonable for an able-bodied individual to do unreasonable for a PD patient to do, and vice-versa. For example, prior to surgery, it would have been unreasonable for Janet to drive at all during moderate to late stages of Parkinson's.³⁴ Her responsibility in this circumstance is clear: Driving would present a foreseeable risk of harm. That harm is easily avoided by refraining from driving. Therefore she should not drive. Her symptoms are unpredictable, and as such her responsibility to avoid driving is always in place; because she could become unfit to drive at any time, she should never operate a vehicle.

An open-loop DBS system would present the opposite circumstance. Here, Janet's tremor would be sufficiently controlled that driving would

³¹ Restatement (Third) of Torts: Phys. & Emot. Harm § 3 (2010)

³² Restatement (Third) of Torts: Phys. & Emot. Harm § 6 (2010)

³³ Restatement (Third) of Torts: Phys. & Emot. Harm § 11(a) (2010)

³⁴ It is, however, difficult to determine if people in mild-to-moderate stages PD have the ability to drive safely. See Heikkilä 1998.

not be an inherently unreasonable action. The fact that the stimulation is constant and of fixed intensity suggests that this would always be the case—she could rely on her symptoms being continually suppressed, and so could drive with the confidence that she would not have some sudden bout of tremor that could lead to an accident. So long as the implant remains functioning, Janet can drive.

Closing the loop complicates things. This is because the fact that the level of stimulation Janet is receiving will change over time, and therefore she must act in such a way as to account for those changes. Despite the fact that a nonvolitional closed-loop system may mitigate tremor as effectively as an open-loop system, her ability to drive would likely follow the case of an untreated Parkinson's patient. While an ideal device would be perfectly responsive to symptoms, in practice, current algorithm designs cannot provide perfect response. This variability between appropriate stimulation level and actual stimulation level introduces unpredictability into the system. Because such an implant effectively decides for itself when stimulation is (or is not) needed, and because Janet has no control over the system, it is foreseeable that the implant might unexpectedly alter her stimulation level while she is driving. This would in turn affect her ability to drive safely. With no means to accurately predict or control her level of stimulation, it would be unreasonable for Janet to drive.

This limitation extends across a broad range of circumstances: with a nonvolitional system Janet's stimulation may cut out and return her to the condition of an untreated Parkinson's patient at any time. Aware of this possibility, it would be foreseeable to her that harm may result if the implant dialed back the stimulation. Janet therefore has a duty to act as if the implant is not in place, restricting her possible range of actions in the same way that an untreated PD patient would.

There are further complications with a nonvolitional implant. Because the nonvolitional implant would control itself, Janet would also have to avoid circumstances where the unexpected *presence* of stimulation, rather than its absence, presents a risk of harm. DBS, in addition to enabling her to do certain things (namely those things requiring fine motor control), reduces her abilities as well; she pays for her enablement by making sacrifices in other areas of her functioning. Janet's ability to communicate effectively is perhaps the simplest of these situations to envision—she would have a duty to avoid any situation in which a sudden, unexpected reduction in her ability to speak could cause harm to others. This is a restriction not incumbent upon untreated PD patients—thus it may be said that implantation of a nonvolitional CLDBS system presents Janet with fewer options than she would have had she decided to

forego surgery! This is clearly undesirable, as the point of DBS therapy is to restore a patient's quality of life by enabling the patient to resume life activities that her disease otherwise makes impossible. This can only happen if the patient can know in advance what her implant will be doing in the future. Realistically, that means either using an open-loop system or giving Janet volitional control of her stimulation.

But with choice comes additional responsibility. A volitional closed-loop system ought to enable Janet to engage in activities that an untreated PD patient cannot do. To continue with our example, she may be able to drive. But, while driving, she will have a duty to not change the level of her stimulation.

