Bard

Bard College

Bard Digital Commons

Senior Projects Spring 2016

Bard Undergraduate Senior Projects

Spring 2016

You Are Who You Choose To Be: Neuroscience and Freedom

Silas Edward Busch
Bard College

Follow this and additional works at: https://digitalcommons.bard.edu/senproj_s2016

 Part of the Philosophy of Mind Commons

This work is licensed under a Creative Commons Attribution-Noncommercial-No Derivative Works 4.0 License.

Recommended Citation

Busch, Silas Edward, "You Are Who You Choose To Be: Neuroscience and Freedom" (2016). *Senior Projects Spring 2016*. 127.
https://digitalcommons.bard.edu/senproj_s2016/127

This Open Access work is protected by copyright and/or related rights. It has been provided to you by Bard College's Stevenson Library with permission from the rights-holder(s). You are free to use this work in any way that is permitted by the copyright and related rights. For other uses you need to obtain permission from the rights-holder(s) directly, unless additional rights are indicated by a Creative Commons license in the record and/or on the work itself. For more information, please contact digitalcommons@bard.edu.

Bard

You Are Who You Choose To Be: Neuroscience and Freedom

Senior Project Submitted to

The Division of Social Sciences

Of Bard College

By

Silas Busch

Annandale-on-Hudson, New York

May 2016

For my family, who taught me self-discipline and responsibility;

For my friends, who taught me happiness and perspective;

For a partner who taught me compassion, morality, and so much more.

Acknowledgments

First, and foremost, my thanks, gratitude, and indebtedness goes to my advisor, professor, and friend, Kritika. She introduced me to philosophy, encouraged me to pursue it further, and has helped me to grow as a student and philosopher over several years. She has read and re-read drafts with care and diligence; her feedback has been impeccably on point, insightful, and helpful for me throughout this whole project. Despite all her input she was always sure to bring out the best in my thinking, not to substitute her own ideas, which is an impressive feat of clarity and is what has built me into the thinker I am today. Kritika made this project fun through laughter-riddled meetings, manageable through structured critique, and possible through her intellect, support, and inspiration.

My thanks also goes out to other members of the philosophy department: to Ruth Zisman, who also advised and encouraged me; to Daniel Berthold, who taught me Hegelian philosophy, which inspired much of this project, though there wasn't space to make that explicit; and to Jay Elliott, whose insight on my midway board pushed me to address a lot of problems with my thesis.

I would also like to acknowledge Barbara Luka, a former professor of cognitive science, who encouraged and inspired my early interests in neuroscience, as well as Arseny Khakhalin, who has taught me almost everything I know about current neuroscience, who has fostered my dual interests in neuroscience and philosophy, and who has provided vital input on my project as a member of my board.

Table of Contents

Abstract

You have just decided to sit down and read this abstract; perhaps, at the end of it, you'll decide to continue and read this whole thesis. Your ability to freely make that choice is simple enough, right? Or is the feeling of free willed choice an illusion? Metaphysicians, moral philosophers, psychologists, and neuroscientists are among the many who have weighed in on this question, offering a complicated variety of answers. Little work has been done to synthesize these approaches. My thesis is an attempt at a synthesis, showing that these perspectives can be reconciled by incorporating ideas from practical reasoning and recent neuroscience. Here, I review prominent philosophical and scientific work on free will, and offer my own proposal. My model of free will suggests that neurological activity can orchestrate indeterministic and deterministic events such that our conscious thought can generate and select novel criteria for neural activity. This resculpts neural pathways, influencing the channels through which future decision-making flows. I argue that this is tantamount to having a capacity for self-generating motivations, which allows both free and conscious choice in the moment, as well as free and unconscious choice in the future. Because the sculpting of neural activity and networks forces some unconscious activity to flow through channels of decision-making that we have already consciously endorsed, I argue that we are responsible for some of our unconscious decisions as well as conscious ones. Future empirical work can determine the extent to which conscious control can affect change in the brain, and to what extent unconscious activity conforms to neural sculpting. This, in turn, will give us greater insight into consequences tied up in our freedom of will: our moral responsibility, legal responsibility, and personal identity.

Introduction to Free Will

Science investigates how things work. Discoveries in the biological sciences provide biological explanations for natural phenomena, such as ecological interactions, cellular activity, or animal behavior. When neuroscience or cognitive psychology provide a biological explanation for the occurrence of a human behavior, it seems to subtract from the strength of our seemingly inexplicable free will to choose our own behavior. Much like how the actual role of gods in bringing about phenomena in the world are pushed further and further back by scientific explanations, the role of conscious free will, without a biological explanation of its own, is similarly reduced as a somehow supernatural force with smaller and smaller domain. Millennia ago, before we understood that the sun rises and sets because of rotations of the earth on its orbit around this ball of fiery matter, we attributed the movement of this celestial body to the flaming chariot of Apollo as he crossed the firmament. When relating this trend to human behavior, the issue becomes further complicated because it's no longer just a question about whether we can chock up the activities of the sun to deities or physical laws. When humans are involved, it introduces a new variable; now an explanation can be biological, supernatural, or human based. A huge question is: is a human explanation equivalent to a biological explanation, and in what way?

Before we understood such disorders as autism, major anxiety disorder, or depression (not that we really understand them today), we thought that behaviors of autistic, anxious, or depressed individuals were indicative of moral depravity, demonic possession, or divine punishment. The first explanation appears heinous in the modern day because it places blame

squarely on the shoulders of the individual for things which we understand now to be outside of their control. The latter two explanations seem to be mislead because they place the blame on supernatural phenomena when today we have at least some of the necessary research to suspect an entirely biological explanation. For some behaviors, such as the poor recognition of social cues by autistic individuals, it seems unfair to blame the individual, and reasonable to provide a biological explanation as an excuse. For other behaviors, like the choice to drive or walk to the store, it would seem fair to blame the individual for their choice, and invasive to chock it up to some series of biological phenomena.

When we enter the realm of scientific explanation, everything seems to come down to deterministic causality, where every event has its directly related, antecedent cause, and that the interaction between an event and its cause can be explained by a more or less complicated physical law. When it comes to human behavior, there seems to be a range of behaviors to which we can accept applying scientific explanations. How we define this range is often unclear, and varied widely based on personal belief and across cultures. Some people have no problem blaming someone with Asperger's for a selfish behavior, for example, while others are comfortable excusing it on the basis of the disorder. This waffling has severe implications for moral theories of responsibility. Does someone with Asperger's have enough control over their behavior in such a way that we can definitively say that they could have chosen a different behavior? If they did have control, then they should be considered responsible for their decision. If they are responsible, then they can be considered moral or immoral on the basis of whatever moral or ethical code we are using to judge.

How can we better define this vague sense of blameworthiness in the face of continuing scientific discovery that explains more and more behaviors? Does the discovery that alcoholism, for example, is attributable to genetic predisposition mean that alcoholics cannot be blamed for their choices? Of course not, they are morally reprehensible for their poor decision to drink excessively and irresponsibly, and likewise, they are morally valiant for their choice to overcome themselves and quit. At least, this is the case for some. It seems equally naive or ignorant to

blame all alcoholics for their behavior as it does to blame none of them. The catch is that not all alcoholics are equally predisposed by their genetics. There are different genetic causes, and magnitudes of cause for such behavior, which indicates that some people have less control than others. In order to assign moral responsibility, we need more information. We need to know the precise genetic background of the individual, and more research is needed to determine which genetic causes can be controlled and which cannot. As such, scientific research can provide support for a moral framework. On the other hand, how do we determine if some genetic component of our behavior is 'controllable'? What does control mean?

This is where free will comes into play, and where a cooperative project between philosophy and science can produce new criteria for free willing control that calibrates our empirical, genetic studies, and that informs our moral framework of which alcoholics are responsible and which are not. Note that alcoholism is not the only problem that these criteria can clarify, similar paths toward a solution can be achieved with other behaviors that are tracked down to genetic predispositions, such as smoking, gambling, and doping. Furthermore, these solutions are applicable to other forms of behavioral predisposition such as problems with temperament as a result of upbringing, environmental, or social ques.

Can we blame someone raised in an abusive household for violent anger, for example, later in life? It certainly wouldn't seem fair or accurate. More and more research points to the destructive and robust effects of traumatic and non-traumatic stress on the physical makeup of our brains, which would thus change how we respond to some stimuli, sometimes irrevocably. For someone who has been broken down until they no longer have control over some decision-making skills, then they cannot be blamed for their behavior, as they have no alternative available to them that they can access. Empirical research can tell us more about what degree of trauma is required to create such an ingrained change in the brain, what degree of trauma someone can recover from, and whether a behavior is driven by a controllable or an uncontrollable pattern of activity in the brain, and therefore whether the behavior is freely willed or not. These are, of course, advances that will take decades to uncover and optimize for practical

application, but having a model and a way forward is an important step toward such discoveries.

As I have shown here, sometimes science comes across as being in opposition to the existence of free will. Science is largely in the business of exploring the deterministic, causal relations in the world, whereas free will and consciousness seem to be concepts that may not be described by deterministic laws of nature (as we understand them currently). Some findings of neuroscience and cognitive psychology point to the idea that consciousness is 'just' the awareness of neural activity in the brain, activity that is unconsciously doing all the work of decision-making. There are three central responses to these claims: (1) compatibilists who respond that free will is compatible with determinism, even if it is true that consciousness is not causally involved; (2) eliminative incompatibilists who hold that free will cannot be compatible with a deterministic system, and that since science has described the universe deterministically, that free will is an illusion; (3) libertarian incompatibilists, who agree that free will is not compatible with determinism, but who also hold that there is indeterminism in the universe, which provides sufficient criteria for free will to exist.

In opposition to the first group, I will hold that determinism does exclude free will. In opposition to the second group, I hold that the success of quantum mechanics and the importance of noise in neural activity shows that we are neither saddled with determinism, nor a lack of opportunity for free will. My position aligns most with the third perspective, my only issue being that many Libertarians hold that the tiniest proof of indeterminism in the brain can pave the way for free will; I suggest that there is evidence for more than a tiny amount of indeterminism, and that free will extends beyond the domain of indeterminism. I hold that a mixture of indeterminism and determinism is at work in the brain in such a way that consciousness wields control over indeterministic neural noise and that through neuroplasticity mechanisms, consciousness molds neural networks, thereby shaping future activity. I suggest that a great deal of deterministic, unconscious neural activity operates within networks that have been molded by consciousness and can be considered an extension of conscious choices. As such, this establishes the possibility of free will during both conscious and unconscious decision-making.

I propose a two stage model of free will during decision-making that can be understood in terms of neuroscience or practical reasoning. The scientific version of my model shows that, in the first stage, indeterminism creates noise in neural activity, which provides a series of neural activity patterns that are available for selection by the agent's consciousness. When a pattern is satisfactorily favored by the agent through the direction of endogenous attention, the pattern molds networks in the brain through neuroplasticity mechanisms. Finally, activity that flows through this new framework deterministically recruits pre-motor areas and thereby precipitates action. This process can also be understood in more intuitive terms of practical reasoning. First, we generate a series of motivations, which are reasons to be or act a certain way. Through deliberation, we weigh the normative quality of these motivations to ascertain which best fits our practical identity and which is the most genuinely moving to us. Once we have selected a normative motivation, this acts as an unbreakable law for us, as we must uphold a law that we hold ourselves to (if we are genuinely motivated by this reason, then there is no way for us not to uphold it, because we are rational we must be a law unto ourselves), and we act on the basis of that reason.

In the chapters that follow, I will be outlining some contemporary perspectives on free will, proposing a model of free will that hybridizes Libertarianism and concepts from Practical Reasoning, and introducing contemporary neuroscience that supports my model. The goal of my first chapter is to set the scene. I will showing how philosophies of free will have been influenced by scientific discovery, and how, in some cases, a contemporary scientific worldview seems to conflict with a belief in free will. In this chapter I will focus specifically on three perspectives that represent the three main camps I discussed earlier: compatibilism, eliminitive incompatibilism, and Libertarianism. In my second chapter, my proposal will combine philosophies of free will and practical reasoning to defend a definition of free will as self-generated motivation. I will unpack my model at the philosophical level in the second chapter, while in the third chapter my goal is to stress the ways in which contemporary neuroscience can support my proposal. Together, this thesis aims to put forward a philosophy of free will that is

agreeable to some currently opposing perspectives, and that can be flushed out through empirical study. With hope, I can show that empirical verification can be reconciled with the seemingly non-empirical aspects of free will, such as responsibility, freedom, and identity.

I acknowledge the magnitude of this project, and I do not claim that this thesis provides a definitive, comprehensive account single handedly. I do not believe that this thesis even comes close to sufficiently substantiating a majority of the intricate or the overarching features of free will, neuroscience, and responsibility. My practical goal is to combine subsets of different disciplines from philosophy of mind to practical reasoning, from cellular to systems neuroscience, and in the end I hope to provide an empirically supported, but falsifiable model of free will. This model could reorient philosophy and science toward a more unified goal in their respective explorations of free will. In my conclusion I will attempt to briefly demonstrate the real world applications of this model.

It will become clear in my conclusory chapter that realizing the application of this model will require a joint effort of both science and philosophy, two disciplines which I consider prerequisites of each other. My scientific discussion will be carefully mediated by philosophical analysis, which I believe can, and should be, a guiding torch for science as much as science can and should be a guiding torch for philosophy. I believe that this approach will establish where we are and where we need to go. I believe we will come to discover how far away we truly are from answering this question empirically, though we can also be encouraged in this pursuit by having a clearer path ahead of us.

Chapter 1
Free Will as Ability, Liberty, and Illusion

Life is like a game of cards. The hand you are dealt is determinism; the way you play it is free will.

~Jawaharlal Nehru

In this chapter, I will briefly sketch some central philosophical positions on free will from both philosophers and scientists. As there are so many perspectives, I will have to truncate my review and focus on a couple positions that will represent clusters of divided views. After describing each in depth, I will also cite the concerns I have about them. The result will be a set of problems that any theory should address. In Chapter 2, I will address those problems when I lay out my position.