Put more abstractly, one can begin at T1, where whatever circumstances Janet finds herself in are *neutral* with respect to whether she is or is not receiving stimulation. Whether her implant is on or off, she is acting reasonably. At some point in the future, T2, the circumstances change to become either *positive* or *negative*. If at T2, the circumstances become positive, and her implant is also in a positive state (i.e. applying stimulation), changing the state of her implant to negative breaches her duty. Alternately, if at T2 the circumstances become positive, and her implant is in a negative state, failure to change the state of her implant to positive breaches her duty. The inverse applies as well, of course, as the circumstances at T2 may become either positive or negative, and Janet's duty will be to adapt her level of stimulation (in whichever direction) to meet the circumstances.

VDBS may present an additional complication to this set of choices. It stands to reason that as Janet gains experience, she will become more adept at signaling the implant to change state. It is currently unknown how quickly a patient might be able to ramp stimulation up or down, or what degree of finesse a patient may have in selecting a precise level of stimulation. There is also no way of knowing, at this point, how much conscious effort Janet will need to expend to change the state of the system. While at first, it may require a deliberate, focused effort to control her implant, eventually it may become instinctual. Could Janet be held responsible merely for experiencing the urge to talk to her grandchildren, if her implant detected that urge and reduced stimulation based on that urge? Janet may, in effect, develop reflexive control of her implant. The law does not hold people accountable for their thoughts and urges alone. Would Janet have a duty to avoid circumstances where she might *think about* doing something that would require her to turn her implant off?

There are other unresolved issues as well. Janet's duty to control her VDBS system mirrors the duty borne by any operator of a potentially dangerous device. The driver of an automobile has a positive duty to use

the device in ways that promote safety, for example by using lights and signals to inform other road users of their presence and intentions. The same driver has negative duties as well, to refrain from using the device in ways that are dangerous, for example by not speeding in the presence of pedestrians. But a VDBS system is not just any device. It is, in effect, an integral part of the user's nervous system. Janet is not being asked to operate some external piece of hardware in a specific way. Rather she is being asked to alter the state of her brain. While in actual fact Janet is controlling a device that in turn causes an effect on her brain (arguably, akin to taking mind-altering medication), from her point of view this process is transparent. She simply thinks, and her state of mind changes along a continuum from maximum to minimum tremor, with proportional side-effects. While the law is unlikely to make such a fine distinction, from a philosophical perspective this raises the question of whether or not a person can be obligated to maintain a given state of mind, particularly in light of the psychological effects that DBS may cause in patients.

INTENTIONAL TORT

Janet returned to work shortly after learning how to control her VDBS system. It felt good to be back at the office, earning a living, and making a meaningful contribution to society. The depression that she had felt since she was forced into early retirement began to lift, and Janet found herself happier than she had been in years. Her time away from her profession, however, had cost her some of her edge, and the other partners at her firm decided that it would be better for her to ease back into the work by managing some of the firm's lower-risk investment portfolios.

While at work, Janet typically had her VDBS system set to provide fairly high level of stimulation. In addition to needing fine motor control to type and write, she felt social pressure to minimize the outward signs of her disease while surrounded by people who were not close friends. This use pattern was in keeping with the general purpose of her implants—to allow her to resume general life functions with minimal disruption from her Parkinson's symptoms.

Janet's performance at work was competent, but after a few years of providing decent returns for her clients, Janet's partners still hadn't allowed her to take on more responsibility. Janet resented this, and vowed to prove to them that she was just as capable as she had been before the onset of her Parkinson's. She began to take a more aggressive approach to portfolio management, engaging in risky trades in an attempt to boost performance. At first, when a gamble failed to pay off, she covered the losses out of her personal portfolio, but eventually she moved up to falsifying records and misappropriating funds from the portfolios of smaller clients. Her coworkers

noted a marked change in Janet’s personality—whereas her office demeanor had once been reasonable and understanding, she became grandiose and resentful of criticism.

But these changes in Janet’s personality weren’t the result of some simple need to prove herself. Instead, they were the outwards manifestation of DBS’ capacity to induce manic or hypomanic states in patients. Her continuous use of high levels of stimulation at work drove her to take more risks than she would otherwise, and blinded her to self criticism. After a time, this behavior spilled over into her private life as well, as she grew comfortable with the changes in her personality.

Eventually, of course, Janet’s luck ran out. After a particularly rough couple of weeks in the market, Janet’s partners finally conducted a full audit of her activity, and discovered her malfeasance. Janet was fired, and the firm’s clients brought suit against Janet and the firm for breach of fiduciary duty.

Whereas our discussion on negligence focused on the issue of whether Janet’s momentary choices to use or not use the VDBS system were reasonable in some instantaneous circumstance, here we recognize that generally reasonable use of the system may still result in harmful consequences to others. Rather than altering Janet’s ability to do or not do certain things, the psychological side effects of DBS can alter Janet’s ability to make decisions, and thus alter her behavior. But is a person responsible for the actions they undertake while experiencing these behavior-altering side effects?

Glannon and Klaming both argue that psychological continuity is key in determining an individual’s sense of personhood and self-identity. If a person experiences a sufficiently abrupt or severe break in the “continuity of the psycho- logical properties that constitute the self and one’s experience of persisting through time as the same person,”³⁵ it is possible that the person will feel “out of step” with what they take to be their authentic selves. Some open-loop DBS users have reported feeling out of step in just this way, and several argue that they feel this way precisely because some feature of their lives with open-loop DBS make them feel less like their authentic selves.³⁶ There is, however, an ongoing debate over what personal identity consists in, and whether or not DBS systems cause significant changes to their user’s identities on alternative views of identity. Francoise Baylis argues for a “dynamic, narrative, and relational” account of personal identity, where a person’s identity is

³⁵ Glannon (2008)

³⁶ Agid, Y, M Schüpbach, M Gargiulo, and L MalleT. “Neurosurgery in Parkinson’s Disease: the Doctor Is Happy, the Patient Less So?” *Journal of Neural Transmission, Supplemental* 409–414 (2006)

constituted by their life's narrative.³⁷ The issue, for Baylis, is not that DBS make a person feel less like their authentic selves by disrupting their psychological continuity—the issue is, instead, that DBS may disrupt the user's authorship over their own lives. Others³⁸ go further, arguing that DBS can disrupt not only the user's self-narrative, it can also diminish the user's "autonomy competence," or the series of competencies we rely on to make decisions in our everyday lives. Some examples of autonomy competencies are the ability to decide between several choices, the ability to interpret one's own emotions, and the ability to imagine the consequences of one's own actions.

Whatever disruption to a person's identity altered brain states cause, we can ask whether a person is responsible for any actions he or she may have made in an altered state. Generally speaking, the law holds that individuals are responsible for their actions regardless of their state of mind. Klaming notes, however, that the law makes some exceptions to responsibility in cases where an individual is involuntarily intoxicated. This may be akin to the case of a DBS patient with an open-loop system. A patient receiving continuous stimulation over which they have no control could be said to be "involuntarily" affected by any side-effects that occur. Assuming that the mental state induced by DBS side effects is equivalent to intoxication, it is arguable that an open-loop patient is no more responsible for her actions than someone who was surreptitiously drugged. This argument, however, is made vulnerable due to the fact that the implantation itself was voluntary, and thus the patient knowingly assented to any possible side effects. Even if it may be conceivable that an individual with an open-loop DBS system could make use of a defense analogous to involuntary intoxication, however, a volitionally controlled system such as Janet's would offer no such protection. With Janet able to turn the stimulation on and off at will, there would be no way for her to argue that her exposure to side effects was involuntary. Furthermore, the variable nature of the stimulation, which would be turned on and off by her conscious decisions, militates against an argument that she was experiencing any severe psychological discontinuity. Unfortunately, it is as yet unclear whether patients would be able to make responsible stimulation choices in all cases—that is, it is unclear if her autonomy competence has been disrupted.