The first major division among the perspectives I will describe is the split between determinism and indeterminism. This division involves a discussion of the fundamental structure of the universe in which we live. Determinism holds that the laws of nature are determinately causal, which is to say that all phenomena occur through a set physical framework, a system of cause and effect wherein every result is directly caused by an antecedent event. For example, the result of bowling pins getting knocked over is caused by the bowling ball hitting them. Indeterminism posits that there are crucial components of the world that exist outside of regular causality, either in being probabilistic, or in being epiphenomenal, and can alter the ways in which the deterministic pathway proceeds. Probabilistic causality is a theory of causality wherein the event (the bowling ball hitting the pins) does not necessarily cause the result (the pins falling over). Rather, the bowling ball is influencing the probability that the pins will fall over, but does not determine it. Determinism as an idea is ancient, but it gained traction with the expansion of scientific and mathematical investigation in the eighteenth century, which clarified facts and

laws in the natural world. In contrast, even more recent scientific discoveries have resulted in the success of theories of uncertainty and quantum indeterminacy. These have provided indeterminism with the claim that determinism cannot take into account the inherent probabilistic nature of sub-atomic motion.

Now, to frame indeterminism in the context of free will. Indeterminism as a position posits that free will is not compatible with a strictly deterministic causal state of nature (bowling ball directly and necessarily causes the pins to fall). Further, some arguments for indeterminism hold that there is no reason to believe the inherent state of nature is deterministically causal. Panpsychism combines elements of determinism and indeterminism to suggest that free will cannot emerge in a purely deterministic world, not even from complex systems, as this would require a miracle. Galen Strawson argues from the position of panpsychism that free will does exist, and suggests that "all physical stuff is energy, in one form or another, and all energy, I trow, is an experience-involving phenomenon (Strawson, 2006; 18). This is to say that the individual components of complex systems that build free willing agents all possess in their natures aspects of consciousness whereby their combination can render consciousness and therefore free will.

Libertarianism is an indeterminist position largely pioneered by Robert Kane, who states that "there must be times at which the agents are able to choose in more than one way, and to do so voluntarily, intentionally and rationally, either way they choose" (Kane, 2014; 56) The libertarian proposal is that free will is an independent source of causality with which we make, what Kane calls self-forming actions in the face of multiple options.

Determinists fall, primarily, into one of two umbrella camps: Compatibilists and Incompatibilists. Compatibilists claim that the existence of free will is compatible with, indeed, some claim it requires, the deterministic causal nature of the world. Incompatibilists, on the other hand, hold that the nature of a truly free will cannot exist in such a causal realm. Among the compatibilists are Dennett, and Velleman, who suggest that consciousness and free will are constant narratives in the existence of our bodies, and McKenna, who holds a naturalist

perspective and rejects some of the assumptions that libertarians propose are required for the existence of a freely willed decision.

Incompatibilists hold that free will is not compatible with a deterministic natural world, but do not make the additional claim that free will exists in the way we feel it does. The scramble over deciding whether free will is or is not compatible with determinism has been largely pushed to the fore by the advancement of modern psychology and neuroscience that has cast what seems to be a discouraging light on the issue. With neuroscientists such as Libet and Haggard, the issue became how free will can exist if it is clear that consciousness is merely the result, and not the direct cause, of activity in the brain.

In his famous experiments, Libet tested the timing of brain activity and conscious awareness proceeding a decision and subsequent action (Libet, 1983). Libet was initially inspired by the presence of brain activity relevant to the action well before (over 500 ms) the motor network executed the action. In his experiment, neural activity was recorded in subjects who were told to perform a simple motor action (flexing their wrist) spontaneously. He found that neural activity in the motor and pre-motor cortex, readiness potentials, occurred approximately 800 ms before the action was executed, whereas the subjects' awareness of the intention to move, which they reported through positioning of a clock hand, occurred about 400 ms before the action. Thus, prerequisite neural activity preceded conscious awareness, a finding that Libet refrained from describing as a blow to free will until later (Libet, 1999), when many who read his work interpreted it as the death of free will.

Many oppose the idea that these experiments annul free will. On numerous occasions Alfred Mele has rejected their relevance, claiming that they do not invoke moralistic action or intention and therefore only comment on the conscious awareness of reflexive action (Mele, 2009; Ch.7). In essence, he claims that experiments trying to pry apart the topic of free will can only claim that the overt actions they investigate "do not have corresponding consciously made decisions or conscious intentions among their causes" (Mele, 2013; 2). Mele goes on to claim

that there's no reason to expect these actions, flexing a wrist, to involve free will. Invoking the relevance of moral responsibility in the topic of free will, he claims that the actions investigated have no moral import; flexing a wrist (in the case of an experiment, as opposed to a case where it might signal an order to kill, for example) does not present a moral decision. There's no reason not to do it, so there's no reason to deliberate, consider pros and cons, and to exercise your free will to make a choice. Mele posits that, since free will isn't utilized in the actions investigated, the studies do not make any claim, positive or negative, about the existence of free will. While Mele might be correct in asserting that heretofore neuroscience has left the issue wide open, there is reason to believe that decisions of any sort, with or without moral implication, will proceed through the same pathway. It's not like if there are two brains, one acting unconsciously, the other acting consciously. Both states of processing, conscious and unconscious, proceed through the same neural pathways. As we will see in Chapter 3, there are just extra mechanisms involved during conscious processing.

Eliminitive incompatibilists claim that the feeling that our phenomenological aspect of free will, the part we experience consciously, has a causal role in our decisions is an illusion. Eliminative incompatibilists claim that unconscious neural activity dictates how the body will react to external stimulation, whereas our conscious experience of free will is an interpretation of this unconscious decision. Importantly, this interpretation occurs just before the action itself, is usually consistent with the action, and is therefore entirely convincing of the idea, the illusion, that the activity is caused by the conscious intention to act. Illusionism, so it's called, is supported by members of the scientific community (Wegner), in addition to philosophers (Smilansky). Smilansky proposes that "illusion has a large and positive role to play in the issue of free will" (Smilansky, 2000: 6; see Smilanksy, 2002: 11). More specifically, he says that "humanity is fortunately deceived on the free will issue, and this seems to be a condition of civilized morality and personal value" (11). This is seemingly in line with an evolutionary biological perspective that we evolved to possess this nearly unimpeachable illusion because it makes for the most productive means, vis a vis moral responsibility and legal systems, to

preserving our species. Similarly, Wegner wraps his version of illusionism in the proposal that we would be wrong to think that the illusion of free will denounces its import and prevalence.

In this chapter, I will focus on the potions of Wegner, Kane, and McKenna. These three will act as representatives of their respective camps (eliminitive incompatibilism, indeterminism, and compatibilism) in order to simplify the process of describing the variety of positions and the problems with each that I will address with my position.

Wegner and Free Will as Illusion

Daniel Wegner uses mental pathology and strange phenomena that distort perceptions of agency (trance states, Ouiji board spelling) as a framework for his theory of free will. He uses cases in which people act without willing it or will it without acting to study that he calls the theory of apparent mental causation.

To begin, Wegner describes a relatively obscure, but frightening pathological case of action without will: alien hand syndrome. He explains that alien hand syndrome is a neuropsychological disorder wherein the patient experiences the movement of their hand as if it had a mind of its own, a result of damage to specific region of the brain on the side contralateral to the alien hand. One patient's hand would:

> Tenaciously grope for and grasp any nearby object, pick and pull at her clothes, and even grasp her throat during sleep... She slept with the arm tied to prevent nocturnal misbehavior. She never denied that her left arm and hand belonged to her, although she did refer to her limb as though it were an autonomous entity (Banks et al., 1989; p. 456)

This example highlights the difficulty of understanding an experiential phenomenon as it can only be self-reported by the individual that they willed an action or not. In the case of someone with alien hand syndrome, it may appear to a third person perspective that the patient willingly

moved their hand, yet the patient would point out, most likely with fear and dread, that they did not have the conscious experience of authoring the movement of their hand. This is an example of a person who can perform an action while the feeling of doing it is absent.

These two variables, doing and the feeling of doing, delineate four ways in which we can experience action and will. Doing something and not feeling that we've done it, like with alien hand syndrome, results in a condition Wegner calls automatism. On the flip side, not doing something and feeling like we have is an illusion of control, like people thinking they control the function of a machine by pressing a button when the button really only works on random occasions (as is the case in some psychology experiments that test the range of our sense of contingency between pressing a button and something happening). Wegner calls these the special cases; normally, our activities fall into one of two other options: not doing something and not feeling like we're doing it (or normal inaction), and doing something and feeling that we've done it (normal voluntary action). Wegner's position is a little weak in that he leaves it unclear how we can and should distinguish between all of these cases, most of them requiring self-reportage, which may not be possible, or may not be truthful. This inability to distinguish clearly is one problem that I will address in my proposal.

What is most interesting, in light of how pathology can alter our perception of normal voluntary action, is how voluntary action works. To determine whether our experience of free will is capable of causing neural activity, we must determine if it can be observed as a path of cause and effect. Causality, Wegner points out, "is an event, not a thing or a characteristic or attribute of an object" (Wegner, 2004; 3), so we cannot rely on the sense that our experience inherently plays a causal role. We must investigate if experience can be empirically identified as a cause of an observable effect.

Philosophers and scientists familiar with the works of Libet (wherein readiness potentials are present and initiating action before the conscious will to do so) and Wegner become startled by the concept that free will requires empirical verification. If this is the case, then surely we

will lose free will entirely because the science seems to imply that conscious experience is a byproduct, and therefore does not play a role in the pathway of cause and effect. In Wegner's precis to his book, The Illusion of Free Will, he includes a section for open peer commentary, and a follow up response section. In his responses, he cites a couple positions from neurosurgeons and psychologists (Bogen, Kihlstrom, and Zuriff) whose peer commentary reflect the interpretation that the experience of free will is an illusion according to Wegner's argument.

Wegner distinguishes between empirical and phenomenological free will. Phenomenological free will is the experience, or sensation, of free will. Wegner posits that the experience does not have a causal role in our actions. In his terms, the experience of free will is not predictive enough of action for it to be a valid causal factor. We've seen how the correlation between the 'feeling of doing' and actually 'doing' can vary as a result of pathology, strange phenomena, or plain trickery on the part of psychology researchers. Wegner makes an even stronger case by elaborating a variety of cases with misconstrued phenomenological free will, such as the patient with alien hand syndrome. By observing the lack of predictive power that the experience of free will has on our actions, we can determine that it cannot play a causal role in all action.

On the other hand, Wegner describes the empirical free will as the "real causal sequence underlying human behavior" in which "each of our actions is really the culmination of an intricate set of physical and mental processes" (5). Interestingly, Wegner is sure to point out that this includes "psychological mechanisms that correspond to the traditional concept of will – in that they involve linkages between our thoughts and our actions" (5).

This is a concise, logical argument that culminates in Wegner's theory of apparent mental causation where people "experience conscious will when they interpret their own thought as the cause of their action" (6). Wegner proposes three key sources of this interpretation: priority, consistency, and exclusivity. Thus, according to his theory of apparent mental causation, we have the feeling of control, as agents of our actions, when our thoughts (A) precede said action,

(B) are consistent with the performance of said action, and (C) when we do not perceive an overwhelming sense of alternative causes of said action. These sensations converge into a strong feeling of authorship, which we experience as conscious will, as "Yeah, I did that." Wegner argues that this feeling is an illusion of control we obtain from the rapid interpretation of many coincidental sensations. As Wegner writes with sobering conciseness, "experiences of conscious will thus arise from processes whereby the mind interprets itself – not from processes whereby mind creates action" (1). He claims that the experience of free will itself does not have sway over actions. The experience is only an interpretation of the events in our brain that cause the action.

Wegner presents a fascinating example of how an action that is clearly not caused by us, the moving of a tree branch, can feel authored if these three factors aligned. If you imagined that branch moving and then it moved, and moved exactly as you imagined, and was not caused by anything else, like a squirrel, you would probably, first, be very taken aback and then, second, begin to consider, no matter how absurd it seems, that you might have caused the movement. Imagine that this predictiveness occurred all-day every-day as you interpreted the actions your body performs: you would no longer consider it strange and would assume, as we do, that we, that is, our conscious wills, have caused the action.

Wegner is very aware that this arouses fear that we will lose free will, and he is very careful to distinguish, again, that he is only calling the experience of free will an illusion, not the empirical fact that mental events can influence our actions. However, I do not agree that he addresses these concerns enough. He asks why this experience exists: "What good is an epiphenomenon?" (8) and claims that phenomenal free will provides us with a sense that our body has authored an action, like how emotions provide us with a sense of how our body is responding to events. As, such, conscious will "serves to accentuate and anchor an action in the body" in a way that "renders the act one's own in a personal and memorable way" (10, emphasis added). Ultimately, conscious will is an important check that "serves as an indicator of mental processes or states, and that thus informs us about the status of our own mental systems" (10). My contention is: who is the 'us' in that sentence. Is 'us' our conscious will? If so, is making our

conscious will aware of this check advantageous even if there's no way for 'us', our conscious will, to turn around and use that knowledge to reorient the body? What makes more sense is that 'us' means our unconscious brain. This would this suggest that consciousness is just a tool the brain uses to check up on itself.

Wegner's claim is that this is enough. We can rest assured that our body has empirical free will, which is to say that the actions our bodies perform are influenced by mental events. Even if our conscious will lacks true causal control, it is, in the end, usually an accurate interpretation of what is actually happening in the brain. Given this position Wegner describes a philosophy of responsibility taking into account the acausal nature of conscious will. He claims that:

> Causality is something you can see in mechanical systems, a relationship between events, and is not dependent on what kinds of events are involved. Responsibility, on the other hand, involves persons – those selves that are constructed through the process of identifying actions as caused by an agent (39).

Thus, responsibility and causality are like apples and oranges to Wegner. Equating responsibility to causality is a mistake to Wegner.

Responsibility is a human question that involves our unique 'selves', and our sense of authorship. Wegner writes that "the allocation of responsibility is dependent on regarding something as a person," because only a person can regard oneself as a person, and only something that regards itself as a person can "take an intentional stance toward itself" (39). According to Wegner, the intentional stance occurs naturally through the power of the illusion of phenomenological free will. When we take this intentional stance toward ourselves the "real work of responsibility is done" because the accused person "accepts responsibility and experiences moral emotions such as shame or guilt (or in the case of positive actions, pride)" (39). While I agree with this perspective, that a great deal of work could be done by such an illusion, to which I will return in Chapter 2, I object with the claim that causality is not relevant

to responsibility.