In one case,³⁹ a 62 year old Dutch man began experiencing manic episodes approximately three years after being implanted with an open-

³⁷ Baylis, Françoise. "I Am Who I Am': on the Perceived Threats to Personal Identity From Deep Brain Stimulation." *Neuroethics* 6, no. 3, 513–26 (2011)

³⁸ Mackenzie (2014)

³⁹ Glannon (2014)

loop DBS system. Therapy with psychiatric medication failed to control the symptoms, which included megalomania and impulsivity. His psychological condition eventually degraded to the point where he was no longer competent to care for himself, and he was admitted to a psychiatric hospital. Adjustment of his DBS system caused the manic symptoms to abate, but in the absence of DBS the patient's Parkinsonism was so severe that he became bedridden. The three-year period between surgery and hospitalization suggests that his condition had not come on suddenly, and that his hospitalization was the punctuating event in a longer decline into incompetency. By the time medical professionals intervened, his decision-making faculty was so compromised that the doctors considered it unethical to allow him to make decisions about his treatment while the DBS system was active. Instead, he was allowed to make decisions about whether to continue DBS therapy only while in an unstimulated state. Eventually, he and his doctors decided that the best choice was for the DBS to be reactivated and for him to be legally committed to a psychiatric facility.

If this patient's symptoms were so severe that he was adjudged to be unable to make meaningful decisions for himself (let alone about his health) while the DBS system was active, what does that say about the ability of patients with VDBS systems to make sound decisions while in control of their stimulation? Had the Dutch patient had a VDBS system instead of an open-loop system, would he have turned his stimulation down when he began noticing adverse psychological effects? Or would his impaired decision-making abilities have driven him to leave the stimulator on? As Mackenzie and Walker argue, "an intervention such as DBS can disrupt a person's autonomy competence to such an extent that he is unable to engage in narrative self-revision."⁴⁰

Of particular concern is the fact we are uncertain whether a VDBS patient will be capable of easily recognizing the extent and severity of any psychological symptoms. Whereas an open-loop system would cause persistent side-effects, which would be observable to the patient's friends and family, the side-effects of a VDBS system would be more transient. They may only appear intermittently, during times of high tremor, or when stimulation passes some threshold. They may also take longer to begin to manifest, or manifest more gradually, owing to the fact that a VDBS patient will spend less total time being stimulated, and may spend much of that time receiving low-amplitude stimulation. While this may be a net benefit in the sense that the likelihood of developing adverse psychological side effects is lessened over the open-loop case, at the same

⁴⁰ Mackenzie at 390 (2014)

time it may make objective detection of those side effects more difficult. As a further complication, individuals in manic or hypomanic states often subjectively enjoy those states, despite the fact that they can observe the negative consequences of their mania. Mania is often accompanied by feelings of grandiosity or invulnerability, and so a patient may refuse to acknowledge the negative consequences of their stimulation. In the worst case scenario, Janet may even become “hijacked” or “addicted” to the system through a feedback loop that causes her to decide to leave the implant on at virtually all times when, in the absence of stimulation, she would be able to see the negative consequences of that decision.

Fortunately, severe psychological symptoms in DBS patients appear to be rare, and Janet, with her VDBS system, could be expected to have less risk than a normal DBS patient because of her reduced exposure to high-amplitude stimulation. This suggests that her day-to-day decision-making ability has a fairly high chance of being unimpaired, and it would be unreasonable to require Janet to turn off her system prior to making major decisions. But even in the event that her judgment was impaired, would it be feasible to require that she deactivate her stimulator prior to making important decisions? How would we enforce that responsibility? How could we draw a boundary around the class of decisions that necessitates that stimulation be turned off prior to making a decision? And what would stop a person from changing their mind when the stimulation resumed?