Just as much as our feeling of guilt in the face of action is based on thinking we have control, an equal and opposite feeling of innocence can be based on thinking we do not have control. Can we merely trust that this illusion of free will is infallible, unimpeachable? I believe that some of the strength of this illusion is endowed by the cultural background in which we are raised. From the youngest point when we can be considered thinking human beings, we are raised with the assumption of choice, and responsibility. Some people end up feeling responsible for different things merely because there were raised to believe they are responsible for some things and not others. This indicates the influence of cultural background over the illusion. If we accept, unanimously, that the experience of free will is an illusion, and this becomes common knowledge that future generations are raised believing, who's to say if they will feel the same sense of intention toward themselves?

As such, my problems with Wegner's view are as follows:

1. He does not provide sufficient guidelines for how to distinguish between voluntary and involuntary action, especially in light of his claim that phenomenological free will is not properly predictive.
2. He does not sufficiently address what the goal of the phenomenological free will is, if it has no causal role.
3. He does not account for the fact that sensations of free will are changed by our social and personal conception of responsibility.
4. He does not account for empirical findings that consciousness does have a causal role in decision making and information processing in the brain.

Kane and Free Will as Liberty

Robert Kane has been the dean of Libertarian philosophy for decades. He has

summarized his and other Libertarians positions on countless occasions, and single handedly written and edited a library of books and articles on the topic of free will, including the Oxford Handbook of Free Will. His most recent article is a succinct 2012 contribution to a symposium on Mark Balageur's book, Free Will as an Open Scientific Problem, in which he, Kane that is, compares and contrasts his position with that of Balageur's, a fellow libertarian. I will start toward the end of this piece, "Torn decisions, luck, and libertarian free will: comments on Balaguer's free will as an open scientific problem," where he more clearly defines his position, and then work my way through the beginning when he distinguishes key aspects of his theory, which should highlight some of the diversity in Libertarian perspectives.

Philosophical Libertarianism, in general, holds that the human capacity for choice is incompatible with the domain of physical determinism of any sort. Thus, libertarianism holds an indeterminist stance about the way our consciousness makes decisions. For the most part, disagreement among Libertarian philosophers centers around what the nature of indeterminism is as it presents in the brain, and at what step in a decision making process it occurs.

Within the range of Libertarian views, Kane identifies Self-Forming Actions (SFAs) and Plural Voluntary Control (PVC) as central pillars of his position. SFAs are those activities in the brain over which the individual has indeterministic control and therefore define the range of activities for which an agent is responsible. It is the process by which an agent self-determines a choice, or decides which option is the most desireable, or best in some way. In being the origin of this decision, not having been deterministically influenced by the exterior world, we would have control over and thereby responsibility for actions precipitated by those choices. Kane is careful to note that "not all acts done of our own free wills have to be undetermined, only those acts by which we made ourselves into the kinds of persons we are—namely, the "will-setting" or "self-forming actions" (SFAs) that are required for ultimate responsibility" (2007, 26). This is to say that the import of SFAs rests in their enabling us to generate a decision about choices that define who we are, or choices with which we identify and of which we can consider ourselves agents.

The libertarian view of free will shares a great deal of qualities with my own, however, I disagree with what Kane distinguishes as most important for the existence of free will, and with the likelihood of certain factors having a role. Kane declares that free will is "the power of agents to be the ultimate creators (or originators) and sustainers of their own ends or purposes" (1996, 4). The reader will note that this relates very closely to my theory of free will as I briefly noted in the introduction: "Free will is our ability to be free to choose, to generate our own decision without the complete control of our biological makeup and external circumstances… Being able to generate our own unique motivation for a given action is free will." Self-forming actions are what make us the ultimate creators, or originators, of our own ends and purposes. SFAs are something that we consciously create in order to precipitate actions toward our purposes.

How do we come to create these SFAs? Kane writes that SFAs "are the end-products of goal-directed or teleological cognitive processes in the brain in which indeterminism in some of the neural events involved in these cognitive processes functions as a hindrance or obstacle to the attainment of their goals" (2012, 56). Thus, he identifies indeterminacy in the brain as a key factor by establishing requisite obstacles. In providing difficulty, the indeterminacy behooves us to make an effort toward one goal or another. As such, SFAs present an effort of will to proceed toward one goal or another that is distinguished from when the brain merely proceeds through unhindered, deterministic pathways toward thermodynamically favored results.

The way in which we create SFAs requires a condition that Kane calls Plural Voluntary Control (PVC). According to Kane:

> Agents have plural voluntary control over a set of options (e.g. choices) when they are able to bring about whichever of the options or choices they will to bring about, when they will to do so, for the reasons they will to do so, on purpose, rather than accidentally or inadvertently or by mistake, without being coerced or compelled in doing so or willing to do so or otherwise controlled in doing so or willing to do so by other agents or mechanisms (2012, 56).

Essentially, Kane is describing the PVC nature of free willing causality as a situation wherein an agent is confronted by multiple possible options and has the liberty to "choose in more than one way," hence plural, and is performing this choice in a way that is not influenced by external mechanisms, or, in other words, is able to choose "voluntarily, intentionally and rationally" (2012, 56), hence voluntary control.

Kane distinguishes between his theory of PVC with a common Libertarian theory that only requires Alternative Possibilities (AP) to satisfy free will. AP allows for what was originally thought of as an ambiguous future where the result is not predictable to one, single, determined solution based on the given causes. Kane argues that AP is not enough (though it is part of the solution), there must be the factor of effort required to achieve one possibility or another, this is the 'willing' component of free willing.

Kane also refers to these goal-directed cognitive processes as 'volitional streams.' This allows him to elaborate further on the nature of control, or liberty, that an agent has over a decision. Kane posits that "one must allow that there can be multiple such volitional streams within the agent," which vaguely appeals to the theory of AP, "the agent must have a certain kind of control over each of these volitional streams," and "the choice that results from either of the volitional streams can be said to be brought about by the agent" (2012, 57), which assures that the agent represents the origin of the choice, the creator of the decision to proceed through one volitional stream or another. Kane concludes that "it is owing to the fact that either volitional stream might succeed in attaining its goal, which would thereby be brought about by the agent's goal-directed cognitive process, that the agent exercises plural voluntary control over the decision itself (the SFA)" (2012, 57). Thus, Kane posits that control over the choice and the effort of willing one choice or another are two pillars of SFAs, which enable agents to act with free will.

Let's look more closely at what Kane identifies as an indeterminacy in the brain. In agreement with Balageur, Kane cites the existence of "microlevel indeterminacies in the timings

of firings of individual neurons that would give rise to some background neural noise." Kane argues that these low level effects, which are currently understood to be a result of quantum uncertainty, are enough for free will. While some claim that the brain as a whole, and the decisions and actions it generates, is too many orders of magnitude too large for these low level indeterminacies to have any noticeable effect, Kane claims that this is okay. "One doesn't need a thunderclap of indeterminism in the brain to get free will," he writes, "just a sprinkle will do" (2012, 52). The precise nature of these microlevel indeterminacies will be examined more in the proceeding chapters. However, given this stance, I will state, now, that I disagree that this "sprinkle" is enough for free will. Kane is underselling the import of indeterminism and how it influences the higher structures of the brain. If the indeterminacies are indeed erased through supersession by orders of magnitude of complexity, and if these indeterminacies only amount to a random, content-less, neural noise, then this would account for the existence of APs, but not of a radical PVC, which Kane claims is necessary for free will. Of course, I cannot blame Kane for not pursuing the science himself, but I object with his accepting that this, as he recounts, is at best an open scientific question. Given his account of what we currently know, it doesn't seem like PVC is part of the equation. Ultimately, I disagree with how he interprets the science; there is scientific support for a great deal more indeterminacy in the brain than a sprinkle, albeit not a thunderclap.

In opposition to Balageur, Kane insists that self-forming actions, which Balageur calls torn decisions, cannot be only phenomenological, or experience based. If they were, they would only represent the way in which our decisions feel like they were not deterministic, while leaving the options open for whether or not they were determined on an ontological level. Kane claims that "one cannot resolve free will questions by appeal to introspection alone" (2012, 53). Balageur only represents his torn decisions phenomenologically, but if he also meant for them to have implications on an ontological level, then, Kane insists, "they would amount to my SFAs" (2012, 53). This is only a semantic disagreement between them, but it reminds us of a key requirement for free will: that it must not be just phenomenological. Recall Wegner's distinction

between phenomenological free will and empirical free will. In order for Kane's libertarian free will to exist, the phenomenology of free will must have a causal influence on the physical actions (the empirical, or ontological aspect) of the brain. We must be at liberty to act upon the world, and ourselves, in the way we wish. Otherwise, we slip back into Wegner's view that the phenomenology of free will is an illusion that convinces us to treat empirical free will in the same way as we would treat conscious free will.

It is important to establish the radical nature of Kane's view as a libertarian. Many libertarians are satisfied by simple two-stage models of free will. Two stage models represent an explanation of free will that dates back to William James, who articulated that free will could be defined as a chance occurrence of alternative possibilities, followed by a choice made by the agent to execute one choice or the other. Kane poses a stronger two stage model wherein practical choices (procedural things like deciding how to get from the attic hideout to the safe house two streets away) are readily explained in this model, while moral choices (should I escape to the safe house, or turn myself in) require a more controlled version of the first, chance stage. His PVC theory describes how the role of indeterminacy in rendering the choice difficult to make forces the agent to exert effort, to will a choice. This added requirement places Kane in a somewhat extreme place as far as libertarians are concerned.

Thus, my concerns with Kane's Libertarian perspective are as follows:

1. Where do we draw the line at actions that have "made ourselves into the kinds of persons we are"? Does this involve intention vs. no intention, endorsement vs. no endorsement? Kane doesn't say.
2. Kane suggests that only a sprinkle of indeterminacy, and thus of the potential for free will, is enough. Maybe a sprinkle would be the bare minimum, but there is evidence for more.
3. Kane discusses goal-oriented behavior as the source of SFAs, but does not discuss how intention and motivation lead an agent to select one goal over another. According to

> Kane, it is unclear whether a free willing agent has control over the goal they pursue with the libertarian free will he describes.

McKenna and Free Will as Ability

The central strength of Michael McKenna's compatibilist position is the defence of Harry Frankfurt's position that refutes the need for Alternative Possibilities. Frankfurt suggests that if a decision is morally significant and the agent has control, then there is no need for alternative possibilities. Opposition from incompatibilists suggests that this implicitly assumes that the entire process is deterministic, and therefore contentious in another way. McKenna defends Frankfurt by claiming that an example can be made wherein only some alternative possibilities are blocked while a decision is in control of the agent and is therefore morally significant.

As a compatibilist, McKenna is interested in diffusing incompatibilist objections to determinism. He writes that the incompatibilist holds that, "determinism rules out free will and moral responsibility by ruling out an agent's access to alternative possibilities," (2003, 201) the assumption being that, in a determinist network, all effects are directly the result of their causes, going all the way back to the origin of the deterministic universe, at which point all fate was sealed. In rebuttal, McKenna is determined to show that "if free will and moral responsibility do not require alternative possibilities, then this incompatibilist argument is otiose" (2003, 201).

McKenna has designed the "Brain Malfunction" example of Casper the philosophy professor as a demonstration of free will without alternative possibilities at the moment of choice. Casper comes across a state of the art 'Make-it-the-Case-Device'. The machine has a large screen, on which flashes a genie who explains:

> Casper, just beneath this screen are two buttons, one marked 'The Morally Good Thing To Do' and another marked 'The Morally Bad Thing To Do'. Let us abbreviate them as 'Good' and 'Bad' respectively. If you press the Bad button you will immediately make it the case that one million dollars are deposited into your

> bank account. The money will be drawn, in one dollar increments, from the savings accounts of one million college professors. The transaction will be untraceable. If you press the Good button you will immediately make it the case that an entire village of people in the Amazon is cured of an otherwise fatal disease. Saving the villagers will not involve any money and by doing so you will not be stealing from your peers. You cannot select both buttons and this opportunity will not present itself again. You have ten seconds to select your option (2003, 209-210).

I will directly quote the remainder of the example because the precise language is important:

> A timer appears on the screen and begins to count down from ten.
>
> Casper pauses to consider these two options, quickly assessing the import of each. He considers the article he read in last Sunday's New York Times on the villagers' plight. He is fully aware of the urgency of their condition. He also considers his sparse salary as a philosophy professor and he squirms at the thought of stealing from his peers. Imagining that shiny, red Mercedes convertible roadster in the window, as the counter ticks away from 3 seconds to 2 greedily he takes the plunge and presses the Bad button. 'Ah, dinner out tonight!' Casper thinks to himself.
>
> As it turns out, although Casper was unaware of this difficulty, and although there is no reason Casper should have been aware of this difficulty, at the time at which Casper greedily decided to press the Bad button, Casper had a small lesion on his brain that blocked the neural pathway constitutive of (or correlated with) a decision to push the Good button during that ten second interval. Casper could not have decided to press the Good button (2003, 210).

This example seeks to highlight the idea that a morally significant decision (pressing the bad button) can present without alternative possibilities (like being unable to press the good button because a lesion prevented that decision) and still be under the control of the agent, because he used reason to consider the options and decide, before lesioning occurred, that the desire for

money is most significant. His relief after the decision ('Ah, dinner out tonight'), in addition to the envy just before (shiny, red mercedes convertible roadster) indicates that he conducted reasonable deliberation that voluntarily and intentionally resulted in the choice to press the Bad button. Thus, even if he didn't have the alternative to press the Good button at the moment of choice, he still had the 'control' to make the choice he did, and it was indeed morally significant.

The thrust of the example lies in the fact that all moments leading up to the decision were unaffected by the lesion. Only at the precise moment when the 'final act of deciding' occurred, I gather from the example and the theory that follows, did the lesion appear, blocking the alternative pathway. McKenna's opposition to the need for alternative possibilities is unique in that it does not try to block the need for all decisions, only those that are morally significant. And nor does it block all components of deliberation. In contrast, the many defenses of Frankfurt's proposal (that APs are unnecessary) have almost exclusively "attempted to pollute all actional pathways other than the actual one in which the agent acts" (2003, 205). From this I understand that McKenna distinguishes between steps, or actional pathways, in the decision making process, such as (1) when Casper deliberates on his income and the fancy car outside, and (2) the step of the exact decision, when he eschews all else and presses the button. All the preliminary moments (deliberations, etc.) can obey the need for APs, because they are not morally significant. On the other hand, during the moment of moral significance, the decision itself, the need for APs shatters and determinism displays a route for free will to exist.

In his words, McKenna claims that a defense of Frankfurt's philosophy that APs are unnecessary should "attempt to close off all morally significant alternatives without attempting to pollute all alternative actional pathways within an agent's control" (2003, 206). In order to pursue such a theory, McKenna insists that a "Frankfurt-defender" should distinguish between "some class of actional pathways comprising morally insignificant alternatives," which, he writes are those that "could not aid in grounding the judgment that an agent in a Frankfurt example is morally responsible for what she does" (2003, 206). In the case of Casper, the alternatives to pressing the bad button were not considered "deliberatively significant" and are therefore

irrelevant to the moral decision. Only those deliberations prior to the decision are relevant. Thus, the existence of alternative possibilities in the wake of these deliberations, is unnecessary, because they, alone, determine what the decision will be. Due to the lack of indeterminism in between the deliberation and the decision, the relevance of any alternatives is negated because there is a direct causal link, and therefore a reliable contingency of the final decision on the deliberation. McKenna's theory is called the Limited Blockage Strategy, because it takes aim at blocking APs for only the morally relevant step while displaying a determinist framework for a decision resulting in moral responsibility.

McKenna's Limited Blockage Strategy refutation of theories requiring APs, or theories opposing Frankfurt's assertion that all APs are irrelevant, credits the decision, and not the prerequisites to the decision (intentions, deliberations) with moral importance over all else. The goal of McKenna's position is to block the need for alternative possibilities for the parts of decision-making that are morally significant, thereby dismantling the indeterminist need for APs during moral decision-making. As such, he holds that it doesn't matter what deliberation leads up to the decision, only the decision and the proceeding action are morally significant, and since he has shown that APs are unnecessary for the decision to be freely willed, his argument appears successful.

Some deliberations will not contribute to the decision, according to McKenna, and therefore do not represent morally salient material. This is tantamount to the idea that we can't be held responsible for thought, only action; it doesn't matter if we hate someone, we cannot be held responsible for that alone. Only if we commit hateful acts are we morally responsible. If this sounds too strong (what about hateful speech?), I would hold, and I'm sure McKenna would agree, that speech and expression of thought counts as action, and therefore as something that behooves punishment if it is a moral wrong. This theory neatly aligns with our legal framework for responsibility in which we can not be punished for our thoughts if we keep them to ourselves.

I object to the idea that the brain lesion does not inhibit free will, and in my objection I

appeal to the goal of this legal framework with which McKenna seems to be aligning. The reason why thought isn't punishable is that thought is subject to change, immediate and unpredictable change, and does not perfectly predict action. Just because someone seems to think in one way, does not mean they will act that way as well. This grants people a kind of second chance. If they hate someone and have every intention to hurt them, a truly morally reprehensible person might go through with their intentions, having no reservations, only malice, whereas a person who may be good at heart will have reservations that would lead them to choose not to hurt another. The person could have been confused or in a state of vulnerability when they reasoned that the subject of their hate deserves to be hurt. The assumption our legal framework makes is that anyone could 'come to their senses' at the final moment and execute a choice that had little to no visible prerequisite beliefs (should I cite something here?). Thus, the alternative option must remain a possibility until the cessation of cognitive processes involved in precipitating the action, and until all actional pathways are completed.

McKenna may refute this, claiming that any act of 'coming to our senses' would involve additional actional pathways that are as much morally insignificant as those prior. If this objection were to be made, I would suggest that our understanding of the actual stages of decision-making, as far as neuroscience has elucidated, are incomplete. However, what we do know is that decision-making includes a great deal of complex recurrent networks, monitoring, and feedback mechanisms. To isolate a precise moment of 'decision' from the neural activity before it disregards the complexity of this network, which also includes activity in the wake of a decision that continues to regulate the subsequent action. There is no point at which a decision is made and all relevant cognition ceases, waiting for the action to be carried through and to receive new information and start over for the next action. Deliberation is not a strictly pre-decision phenomenon, it continues to influence the nature of a decision and how we execute the appropriate actions after a 'decision' has been made.

If McKenna considers deliberation to be morally irrelevant, then what distinguishes deliberation from decision? Willed action from reflex? It seems like the causal involvement of

consciousness in the decision making process is pivotal for an action to be considered morally salient. At the very least, conscious intent is pivotal for legal ramifications.

It would be foolish of me to claim that all thought has moral significance. It seems clear to me that there is a clear line between ‘idle thought’, in which our mind wanders and experiments in ways we have no intention of bringing about, and ‘intentional thought’, in which we deliberate on our precise intentions in a way that influences our subsequent action. I propose that intentional conscious thought is morally significant, and under control of the agent, and required for a freely willed decision. These two kinds of thought are distinguished by the presence of intention.

1. Thus, my objections to McKenna can be summarized as follows:
2. He inaccurately isolates a decision from its intentional deliberation.
3. He claims that APs are unnecessary for free will, but allows them for the intentional deliberation leading up to the decision.
4. He claims that the lack of APs is not tantamount to predetermination.
5. He does not discuss the need to generate the decision consciously, or discuss why deliberation isn’t considered more relevant to the decision.

Chapter 2
Self-Generating Motivation

You are who you choose to be.

~Hogarth Hughes, from "The Iron Giant"

A discussion of free will involves a great number of words that are used in everyday speech (e.g. intention, reason, and choice), as well as some large, but familiar concepts (e.g. responsibility, morality, and freedom). In order to broach the topic of free will, we must first agree on the meaning of our terms. There is no way to discuss each term and then the next in a linear fashion, so I will have trouble describing this constellation of terms clearly. First I will indulge in a brief etymological analysis, then we can dive into the terminology of practical reasoning, and finally I will summarize my position and the specific, philosophical concerns that it can address.

The import of free will in human life is obvious. Terminology surrounding free will can be traced back to some of the oldest recorded languages, which highlights the fact that the concept of free will has been central to human life and communication for millennia. We can trace the etymological history of volition, a term closely related to free will, from vurita and varam in ancient Sanskrit languages, to Old Slavic veleti, to command, volt in Old Latin, and volo, I wish, in Latin. From volo come the words "volition", "voluntary", and "malevolent/benevolent", to name a few. Volition describes the sense in which we have a wish in mind, and we order ourselves to perform it.

A related word is "will", whose etymological roots return to the Old Slavic veleti and go down the path of the Gothic wili and wiljan, to wish. From wiljan comes the German wollen,

Old English willa, Middle English wille, and our modern day word, will, whose definition I will compare directly to that of volition in the following section. Will, alone, doesn't suggest that I don't just have a wish or a desire to see something happen, will implies the wherewithal to bring something about for a reason. With the invocation of reason comes the invocation of practical reasoning.

Free Will and Practical Reasoning

Practical Reasoning is a field of philosophy that parses out how we think about things. As the name implies, the human ability to reason is the center of exploration. Practical reason is the human ability to apply reason through reflection and deliberation in order to decide in what way to act. While this pursuit seems to directly invoke the question of free will, it is interesting to note that the more metaphysical theories of free will do not strike philosophers of practical reasoning as a primary concern. This is because 'practical' reason is about how reasoning is applied and the consequences of its application, not how or whether reason can originate freely. It is assumed that our reason brings about our action through deliberation, and that our reason is either rational or irrational. Whether our rationality or irrationality is freely willed or not is not up to debate between philosophers in this field; however, I believe that practical reasoning has a great deal to offer questions of free will because of the fact that it addresses the topic of what comes next (of what we do with our free will) without being stymied by debate over the fundamental question of whether or not we have free will. I also find that some terminology and questions that are used in practical reasoning are more readily translatable into empirical study, which bridges the mostly abstract conversation about free will and the observation based process of science.

When a reason is sufficient enough to precipitate an action, it is a normative motivation. A reason is normative if it doesn't just explain why I am doing something (ie "I'm escaping through this trap door because the police are hot on my heels), but can verify what is moving me

to act, which tells us what about the action we endorse as right (ie “Having the police hot on my heels is what motivates me to snatch what little remains of this burned manuscript and escape through this trap door”). In the field of practical reasoning, we endorse things because we value them. We value one thing over another based on our identity. We identify ourselves as a certain kind of person, which behooves us to hold certain things of value, to endorse them as genuine motivators, as normative, and thereby act on them. Ultimately, motivation comes down to our identity, and how it compels us to uphold the values it entails. Reasoning is the process by which we measure reasons against the values of our identity to decide which reasons we endorse as normative, as genuinely motivating us to act.

Christine Korsgaard has written a great deal in the field of practical reasoning on motivation, identity, and morality. In “The Sources of Normativity,” she discusses the basis of ethics in human life as well as how, and why, we should and do obey ethical laws. She writes that “the reflective structure of human consciousness requires that you identify yourself with some law or principle, which will govern your choices. It requires you to be a law to yourself. And that is the source of normativity” (103-104).

The human ability to be self-aware is universally recognized (I can recognize that I am writing this, and you are aware that you are reading it, unless one of us has a severe neuropathological disorder, which I hope and assume is not the case). We are even aware, frequently, of the precise way in which the stages of deliberation proceed (I need to prevent the police from reading these papers, but I don’t have time to burn them with this lighter one at a time, but being caught with them is worse than being caught without them... maybe I have time to cross the room and toss them all in the dying embers of the fireplace, etc., etc.... until we decide we don’t have enough time to deliberate on the best route and simply toss them in the fire, making sure to pull the rug back over the trapdoor as we steal away). Science has yet to elaborate the physical mechanisms behind this capacity of self-awareness. Whatever the mechanism, given our reflective psychology as we understand it now, we have the ability to distance ourselves from ourselves, to segregate between our thoughts/deliberations and our choices.

Philosophers of practical reasoning would call this a segregation between a thinking self and an active self. This is not to suggest that we are split between two selves, rather, these are two sides of the same self. I find that this segregation splits our self roughly in the same way that I do later in this chapter (when I address Alternative Possibilities): between (1) the mechanisms of deliberation where I propose indeterminism is amplified, and (2) the mechanism when the "final" choice is made and determinism regains control over the causal network in order to ensure a straightforward causal link between the prior deliberations/intentions and the action. As such, I will proceed to use the terminology of thinking and acting selves because it simplifies the conversation.

According to Korsgaard, our active selves must unquestioningly submit to the legislation of our thinking selves. She writes that, "we are in a position to require things of ourselves... we have legislative authority over ourselves" (151). This means that your reflective nature dictates your choices according to the most legitimate authority you have: yourself. Korsgaard's argument invokes a Voluntarist position, which suggests that motivation flows through a pathway of obligation, which is assigned by a legitimate authority figure (e.g. God, the government, parents), who has the power to enforce the law. In contrast to the Voluntarist position, Korsgaard claims that being normatively obligated by a legitimate authority is considered free will because the only legitimate authority is yourself, not God, not the government, not even your parents (as hard as the notion is to shake). Of course, we can acknowledge that the laws of a government, or the principles of your parents' household are just and right, but Korsgaard would hold that we are obligating ourselves to uphold that law, and only we can do that. If you obligate yourself, you are appealing to your autonomy, which governs your choices freely. The capacity to appeal to your own autonomy through reflection with your thinking self is what Korsgaard means by being a law to yourself. Being a reflective human makes you follow a law or principle that is dictated by yourself.

Our practical identity dictates what laws and principles we will follow. The way you advocate for a law, under which to rule yourself, is by valuing one over another, a philosophical

position that almost assumes Alternative Possibilities. The way you, as a reflective human, advocate for "the principle or law by which you determine your actions" is by assessing which one is the law that "you regard as being expressive of yourself" (100). In this way, you value a certain law that is expressive of your identity. Your identity is the way you define yourself, and your practical identity includes the requisite principles you must follow in order to have assumed a given identity. You choose what practical identity you associate with, and follow the principles that such an identity is bound to follow. Being a rational agent means that you follow a law that you have valued by defining yourself as someone who is disposed to follow such a law.

A rational agent is one who is disposed to make rational, law abiding (personal law, not necessarily communal law abiding), decisions based in reason. As Korsgaard put it, "because the will is a causality, it must act according to some law or other... since the will is practical reason, it cannot be conceived as acting and choosing for no reason" (97-98). Thus, being normatively obligated to make a choice by yourself is a force of causality that follows a prescribed set of laws. I contend that free will is both a causality and, at some juncture, an independent source of indeterminism, which enables our freedom up to the point when we select a practical identity and thus generate a reason for our actions ourselves. After this point our rational nature kicks in and enables our will to be a causality in the way that philosophers of practical reasoning describe.

The importance of identity in practical reasoning is part of what makes the use of philosophies like Korsgaard's so advantageous for my position on free will. For me, and this will become clear in the next section where I lay out my position, the importance of identity in calibrating decision-making during complex problem solving demonstrates the importance of a conscious mechanism for deliberating on our choices. Further, it establishes how our unconscious decisions can fall under the umbrella of free willed choice when it aligns with values we endorse as normative for our practical identity. Metaphysical positions on free will, such as those described in Chapter 1, rarely focus on the importance of identity. Wegner doesn't explicitly address identity in the wake of his philosophy, but he does hint at it through his response to peer reviews on the topic of responsibility. For Daniel Dennett, whom I will address

in Chapter 3, an illusionist position such as Wegner's doesn't necessarily exclude the capacity for free will, and he reconciles it by holding that we have a personal narrative that acts as a practical identity that is continuously shaped by our brain for our consciousness to observe.