It may also be possible that there are some cases in which side effects become so severe that rather than a patient being simply unwilling to turn down stimulation, they may be unable to do so. In another case,⁴¹ a DBS patient with Tourette’s experienced severe dissociative symptoms approximately one year after implantation. This patient used a hand-held adjustment device to control his level of stimulation, but at a certain amplitude the patient began to experience effects similar to those seen in individuals with Dissociative Identity Disorder. Above a given threshold of stimulation, the patient regressed to an anxious, child-like state, and cowered in a corner of his room. Reducing the level of stimulation alleviated the dissociative symptoms, but the patient was unable to recall any memories from the high-stimulation period. It is unclear whether Janet, if she suddenly began experiencing dissociative symptoms, would be able to turn the implant off if she wanted to. Even in the event that she did turn down her stimulation level, she may not remember that anything

⁴¹ Klaming, Haselager, “Did My Brain Implant Make Me Do It? Questions Raised by DBS Regarding Psychological Continuity, Responsibility for Action and Mental Competence” *Neuroethics* 6, 527–539 (2012)

had gone wrong, and may not remember her actions during the time she was experiencing dissociative symptoms. Responsibility for such a situation could still be traced back to Janet's initial decision to ramp stimulation up to a point where a dissociative threshold was crossed, but Janet may not be able to know in advance where that point is (or if it even exists in her case).

There are several possible technical solutions which may help address these issues of side effect detection and impaired patient decision making. It may be possible to program a VDBS system to be at least somewhat responsive to tremor signatures without patient input. This will ease the patient's workload to some extent, making use of the system far more convenient. It may be possible to give Janet "budget" of stimulation to use over and above what the system deems necessary for tremor control at any given moment. Ideally, this could prevent Janet from applying constant overstimulation. This type of programming would be vulnerable, however, to a patient who intentionally induced tremor signatures by constantly attempting fine motor tasks. As a counter to this vulnerability, it may be beneficial to log Janet's use of the system. This would allow her clinician to determine how much stimulation she was receiving at any given point in time or over a period of time. These records could also be helpful to device manufacturers and researchers, providing a wealth of data about patients' brains, symptoms, and use patterns. Datalogging would also be use to Janet, helping her make objective decisions about her stimulation level. It may be possible to create an application that monitors and records the unit's stimulation level, as well as any tremor signatures detected by the system's sensors, and transmits that information to a smartphone or smartwatch. An application such as this could also gather subjective feedback from Janet by periodically asking her about her mood and other psychological criteria, and plotting those responses against the stimulator and tremor logs to develop a more complete picture of Janet's stimulation choices. The application could then make suggestions to Janet, making her mindful of her use patterns and helping her recognize problems before they become severe. This log would also provide an objective means for third parties to determine what Janet's system was doing at any given time. Log data would likely be discoverable in a legal proceeding, in the same way as other medical records or, perhaps more analogously, text-message or phone records. This would further restrict Janet's ability to assign blame to the implant in the event she causes harm to others.

CONCLUSION

Each of the issues we have discussed in this paper could easily be the subject of its own, far more exhaustive investigation. We regret that at this point in time we have neither the time nor the space to dive deeply into the many questions posed in the preceding pages. Hopefully, other papers will follow this one, and solid legal and philosophical frameworks can be developed for solving the problems posed by volitional DBS, while also promoting its adoption and future development. VDBS represents another step towards proving that significant integration of computer systems and consciousness is possible. We should bear in mind that humans are remarkably adept at turning “possibly” into “actually.” Put another way, we’re very good at turning fiction into reality.

Our culture’s fiction contains numerous explorations of what is possible when brains and computers are directly integrated. Many of these explorations present visions of the future where humans are empowered by this class of technology. Others, however, provide illustrations of the potential negative consequences of advanced neuroprosthetics. The truth undoubtedly lies somewhere between these two presentations. We should be proactive in ensuring that VDBS and its successors are designed and distributed in ways that are responsive to legal and ethical requirements for responsibility and autonomy.