Korsgaard makes the strong claim that free will and Immanuel Kant's categorical imperative exist by necessity, and that this is not really up to debate. As she writes, "occasionally one meets the objection that the freedom that we discover in reflection is a delusion. Human actions are causally determined. The philosopher's bugbear, the Scientific World View, threatens once more to deprive us of something we value" (94). However, she concludes that "Determinism is no threat to freedom" (95), though she also notes that this "doesn't mean that I am claiming that our experience of our freedom is scientifically inexplicable" (96). According to Korsgaard, freedom is ensured by the categorical imperative, which "tells us to act only on a maxim which we could will to be law" (98) The only problem that freedom entails is that "the will must have a law, but because the will is free, it must be its own law" (emphasis added), which Korsgaard claims isn't an issue because "the categorical imperative merely tells us to choose a law. The only constraint on our choice is that it has the form of a law. Nothing determines what that law must be. All that it has to be is a law" (98). This position assumes that our choice of law is not determined, a common assumption in practical reasoning. While assumptions can be dangerous things to make, every field must make some, and I don't find this one to be a problem. I find that we do have non-determined choice, which I will argue in the remaining sections of this chapter. Furthermore, as I will show in Chapter 3, the 'scientific world view' may agree with this philosophical premise.

A New Model of Free Will

I propose a two stage model of free will during decision-making that can be understood both through the lens of neuroscience and practical reasoning. There are two central contentions in my model: determinism vs indeterminism, and conscious vs unconscious. These do not

necessarily map onto each other on a one to one basis. In my model, indeterminism is always a factor. In biological terms, indeterminism constitutes the random variations of neural activity that create noise. Neural noise is the random fluctuation of electrical signals that pass through a neural network that can be instigated by the effects of quantum uncertainty. Quantum uncertainty, which will be explained further in Chapter 3, enables the random variation of position and velocity of particles that make up chemicals and proteins in our brain. This low level variation enables slight differences in electrical conductivity of neurons in a circuit. Fluctuations in one cell can affect the fluctuations in others, which ultimately creates a system of noise. More and more contemporary studies are showing that neural noise is what enables plasticity, changes to the way a cell in the brain reacts to stimulation, which underlies learning, memory, and other adaptations that are vital to life.

That's indeterminism in the brain. We can think of determinism as a balance for the random, indeterministic noise in the brain. While cells are capable of fluctuating in specific properties, other characteristics remain deterministic. For example, neurons fire on an all-or-nothing basis. They have a specific threshold of activation that they must receive before an action potential (more on that in Chapter 3) is initiated, which passed the stimulation signal to subsequent neurons. The nature of a cell, wherein it follows laws of electrochemical physics, that follows these strict rules of firing upon reaching the activation threshold, is an example of determinism. If you know the thresholds and the activation, you can predict the result because it follows a deterministic pathway.

That was an example of determinism in the brain. These two can take place at higher levels than just a single cell. Neural noise, indeterminism, can take place at a network level if cells are active enough to transmit enough signal as to cause fluctuations in subsequent cells, thereby amplifying variation in network activity. On the other hand, strict patterns of activity, determinism, can also take place at a network level if the indeterministic fluctuations are not amplified.

Thus, we have our mechanisms of indeterminism and determinism, which will be instantiated and supported in Chapter 3. Now, we must broach the second contention: conscious vs unconscious involvement in free will. I submit that both have a role, though they may be unbalanced in terms of quality and quantity of contribution. Consciousness, I propose, can direct attention in the brain to amplify specific neural noise so as to focus on specific variations of activity. Through directing attention to specific networks and providing feedback to neurons within the network telling them to follow specific activity patterns, our conscious attention can sculpt neural noise into novel neural pathways and activity patterns. What I propose this does is change the criteria, or rules, that neurons in these networks follow to respond to specific kinds of stimulation in the future. This can be milliseconds, hours, or even years later.

The unconscious, on the other hand, does not consult higher level cognitive mechanisms, such as attention, when directing neural activity. When a stimulus is given to an unconscious framework, the resulting activity obeys deterministic rules that define how the neural network will respond, what premotor areas it will recruit, and how a response will be precipitated. This is a fast and easy way to solve simpler problems. Most animals rely entirely on unconscious responses; they are not confronted by enough complex problems for them to have evolved a problem-solving capacity through higher level cognitive mechanisms, they simply need to run when chased, and feed when hungry. Of course, this is not to say that unconscious systems are simple. They are incredibly complex, mostly because they are in charge of most of our lives. We are understanding more and more that unconscious systems don't just control our day to day breathing, heart beat, the precise motions of my fingers to maneuver a screwdriver, but also things like judgement through implicit bias.

How do these two contentions map onto each other? Unconscious networks rely entirely on deterministic neural networks. They do not take the time to consult higher level cognitive mechanisms and amplify indeterministic neural noise, they just proceed through the path of least resistance, as in, the simplest, fastest way to dodge the attack of a predator from the bushes, or a broken limb falling from above. This is not to say there is no indeterminism in unconscious

activity, only that the indeterminism isn't amplified, and ultimately doesn't end up contributing an appreciable amount to the overall activity. Most determinists who acknowledge the success of quantum mechanics in describing indeterministic quantum activity in the universe maintain that such fluctuations occur at such a low level, quantum level particles, as to be subsumed by the patterns of activity that emerge at a higher level from averaged fluctuations. I hold that consciousness, on the other hand, applies higher level cognitive mechanisms to amplify indeterminism in order to sculpt novel network activity patterns and connectivity. But it isn't so straightforward. When a conscious mechanism changes a neural network, activity through that system can progress deterministically such that the changes to cognitive processing map directly onto changes in how this processing recruits premotor areas, thereby precipitating a new response that reflects changes in cognitive processing.

This process can also be understood in the more intuitive terms of practical reasoning. First, we consciously generate a series of motivations (directing attention toward several variations of neural noise), which are reasons to be or act a certain way. Through deliberation, we weigh the normative quality of these motivations to ascertain which best fits our practical identity and which is the most genuinely moving to us. Once we have selected a normative motivation, this acts as an unbreakable law for us, as we must uphold a law that we hold ourselves to (if we are genuinely motivated by this reason, then there is no way for us not to uphold it, because we are rational we must be a law unto ourselves), and we act on the basis of that reason.

So, where does free will fit into this picture? Free will, as I have defined, is the ability to consciously self-generate a motivation to act. In being able to self-generate motivations, we can act in a way that is free from predetermination by our environment, upbringing, and biological makeup. Of course, we can not fly because we will it, and it is true that these factors have a large and important impact on our deliberations, but we still have the ability to freely generate novel motivations, which can provide the freedom to choose how to act.

Further, I propose that some unconscious decisions proceed through networks of neural

activity and connectivity that have been shaped by our conscious deliberations. As a result, once conscious deliberations have been made, this establishes what laws we hold ourselves to, and many unconscious decisions are made using that framework. Let's say you dig a ditch so water flows harmlessly alongside the road instead of covering the street in a deep puddle. You can take credit both for having dug the ditch to make the first trickles of water flow away from the road, but you can also take credit for the continued function of that system you put in place a year later. Every once in awhile you may need to fish out a damn of leaves, and thus you can feel even more instrumental in the continued function of your beautiful ditch. Similarly, if we decide that we absolutely refuse to harm another individual, and then someone surprises you so that you unconsciously evade them instead of striking back, you can be proud to have shaped your unconscious orientation away from violent responses toward a more passive self defense. I will revisit this idea of the unconscious as an extension of the conscious, and therefore part of the framework of free will and responsibility, at the end of this chapter and in Chapter 3.

I have four central concerns about those positions that are currently held about free will coming out of Chapter 1:

1. What is the nature of indeterminism in the brain? What is it and how does it influence large scale cognitive activities?
2. Can alternative possibilities be obtained through indeterminism? What does it mean to have been able to do otherwise at a cognitive level?
3. Can we control this indeterminism in a way that does not go through deterministic pathways, or merely generate random noise?
4. Can our conscious mind control indeterminism in a way that establishes a middle ground between determinism and randomness?

Next, in the remainder of this chapter, I will focus on three key necessities of my philosophical position in response to those concerns: the need for alternative possibilities, control over indeterminism, and that this control be conscious. These are mainly in response to McKenna,

Kane, and Wegner, respectively. My theory of free will responds to these concerns and thereby seeks to establish the following requirements for free will, and defend their verity:

1. Alternative possibilities are available through indeterminism.
2. I am the origin of my choices in a way that is neither deterministic nor random.
3. The reasons, or motivations, that guide these choices are generated by my consciousness.

In Chapter 3, when I address the scientific background of my proposal, I will return to the nature and role of indeterminism in order to flush out the plausibility of my philosophical claim.

Alternative Possibilities

McKenna and Frankfurt's compatibilist attempts to block the need for alternative possibilities is convincing; however, I propose that their attempts do not properly account for free will. As such, their compatibilist position fails to defend free will in the face of determinism. The failure of Frankfurt's proposal lies in the fact that it does not allow for alternative possibilities for an agent's intention at all. According to this proposal, intentions are the direct and only result of the series of external stimulations, or experiences, to which the agent has been exposed prior to decision-making. The external stimuli are not under the control of the agent. The agent merely encounters them while navigating space. One can argue that the choice of where to navigate is under the distinct control of the agent, that we choose our environments (e.g. whether we listen to gypsy jazz or death metal, make friends with athletes or poets, eat nutritiously or not, etc.). Ostensibly, having this control over the many aspects of our environment would translate to control over that which stimulates us, and, thus, that which influences our intentions and therefore our actions.

I agree with this argument, our free willed choice of environment does influence our development and identity. However, this scenario cannot be supported by a determinist stance. It assumes that we are always in control of our environment, since birth (or maybe even

conception) and at every point after. If it is not the case that newborn infants have control over what they are and are not exposed to, then it cannot be said that their subsequent navigation is freely willed. If an action corresponds directly to an antecedent environment, then the agent must have control over their original environment in order to have control over anything. Even given control over this original decision, determinism implies that all subsequent experience and action would be the predetermined effects of this original antecedent cause, which would eliminate control from the hands of the agent the instant after they've made their first free willed choice. Furthermore, if we don't assume that an agent can have this miraculous moment of free will at the very beginning of their life, then the predetermining choice is made by, say, the parent. As such, all subsequent experience and action is the sole cause of the parent, which is the sole cause of the grandparent, and so on until the origin of the universe.

For any given action, if the agent does not have control over the antecedent environmental stimuli that precipitates their thoughts and intentions, which it would do in a deterministic universe, then the agent does not have control over the subsequent action. Further one cannot hold that an agent has control over the antecedent environment, because the actions the agent underwent to navigate space into this precise environment were actions that the agent did not have control over according to the prior argument. The the only way one can assert, through an avenue without alternative possibilities, that agent has control over their environment is if you also assert that the origin of this cycle of environment shaping action and action shaping the environment, the moment of birth, the newborn agent assumes absolute control over their environment. This is an extremely difficult position to defend, especially given our knowledge about the environment of the womb influencing development of the child prior to birth. Thus, it becomes unclear where to draw the line for the moment when an agent becomes an agent and obtains control over their primary, germinal environment and can thereby be said to have control over what influences them to intend to act a certain way. This is what is required in order to defend this compatibilist position

McKenna accounts for this opposition to Frankfurt. He allows that deliberation prior

to a choice can include non-deterministic alternative possibilities, which breaks this chain of deterministic cause and effect that removes control from the agent; however, he differentiates between the period of deliberation and the period of choice-making in order to delineate between periods when alternative possibilities are needed and when alternative possibilities are not needed.

It is important to note that I agree in a theory of mental causation wherein our intentions result in our actions. I hold that our mental considerations directly control our actions, at least in non-pathological cases. As such, for a non-pathological individual, our intention to act can be said to determine our actions. The concern I have with Frankfurt relates to the manner in which our intentions our formed. Given a scenario without alternative possibilities, compatibilism suggests that our intentions are determined by our environment, which perpetuates the aforementioned cycle. However, given indeterminism, we can assert that the environment does not directly determine our intentions. I do not suggest that indeterminism results in a lack of determinism between our intentions and our actions, much like McKenna, but solely between our environment and our intentions. I suggest that our brain utilizes indeterminism in such a way as to create this possibility. Thus, I propose that our free will is our consciousness wielding our brain as a tool to amplify indeterminism in order to cause the generation of undetermined intentions, which I call self-generated motivations. By necessity, the self-generation of motivations requires alternative possibilities, or else there would be nothing undetermined to generate. As I will discuss further in Chapter 3, indeterminism gives plausible support for such alternative possibilities.

As we saw in McKenna's example of Jasper the professor, Jasper was fully capable of deliberating on the multiple choices he had available such as to accept the money, help the villagers, or to back away from the situation entirely. However, when it came time for Jasper to make his decision, a miraculous lesion appeared in his brain that prevented him from making the choice to help the villagers. McKenna blocked the alternative possibilities that could have been available to Jasper at the time of decision. If it is the case that Jasper could not have chosen

otherwise as a result of this lesion, I hold that this was not a freely willed decision. As I have said, my central objection to McKenna's proposal lies in his attempt to cleanly differentiate between a period of deliberation and a period of choice in a temporal dimension. I believe that this does not account for the complexity of decision-making as it occurs in the brain. Deliberation is very closely entwined with the actual decision. I would suggest that there was still the possibility that Jasper could have chosen otherwise than he did, even though it seemed clear that his intention was to press the bad button. It has been found that deliberation continues to occur even after the pre-motor areas of the brain are engaged and begin to precipitate the motor response, which causes an action to be taken. These further deliberations act a as a variety of regulatory mechanisms.

Given the indeterminism he admits into the model for deliberation, it's entirely plausible that, somewhere in the deterministic chain of events that McKenna describes, the insertion of an indeterminate intention, during or even after the action is being precipitated, could have provided Jasper with the opportunity to have chosen otherwise. As a result, it could not be assumed that he would necessarily follow through with pressing the bad button. By blocking this avenue at what McKenna presumes was a safe time, he actually did prevent Jasper from having control over which choice he made, which removes free will from the scenario.

This argument relies on disagreeing with how precisely Mckenna blocks Jasper's alternative possibilities. Since my model off free will also includes a compatibilistic level of determinism in the connection between a choice and the action it precipitates, it behooves me to establish precisely why blocking alternative possibilities with determinism in one way allows for free will while the other, McKenna's proposal, does not. I propose that indeterminism is present during conscious decision-making. I would suggest that during unconscious decision-making the brain does follow a strictly deterministic pathway, which would accommodate McKenna's block of alternative possibilities. If Jasper had made the decision unconsciously in this scenario, then I would agree that blocking the alternative route (pressing the good button) would not have made a difference in the causal chain of events and would therefore have not inflicted

on his authorship of the choice he made. This is a strong proposal that segregates between the fundamental mechanics of consciousness and unconsciousness. Later, when I address the fourth concern (i.e. the conscious generation of choice), I will reflect on the nature of the difference between conscious decision-making and unconscious decision-making, especially as it relates to the concept of responsibility.

For now however, I propose that the failure of McKenna's system is in his inaccurate segregation between the process of deliberation generating a choice, which must be indeterministic, we both agree, and the process of a choice generating an action, which must be deterministic, we both agree. He segregates the two according to their temporal qualities. According to McKenna, if we can just wait until deliberation generates a choice and ceases, then we can block alternative choices and avoid affecting the deterministic stage of the process. Since we cannot segregate based on time, we should segregate based on what parts of the process are conscious. If Jasper was still consciously deliberating even while he pressed the button, which I expect he would, then it would have been wrong to block alternative possibilities then. Ultimately, instead of dividing up these processes temporally, like if one occurs, and then the other in a linear order, we should divide the particular neural pathways involved in either conscious deliberation or unconscious pre-motor preparation. If we block alternative possibilities through highly localized lesions in the latter pathways, then I would agree that we have agreed on a place in decision-making when alternative possibilities can be blocked without disrupting free will.

Controlling Indeterminism

At this point I will be transitioning into the use of concepts and terminology from the field of practical reasoning. As an introduction, I refer back to our earlier discussion of volition. Volition is the capacity to make a choice after consideration and deliberation. This alone does not seem to describe free will as I define it. It could be argued that the cognitive mechanisms

of all voluntary deliberation are hardwired by our biological makeup, but if the process of all 'voluntary action' is reduced to neural activity that is pre-defined and out of the control of our consciousness, then free will, as I define it, does not exist. I suggest that, in order for free will to stand a chance, we need to have a way of controlling this activity. We must be able to make a choice about what to do, in a way that is not hardwired, which I translate to mean pre-determined. Thus, we must be able to generate a new option, one that was not presented through a strictly causal network. I will sketch how, in essence, this means generating a new motivation to act.

The concept of self-generation is central to my theory. Philosophers such as Belageur and Kane have engaged in conversation about such an idea, which they've called torn decisions or self-forming actions. Philosophers debating this component of free will look for an example of an instance wherein, given two or more equally compelling options, each with their own calculated weight and normative value, we can make a choice to go in one direction or another. Libertarians suggest that this choice can be made nondeterministically, which is to say that given a set of options an agent can choose one without the decision being determined, or reducible to a hardwired mechanism. One of the strongest oppositions to this theory of free will, in which we have control over nondeterministic choices, is that such choices would be based on luck or chance. One of the largest hurdles that a libertarian perspective must overcome is how an agent can have control over indeterminism so that it does not manifest as purely random noise. According to Kane's definition of free will, with which I agree, an agent must have control over their decision in a way that is not determined and not random. Thus, free will requires indeterminism in order to break free of determinism, but also requires that indeterminism yields to some force possessed by an agent in order for it not to be random or, as Balaguer and Kane say, in order to have appropriate non-randomness.

In response to these requirements I suggest a theory free will that invokes novelty. An agent must be inventive; an agent must be capable of being handed multiple choices and inventing a novel response. But how does novelty come about? Is it possible to merely invent

something out of thin air? The way that nature arrives that novel events is through random mixtures, the worst of which are purged through selection pressure and the best of which are promoted. For example, the evolution the new and improved faculties of an organism come about through millions of years of genetic variability. Each genetic variant has the opportunity to test its advantages or disadvantages against the trials of the environment. Through selection pressure, those genetic variations that result in the greatest advantage survive and reproduce. Is this how novelty comes about in the brain? Does a brain, in rapid time, generate a series of random possibilities so that the agent can merely select which one is best? In this way the brain would be a tool that develops random variations of putative neural activity in response to given inputs. In order for an agent to have free will, they would then choose which pattern of neural activity was best, or most appropriate according to the agent's moral and practical values. A key question is how an agent can have control over these random neural variations, and thus over indeterminism, in such a way as to select a "best" option and allow it to precipitate a physical response in the body.

What I have said so far doesn't answer this question because the agent can seemingly only make this choice in a deterministic or random way. It is clear enough to see how an agent could make this decision through determinism. Given an indeterministic set of putative neural responses, the agent would merely invoke a series of rules in order to calculate the optimal choice. Ostensibly, these rules would act like selection pressure driving the most advantageous output to make the cut and precipitate action. According to a determinist, these rules would start as instinctive programs hardwired into the brain, and would be honed over time through trial and error. In making mistakes on some occasions, using some rules, and being successful on others, using other rules, an agent's brain would, over time, hone their preferences and beliefs into what appears to be a very personal and unique approach to every problem.

However, the determinist argues, in actuality a complex computer could perform the same function given the same predetermined rules at the outset and by following through with complex calculations and adjustment of rules for these calculations in response to the results.

This is called a feedback loop when feedback, the success or mistakes of an action, is provided to the system and the system uses a simple, unwavering, hardwired regulatory program to adjust the central calculating procedures in order to generate a different, presumably better, response on the next try. An example of a feedback loop can be seen when kicking a soccer ball. You intend to hit the ball low to the ground and with a straight, powerful trajectory, but instead it lifts high in the air and flies over the goal. This represents a mistake in the body's movement when kicking, and thus a fault with the calculations the brain made when planning the strike. Seeing the ball clear the crossbar by 20 feet, the feedback, which tells you that you have to adjust to hit the ball lower, your brain changes the way in which it calculates the next strike. You retrieve the ball, and try again, this time with your brain programming the body to lean over the ball more. You successfully drill the shot straight, low, and with great speed. Observing the success of your correction, another form of feedback, your brain initiates memory of the new program and, through repeated practice, will be able to enact the appropriate program in the future.

The vast complexity of this trial and error, which takes place constantly all day every day, would be the basis for an illusion that an agent's rules are personal and unique, which would confirm their basis in free will. However, if we can reduce the seemingly free willed soccer practice session to a complex feedback loop with prescribed calculations, tolerance levels, corrections, etc occurring in your brain, then we can also trace the origins of every program that was used to the inchoate, hardwired programs you were born with, which were modified to the point they are now by trial and error feedback in the environment. Free will as I define it requires an additional feature of intention, in which an agent has conscious, meta-regulatory control over changes in her programming. Otherwise, it is relatively straightforward, though admittedly not easy, to argue that our perception of ourselves as free willing agents is an illusion.

Therefore, as agents with free will, not only must we be inventive, we must have control over the rules of the game. Instead of being motivated toward one choice or another by rules derived from predetermined instincts honed by reducible feedback loops, we must have the capacity to generate new rules for our mental calculations, which I refer to as new motivations.

I do not define a new motivation as one which has never been created prior, but rather as something of which we are the sole origin. Of course, even a free willing agent cannot decide to pursue a new motivation that defies the laws of physics; we cannot simply fly because we have the motivation to do so. By novel motivation I do not even suggest that the eventual choice can be distinguished from what it would have been had we pursued an automatic instinct. The generation of a novel motivation is merely the amendment of the current rules dictating what we are motivated to do (how calculations should be made), and why we are motivated to do it (how these calculations can be applied to present and other decisions). This is to say that the generation of a novel motivation involves a change in our intentions as well as that which we find normative so that we may be compelled to take one course of action and not another.

Consider the case of a simple choice that displays the complexity of motivations that can inspire a seemingly simple decision. One morning you are given the choice of orange juice or coffee to drink. Most mornings you have coffee, but on this occasion, for whatever reason, you are confronted with two choices, to maintain your coffee habit, or to try something new. You could have a range of reasons for performing either action. You may take the coffee because you have learned from past experiences that coffee wakes you up more. On the other hand, you may be in the mood to switch things up and try the orange juice. Or you may feel embarrassed, for whatever reason, to make one choice. Or perhaps there is very little coffee left, so you take the juice so others can have coffee if they so prefer. Maybe you even know that a certain person will want coffee, so instead of being impartially altruistic, you're selectively considering one person. Maybe this condition is not entirely selfless, but may curry the favor of this person to your benefit. My point is: the variety of motivations driving any one of the choices is extensive.

While the original components of our motivations can be logically reduced, you will see that, in the end, we have the ability, when making freely willed decisions, to wield indeterminism in a way that invents novel variations of those reducible components as a set of options. This indeterminism breaks the chain of determinism. Suddenly, possibilities enter the picture that could not have been 100% anticipated, even if we knew everything about the agent leading up to

this point. It breaks the need for all decisions to be hardwired and predictable (though they may still end up aligning with what may have been hardwired and predictable). Now we need a way to control the set of randomized options. I suggest that meta-regulatory elements in our brain, which is our consciousness at work, choose the motivation that suits us from this pool.

We must be careful to remember that in free will there is a component of determinism. The precise mechanism of the decision proceeds through a causal mechanism that weighs how much a choice stands up to requirements for a "good decision." The choice that is best according to our requirements, or our principles, is allowed to precipitate action. This is not a facet of free will that I grudgingly admit to, the determinist argument does not win more ground because I admit there are components of determinism. I suggest that determinism grants free will more strength by linking our reasons to our actions. If the process of performing an action based on a reason was not deterministically causal, there would be no reason to claim we have control over our reasons, because we wouldn't be sure that they caused the appropriate action anyway. Indeed, as Kant writes, the will is a causality, so "freedom is by no means lawless" (65, for beck). Korsgaard provides some commentary on this subject as well, pointing out both the causal and the non-causal elements of free will: "Since the will is practical reason, it cannot be conceived as acting and choosing for no reason. Since reasons are derived from principles, free will must have a principle. But because the will is free, no law or principle can be imposed upon it from the outside. Kant concludes that the will must be autonomous: that is, it must have its own law or principle" (93-94). Thus, I claim that while the calculations themselves are causal, thereby linking our reasons to our actions, the way in which we calculate is through principles, what I call rules, that are our own.

Thus, an agent's consciousness uses meta-level regulatory mechanisms in the brain to change the rules of the game, to alter why we are compelled one way or another, to alter what it is that motivates us in that moment as an agent and not as a machine that merely responds in the way it should given feedback-honed hardwiring. By consciously changing the rules we change the way our brain calculates weight, and therefore influences the decision-making pathway

through our free will.

Conscious Control

As I have said, consciousness controls the way our brain calculates the value of things we consider when making a choice. In other words, our consciousness dictates what motivates us, and these motivations inform our brain to assign more or less weight to specific considerations we have in order to sway our decision-making toward the choice that represents our normativity, or that which compels us. Importantly, I believe that the phenomenology of consciousness, the experience of it (not the mechanisms themselves but how they feel to us), is vital to this process. In Chapter 3 I will outline how the phenomenology of an agent's consciousness has control over regulatory mechanisms in the brain, which dictates the structure of our decision making, as well as how current neuroscience supports this claim. Here, I will discuss what it means, philosophically, for the phenomenology of consciousness to have this control.

One major implication is in distinguishing between conscious and unconscious free will. If the phenomenology of consciousness is required to control our regulatory mechanisms, then this would suggest that unconscious decision-making is lacking in free will. The vast majority of our decisions are unconscious, which means that the vast majority of our actions are generated unconsciously. This would leave us to suppose that a majority of what we say and do is not the subject of our free will and therefore not something we are morally responsible for saying or doing. I propose that many unconscious decisions are based on the motivations that have already been endorsed by our consciousness. Essentially, our unconscious decisions are an extension of our conscious ones.

"Trying Not To Try" is a recent book by Edward Slingerland in which he documents the concept of wu-wei (oo-way) throughout ancient Chinese philosophy, and relates it to findings in modern cognitive science. Wu-wei is a state of effortless or spontaneous action, which is complex activity we perform without being entirely aware of it. Some people might experience

wu-wei when cooking, for example, when you gracefully combine a series of ingredients, focusing on multiple things cooking at once, and it magically comes together. Perhaps you're a gardener and can relate to those moments when you 'wake up' and realize you've weeded the whole garden, or a writer and you're familiar with those times when you 'wake up' and find that it has grown dark outside and ten pages of beautiful prose have leapt out of your fingers. Those times of thoughtless, effortless activity are called wu-wei. Athletes commonly encounter this state in the intense moments during a match when they stop over-thinking their actions, allow their instincts to kick in, and they perform better than normal.

Over-thinking is a key word in this philosophy. Several schools of ancient chinese thought held that trying too hard, consciously, was counterproductive, and that letting go and trying not to try, so to speak, was the best way to achieve a state of contentment and productivity. Slingerland distinguishes between what he calls hot cognition, which is our unconscious mind, and cold cognition, which is our conscious mind. What is interesting is that wu-wei, a desirable state wherein we acquire a robust charisma and sense of power, or authority, is achieved through the submission of our cold cognition, consciousness, to our hot cognition, unconsciousness. Not only is a majority of our decision-making unconscious, but now submitting to unconscious decision-making seems to be encouraged! How is free will relevant if the our cold, conscious cognition is pushed so far on the back burner?

Well, both the ancients and modern cognitive scientists advocate for a period of training and calibrating. It is widely accepted that we shouldn't just act on our every instinctual desire. A couple of the ancient chinese schools of thought (eg Confucianists and some Daoists) believed that we had to cultivate our hot cognition in some way. Strictly Confucian thought held that a strict regimen over a long period of time would carve our unconscious mind enough for us to leave the majority of our activities to spontaneous action. Less conservative Confucianists following Mencius believed in a more liberal, but definitively cultivating approach, which involved training and care until we could reach a desirable state. More conservative Daoists led by Zhuangzi believed in going with the flow, but after having selected a set of values that we

wish to pursue. What all these philosophies have in common is that they seek to train the hot cognition to act in a desirable way automatically so that our cold cognition, which is slower and less efficient, can take the back seat during wu-wei. Equally so, all of these philosophies imply the need to choose a value-set that we either entrain, or hope to follow. This has to be done by the cold, conscious mind beforehand.

Part of what this book points out is: if it is the case that our unconscious mind can handle matters just fine on its own, then why do we have big, bulky, slow, and inefficient consciousness? This question complements a similar question about why we bothered to evolve a phenomenological aspect of consciousness if it has no direct causal role in our behaviors. Much like what we saw with Wegner, one answer is that phenomenological consciousness is a tool that the brain, not us, not our conscious self, uses to model other minds, and as a metric for how good or bad an action may be. Slingerland discusses the theories in cognitive science that suggest a role of consciousness in modelling other minds in a way our unconscious mind cannot:

> Our unconscious minds are very good at quickly detecting agency, identifying threats in our environment, and reading emotions in faces. Only our conscious processes, however, seem capable of complex modeling of other minds. Consciousness creates a virtual representation of the internal thoughts of others so that we can figure out how to interact with them… (59).

Additionally, Slingerland mentions how "the virtual world of consciousness is also where we get to practice things without actually having to do them… this helps to make the artificial, novel movements involved more fluent and effective" (59).

Slingerland is not a philosopher of free will, nor does he try to make a statement about free will, so I will not be critiquing him. Rather, I want to display a fascinating example of how convergent theory from ancient philosophies to modern cognitive science supports the heretofore uninvestigated importance of the unconscious mind, but still supports the need for our consciousness to act as a referee. It is the conscious mind that is making choices with free will

and modifying the way our unconscious mind approaches decision making. This process, I claim, is tantamount to our unconscious decisions being freely willed as we have a substantial measure of control over modifying them to fit our personal moral beliefs.

It could be argued that I am overestimating the impact of consciousness on our day to day activities. Quite often our intentions seem to elude us, though we still feel that we are responsible for our actions. Many of us will feel culpable for something we did in the throes of passion, like yelling at our closest friend because we're in a bad mood. Why do we still feel responsible? Because a bad mood shouldn't have control over us? Because we could have chosen to overcome that problem and not let it overwhelm us? These rationalizations are based in conscious decisions we've made. We may not explicitly say to ourselves: "I identify as someone who does not allow my emotions over a problem to hurt people who had nothing to do with said problem, therefore I think it is wrong to harangue my friends when I am in a bad mood." Such over-analytical self-awareness is probably just a symptom of being a philosopher, and is not requisite for being a regular human being.

However, even if we do not have the same terminology and cohesive, rational thought, of which that quote would be an exemplar, we still go through these logical steps, and I would say that we do so consciously. At some point prior to yelling at our friend, we have made the conscious decision that that is something we feel is wrong. Maybe we saw someone in a movie behave that way and we thought it was wrong. Maybe someone did the same thing to us and we committed ourselves to never being like that. Deciding that we feel something is right and something else is wrong, or, more simply, that we like one thing and not another, is entirely satisfactory for establishing an identity and endorsing requisite values, and regulating our behavior (conscious and unconscious) through a robust mechanism of free will.

We are still pestered by the findings of cognitive scientists that a majority of our decision making is unconcious. Some figures rise as high as 99% of our activity being mediated through unconscious decisions. For one thing, this is slightly misleading as it includes unconscious

control over such things as heartbeat, breathing, moving our fingers in the precise way it needs in order to type these words, walking in a straight line while also navigating around that person who is moving slowly ahead. The amount of neural activity needed to regulate our day to day activities requires that a great deal of it occur automatically through highly efficient, machine-like processing in our brain. Otherwise, we would be too distracted by menial tasks to be able to execute long term planning and moral decision making, which probably evolved because they have an essential, and irreplaceable impact our livelihood (how much others like and will cooperate with us, how many extra resources we can secure without exerting so much energy, etc.). For decisions that are morally relevant, or related to our long term goals, a much higher percentage is conscious. Not all, but much more. And really, this still makes sense as a balanced structure for conscious and unconscious activity.

Granted the case that only a small percentage of our decisions are conscious, I believe this still leaves quite a bit of room for what we would consider thorough application of free will. Let's say 10% of our decision-making is conscious and therefore capable of being directly controlled by free will. This is like saying that for every hundred days we spend making decisions, ten of them are spent figuring out exactly how we want to make these decisions. The remaining ninety days are spent executing our decisions based on those rules we decided upon. Ten days is a lot of time to decide what we think is right and wrong. It's important to distinguish between deciding what we think is right and wrong in the moment (as a kind of gut reaction, but one which we consciously endorse), and deciding what we think is right and wrong in the most optimal sense. People spend their entire lives being philosophers, lawmakers, and policymakers in order to find the 'most optimal option' for what is right and wrong. However, we don't need a 'most optimal option' in order to have an opinion about what is right and wrong. In the moment, I might think that I am perfectly justified in yelling at my friend simply because I'm in a bad mood. In that moment I might think that expressing my anger is better than not doing so. Given enough time I may decide the opposite, that succumbing to that anger only made the problem worse.

Thus, we can decide right and wrong, in a way that is entirely mediated by our free will, rapidly. Even if only 10% of our actions are mediated by those moments when we first decide how we feel about something (like when we first encounter someone wearing a green suit, or see a parent pushing their child to get them somewhere they have to be), it becomes clear that the number of things we endorse as right or wrong during that 10% can certainly account for all our subsequent actions. It only takes an instant to have a true gut reaction, and probably not much longer to decide on a temporary moral position. It's a lot like riding a bike, we spend a couple weeks learning and then we are able to ride for the rest of our lives. In that sense free will is like muscle memory; we only need a short period of concentration to master a certain level of skill that we can continue to execute thoughtlessly for a long time. This doesn't mean that we will become internationally renowned cyclists in a couple weeks, or revered fonts of moral wisdom in ten days. That kind of mastery takes years and years of practice. On the other hand, humans have an amazing capacity for memory. I suggest that we use this memory to store our conscious decisions in a way that can be accessed later without necessarily requiring our conscious endorsement. Indeed, we don't need our conscious endorsement because we already had it when we endorsed it in the first place.

Chapter 3
The Neuroscience of Free Will

So far I have outlined some positions on free will and made my own claim from the stance of practical reasoning. Here I will show how I find current neuroscience can support my position. Additionally, I will try to highlight how and why my interpretation of contemporary scientific theory differs from the central perspective from leading scientists and philosophers, who most often chock up free will to an illusion.

First, let's review my position. What does free will mean? Free will is the ability to make a choice without external influence. A choice is an option for action that we select from a set of alternatives. The basis of this selection must exist in a decision-making pathway that can exist on its own, which is to say, without being determined by the events, what I've called external influence, leading up to it. By external influence I mean anything that has happened prior to a state that has exclusive, causal influence on the eventual result of a choice. Free will does not rely on such a mechanism. Free will constitutes a choice that is self-generated. A freely willed decision can never be fully predicted because it has no deterministic analog, which means that there is no antecedent state that has, according to a defined law of physics, determined the subsequent existence of the choice state. Free will does not require intrinsic programs that are hardwired in our brain, otherwise it is determined by a program that relies on antecedent states that directly, and therefore predictably, cause future states.

Now, the goal of this chapter is to place that proposal into a scientific context. Here's one attempt at explaining decision-making. We encounter something, and we receive sensory

information about it (e.g. the shape, color, smell, movement etc of an object, or of someone's face in response to something we say). When we perceive it the sensory information gets processed in various areas of the brain according to the modality (e.g. touch, taste, smell, sight, sound) or the valence (i.e. what kind of attention is needed). Processing alerts parts of the brain that are required to decide how to respond. Several patterns of neural activity, which represent possible reactions, get some traction and whichever presents with the most support, and therefore has the highest sway, is selected to precipitate activity in pre-motor areas that are in charge of designing the response and executing the excitation of relevant motor areas.

This scientific model is the simplistic "sense-think-act" cycle that is the groundwork of the roboticist perspective. Roboticists are the eliminitive incompatibilists of the lab. In his book, "Mental Biology: A New Science of How the Mind and Brain Relate," lifelong neuroscientist W.R. Klemm defines roboticists as those scientists and philosophers whose position "requires the assumption that humans are biologically programmed robots… that our brains have been programmed by our genetic endowment and by past experience–thus, the conclusion follows that we are biological robots" (172-173). In essence, the roboticist belief is that (a) the activities of the brain are reducible to matter that makes it up (cells, subcellular organelles, molecules, and all the way down to the particle physics of atoms), and (b) that the structure and function of the matter in our brain, at all levels, is reducible to hardwired mechanisms and direct stimulation from the outside world. While a roboticist might accept that there is a degree of uncertainty, or indeterminism, when you get down to the quantum level, they would suggest that this would only modify their claim if it turned out that activity in the brain is thereby random. Nonetheless, whether determined or random, choice would still be outside the grip of anything like free will.

Before moving forward, I should take some time to discuss the physics of indeterminism at the quantum level. What do I mean when I use the term quantum uncertainty? Such phrases like "quantum uncertainty", or "Heisenberg's principle" may strike fear in our hearts, but it is not necessary to understand the complexity of the topic in full. I will endeavor to summarize the relevant material so that we can understand quantum indeterminism, and how it plays a role in

the capacity for free will. In 1927, German physicist Werner Heisenberg introduced his principle of quantum uncertainty, which suggested that the more we know about one quality of a particle, such as its location in space, the less we can know about other qualities of that particle, such as its momentum. This is to say that there is ultimately a limit to what we can know about a particle. Imagine you're watching a film at the theater but the screen is very dim so that you can barely make out what's happening. The attendant tries to correct the error, and successfully brightens the image, but simultaneously, the sound dies out so we can't understand what's being said, even if we can see people talking. This helps us to understand the way in which quantum states are measurable. One question we may ask, however, is whether this limit is epistemological, bound by insufficient technology, or ontological, an intrinsic fact of nature?

Originally, Heisenberg attributed some uncertainty to what's called the "measurement-disturbance relationship" wherein the act of measuring a particle changed the quantum state. As such, we can not ascertain the precise state of a particle. Some recent studies have been finding that the measurement-disturbance relationship doesn't account for as much uncertainty as was previously thought (Rozema et al., 2012). However, the team still notes that "in the end, there's no way you can know [both quantum states] accurately at the same time," which supports the fundamental aspect of the principle of quantum uncertainty, which purports to describe an intrinsic aspect of nature that cannot be known.

Why can't it be known? The reason for this is that these qualities of a particle are seemingly indeterministic. They follow no definable laws of deterministic physics that we know. They behave randomly. There are clouds of probability wherein a particle may reside, or ranges of momentum the particle may have, but where precisely in the cloud or in the range the particle may be is a matter of chance. Thus, quantum uncertainty, uncertainty about the quantum state of a particle, renders physical systems indeterminate, at least at a small scale. Given the success of quantum mechanics in displaying this indeterminate nature of particle physics, we are confronted by two problems for free will. First, how can this indeterminacy have an effect on such a large, complex system as the brain without being entirely superseded? How could quantum states really

have a meaningful effect on atomic structure and activity, let alone molecules, cells, circuits, and networks that eventually conglomerate into a single brain? Second, even if there were a way in which quantum uncertainty could translate on a larger scale to the brain, how can we exert free will in a way that controls randomness, but isn't determined? In other words, how can there be a middle ground between random, uncontrollable indeterminism, and pre-ordained, predictable determinism?

These questions will take some time to even begin to answer here, not to mention the time it will take for us to really develop sufficient empirical (or even theoretical) support for a middle ground. In this chapter I seek to support the plausibility of this middle ground given our current understanding of physics and neuroscience, and how science and philosophy can team up to explore these avenues moving forward. To be clear, I am not attempting to prove the existence of free will with our current evidence. Instead I am supporting the idea that it is an open scientific and philosophical question, that it has not been disproven, and that it is plausible that we may find an explanation in the direction of a middle ground between indeterminism and determinism

For now, let us endeavor to answer these questions by returning to the "sense-think-act" cycle. Philosopher and cognitive scientist Andy Clark, in "Mindware: An Introduction to the Philosophy of Cognitive Science", writes that this cycle is supported by roboticists who suggest that there is "a simple division of labor in which perceptual processing yields a rich, detailed inner representation of the 3D visual scene, which is then given as input to the reasoning and planning centers, which in turn calculate a course of action and send commands to the motor effectors" (88). In this cycle we see the kind of distinction between deliberation and the decision that was assumed by McKenna's thought experiment.

It has become clear that this simplification is not true in many ways, a couple of which relate to problems with McKenna's model of the decision-making pathway. According to the roboticist model, and according to what McKenna implies to be a model for cognition, there is a period of deliberation and then a final, cumulative decision, which then initiates action. More

recent findings suggest that this cycle ignores the interactive quality of this cycle wherein, as Clark summarizes:

1. Daily agent-environment interactions seem not to depend on the construction and use of detailed inner models of the full 3D scene.
2. Low-level perception may "call" motor routines that yield better perceptual input and hence improve information pick-up.
3. Real-world actions may sometimes play an important role in the computational process itself.
4. The internal representation of worldly events and structures may be less like a passive data structure or description and more like a direct recipe for action (88).

As such, what Clark calls a "compelling hypothesis" is the new model which suggests that "the visual system is not even attempting to build a rich, detailed model of the current scene, but is instead geared to using frequent saccades to retrieve information as and when it is needed for some specific problem-solving purpose" (91). All of this means that decision-making is not really a straightforward cycle. It can be more accurately described as an interactive system that monitors its progress, adjusts to specific demands, and always regulates the computations it performs, even after the "decision" is made by processing centers.

What does this mean for my philosophical claims? This suggests that I was right to question the conclusions McKenna draws based on his thought experiment. The act of decision making is too complex for us to be able to say that the lesion in Jasper's brain played no role in his final decision, perhaps he could have done otherwise. It isn't possible to defend Frankfurt's objection to alternative possibilities by selectively limiting alternative possibilities at the "moment of decision" instead of for all deliberation, because decision-making and deliberation are intertwined, and may even be identical processes for some decisions, as Clark's point number four above addressed. The remaining features of the interactive hypothesis clearly establishes the complexity of the decision-making process. Now, such a complex process, with numerous

variables that need to be controlled constantly, comprehensively, and with some degree of consistency, needs to be under control. Much like businesses, governments, or individual tasks, the bigger, or the more complex our decision-making process becomes, the more need there is for a meta-regulatory network to oversee the completion of deliberation and choice. This opens the door for conscious intervention on the whole causal pathway of decision-making.

Many books on the topic of consciousness are at some point confronted by the same question: if phenomenological consciousness plays no role in the causal pathway of decision making, then what on earth is it for? We know that consciousness is clunky, time-consuming, and inefficient in comparison to the unconsciousness. This is why so much of our lives is run by unconscious elements of our brain; if we had to keep track of every heartbeat, every breath, every step, every movement of our mouth to articulate a word, so on, we would be incredibly overwhelmed. We wouldn't be able to handle it, and there would be no time for sports, or love, or philosophy. Why, then, would we have evolved to have such a rich phenomenological experience if it did not impact our decision making in some way? Why, when our unconscious mind seems fully capable of controlling all aspects of our basic life, as recent research is showing, do we have this inefficient, and ultimately inconsequential experience?

"A Brief Tour of Human Consciousness" is V. S. Ramachandran's analysis on the topic of the mind, both conscious and unconscious. In his exploration of Libet's experiments, which we have discussed earlier (where it was found that the feeling of willing something occurs after the initiation of relevant activities in the brain), Ramachandran is confronted by this why question: "Why might this be happening? What might the evolutionary rationale be? … If the subjective sensation of willing is like a shadow that accompanies us as we move but is not causal in making us move, then why would evolution bother delaying the signal so that it coincides with our movement?" (88-89). He suggests that the delay of conscious experience from the initiation of relevant brain activity may have developed so that the sensation of willing an action would "coincide not with the onset of the brain commands but with the actual execution of the command by your finger" (88).

From a cognitive science perspective, this makes sense as there needs to be consistent, temporal contingency between two events in order for humans to suspect that they are causally related. As we saw with Wegner's tree branch example, if a tree branch moves exactly as we imagine, at the same time as we imagine it will move, and does so consistently, then, and almost only then, is it relatively easy to think that we caused it to move. Likewise, if our feeling of willing our action was not temporally proximal to the action itself, then perhaps we wouldn't feel so strongly that we had willed it. Imagine if we thought about moving our finger, and it happens almost a full second later. Imagine now that all our actions come with this delay, while we might get used to it, it may not convince us that there is an unimpeachable causal link.

It could be argued that that this means nothing about whether or not this feeling of willing is causal. It's just something our brain makes us experience in order to, perhaps, inform future decisions that the brain makes for us. The same can be said in the other direction: if it's just a matter of our brain trying to give us the experience of willing, then it stands to reason that, instead of delaying the signal to the time of action, the brain could have evolved a different capacity for the feeling of contingency such that we would still have the sensation of willing even if the awareness occurred without delay, well before the action. In other words, why come up with the contraption of delayed experience instead of more liberal contingency? Would a more liberal sense of contingency alter our feeling of control over events in our environment? It's possible, but if the contingency were to be a small range at the exact time when action occurs, then it's unlikely any environmental event would consistently fall within that range.

Ramachandran identifies a paradox in Libet's findings, which stems from our intuitive understanding of causation:

> On the one hand, the experiment shows that free will is illusory: it cannot be causing the brain events because the events kick in a second earlier. But on the other hand, the delay must have some function, otherwise why would the delay have evolved? Yet if it does have a function, what could it be other than moving (in this case) the

> finger? (89)

He notes the intuitive conclusion from Libet's experiments (that free will is an illusion), but also points out that there must be a functional purpose in the precise delay of the feeling of willing. It had to occur closer to the event to be contingent, but not too close, otherwise we wouldn't draw the intuitive conclusion that our will must have caused the action because that which comes first would cause the latter event. He concludes that "perhaps our very notion of causation requires a radical revision… as happened in quantum mechanics" (89), which strongly supports the direction that my philosophical claim has taken, suggesting that a specific amplification of quantum uncertainty introduces a non-deterministic quality to decision making, which could solve this paradox.

It may seem like the only solution is to find a way that the universe can function wherein events can cause their antecedent events (B can cause A even though it occurs later temporally). This sounds like a demand for a time traveling solution; event B must be able to inflict causation backwards in time, or something crazy like that. Now, it may be the case that time is not linear, that causation may have a completely different temporal profile than we expect, but this is not the necessary revision to our notion of causation that is required to solve this paradox.

In his attempt at a solution, Ramachandran connects matter in the brain to qualia, or phenomenological experiences like "the redness of red, the flavor of Marmite or paneer tikka masala or wine" (96). To start off, he says that qualia "must have evolved to fulfill a specific biological function–they cannot be mere by-products (an "epiphenomenon") of neural activity" (98). But what is the biological relevance, and how does it support the notion of free will? Ramachandran writes that:

> Sensory representations that are themselves devoid of qualia might acquire qualia in the process of being economically encoded or "prepared" into manageable chunks as they are delivered to a central executive structure higher up in the brain. The result is a higher order representation that serves new computational goals (98-99).

To some extent it can be argued that the biological purpose might just be to create an experience as a tool to better compute a decision. The idea that the brain is creating this sensation for itself to feel, instead of us, as agents, experiencing it, seems to degrade any chance at reserving free will, but some philosophers have found ways to argue from this structure that Ramachandran describes that there is free will.

This brings us to Dennett's narrative consciousness philosophy, which is supported and excellently described by Andy Clark. He writes that Dennett's idea is that "a system has a belief just in case its behavior is well predicted by treating it as a behavior" (178), which is called the "Intentional Stance" of consciousness. The brain is considered "the locus of multiple, quasi-independent processing streams. There is no single, ultimate judgement issued by the brain in response to an input–no decisive moment in space or time where the system settles on a unique definitive content fixing the conscious state" (178). As such, Dennett writes in his book "Consciousness Explained" that consciousness is a constant narrative that is "not anything we are born with, not part of our innate hard-writing, but in surprisingly large measure, an artifact of our immersion in human culture" (346). This organization allows us to make "cognitive objects of our own thought processes and to weave a kind of on-going narrative (about who we are, and what we are doing, and why we are doing it) that artificially "fixes" the cognitive contents" (Clark, 179).

Personally, I don't know how to feel reassured by this philosophy of self, and of free will as a sensation requisite of this narrative self. This concept is asserting that all our experiences of agency, of beliefs, of self, are complex narratives that our brain conjures for us. Imperatively, the "for us" is not 'us', as in our phenomenological experience of ourselves, but 'us' as in our brain. This illusionary narrative is not generated for the benefit of some small, weak agent to watch like a movie, the agent is the movie that the brain is watching, presumably as a monitor to better inform itself about how to proceed in its day to day activities of determining everything our bodies do and everything we will feel subsequently. This makes experience feel like a tool for our brain to monitor its progress, which disagrees with my definition of free will. While

the narrative is self-generated, Dennett's claim is that the self is our unconscious body, which, without a causal conscious influence, can only be shaped by the factors outside of our control. In the next chapter I will provide evidence that this is not true.

Returning to Ramachandran, he continues to describe this structure of consciousness, but takes a slightly different approach. Recall that he has established the existence of "a higher order representation that serves new computational goals." He expands on this representation and says:

> Let us call this second, higher-order, representation a meta-representation… One could think of this meta-representation as a second "parasitic" brain–or at least a set of processes–that has evolved in us humans to create a more economical description of the rather more automatic processes that are being carried out in the first brain. Ironically, this idea implies that the so-called homunculus fallacy–the notion of a "little man in the brain watching a movie screen filled with qualia"– isn't really a fallacy (99).

So far, he seems to be describing a similar scenario of a brain that creates a movie-like narrative for another part of the brain to watch. However, Ramachandran elaborates on the concept of the homunculus in a way that suggests not that it is generated by meta-regulatory activity, but that the homunculus generates the activity itself, or at least that they are one and the same. As he writes: "I suggest that the homunculus is simply either the metarepresentation itself, or another brain structure that emerged later in evolution for creating meta-representations… it doesn't have to be a single new structure–it could be a set of novel functions that involves a distributed network," which then "serves to emphasize or highlight certain aspects of the first [representation] in order to create tokens that facilitate novel styles of subsequent computations, either for internally juggling symbols sequentially ("thought") or for communicating ideas to others through a one-dimensional sound stream ("language")" (99-100).

This direction of thought takes us closer to what I conceive to be a new model of cognition that supports the existence of free will. I proposed in Chapter 2 that we are capable

of self-generating motivation. What I meant by that is that free willing agents are capable of altering events in their brains in order to favor one set of calculations over another, and thus to favor one decision over another. The change of calculations is a change in motivations. In terms of practical reasoning, we generate new motivations by creating new reasons to act one way and not another. In order to achieve this we select a new way of calculating what decision to make. The action of calculating, in brief, consists of neurons firing and activating each other in a pattern that initiates later activity, which eventually will result in activity in motor areas that stimulate bodily responses. This mechanism proceeds deterministically, otherwise there would be a disjoint between what decisions cause which actions, and this would be biologically disastrous.

However, the question of free will arises earlier. In Chapter 2, I proposed that free will is the ability to change which parameters are important for these calculation; free will is the ability to change the rules, and it must proceed on a middle ground between determinism and indeterminism. For more scientific support of this claim, I will turn first to Peter Tse for a view of free will at the level of the synapse, the communicative connections between neurons, then I will refer back to Klemm for a view of free will at the level of meta-regulation, which will bring together the effects at the small level of the synapse to display how consciousness exerts brain-wide control over neural events, thereby exercising free will over choices and actions.

Synaptic Free Will

"The Neural Basis of Free Will" is Dartmouth professor of cognitive neuroscience Peter Tse's unprecedented attempt to answer the free will question by looking at the activities of individual and small groups of neurons. Through careful analysis of the physics, chemistry, physiology, molecular biology and computational aspects of the brain, he finds that a strong free will, much like I have defined, is not only plausible, but empirically verifiable. In the following section I will distill the many components of his argument to a concise explanation of how indeterminism enables free will. I will also attempt to convince the reader that this scientific

proposal supports precisely my philosophic claim that free will constitutes self-generated motivation.

Tse argues that his philosophy of free will applies regardless of whether we accept or deny indeterminism. However, he generally assumes that the universe has been found to be at least sufficiently indeterministic for a "strong free will" to exist, because it supports the fact that "things really could have turned out otherwise" (2.1). As we have seen, there has been some confusion about whether observed indeterminism was an issue with our method of measurement or an ontological reality. As Tse points out, "determinists are loath to give up the idea that unmeasured properties nonetheless occupy a definite state at each moment," (A1.6) which is to say that determinists support the idea that unmeasured quantum states of a particle still exist in a set form, they are just inaccessible. Imagine you've walked into a dark room without lights. You use a flashlight to determine first the location of one wall, and then the wall opposite to it. You can't see both clearly at the same time. The determinists argue that, regardless of which wall you're looking at, the other exists in a set location. Of course, this is the most intuitive understanding of the universe, but indeterminism suggests that at the quantum level, the precise qualities of quantum states, like the nature and location of the opposite wall, are not set, they do not "occupy a definite state at each moment".

Tse also writes that many determinists have tried to suggest a variety of complex alternatives to quantum indeterminism, but ultimately "there is no need to invoke quantum nonlocality, superposition, entanglement, coherence, tunneling, quantum gravity, or any new forces to understand informational causal chains in the brain" (A1.8). He writes that some of these theories are inapplicable because they require circumstances that are physiologically irrelevant ("the brain is, simply put, too "warm" to support this kind of quantum-domain coherence, and synapses are too wide to support electron tunneling"), whereas the rest are improbable in comparison to quantum indeterminism ("I doubt that quantum-domain effects— beyond the variability in neural dynamics introduced by amplification of microscopic fluctuations—are required to account for how information is processed by neurons").

He is careful to note that while indeterminism introduces a kind of gap in the causal chain, as Kane has suggested, free will doesn't require that consciousness "play a necessary role in determining which possibility is actualized," which may sometimes seem unlikely, but rather that "consciousness, and the entertainment of possible scenarios and courses of action in working memory, plays a role in changing the criteria for firing on neurons that might lead to future mental events" (A1.9, emphasis added). This classification of free will supports what we have already discussed wherein consciousness changes the rules of the game in order to influence decision-making.

Having established the premise of indeterminism, Tse goes on to acknowledge the apparent problem of mental causation. The question is how mental events can change the physical substrate that encodes that mental event. "Higher-level facts in the present are realized in physical facts in the present. In this regard, any information that is realized in neurons that are active now cannot change the physical basis of that information itself. That would be impossible circular causation" (2.3). Thus, we know that mental events arise from coded information in neuronal activity, but how can you reverse the direction of causation such that the mental event influences the activity that created it? Tse writes that we don't need to reverse the direction:

> Neurons do not only make other neurons fire; they also recode the informational criteria realized in the physical criteria that will make other neurons fire at some unknown time in the future when the right inputs come along. The central thesis argued here is that physically realized mental events can change the physical basis not of present but rather of future mental events by triggering changes in the physical and informational criteria that must be met by future presynaptic inputs before future postsynaptic neuronal firing occurs (2.3).

This is a lot to digest, but if we pick apart the terminology the central thesis becomes clear. First we will look at how a neuron works, in brief, and then his philosophy of criterial causation, then we will pick apart the neuroscience to understand how criterial causation is involved in neural

events on a cell to cell level.

Neurons are cells in the brain that conduct charged ions (e.g. Na^+, K^+, Cl^-, Ca^{2+}) across its membrane through ion channels. Cascades of moving ions, such as rapid influx of Na^+, result in flow of electrical current and changes in membrane voltage, which is the difference in electrical potential between the outside and the inside of the cell. These changes can culminate in a hyperpolarized (less excited) state, where the difference in potential between the inside and the outside of the cell is large, or in a depolarized (excited) state, where the difference is minimized. During depolarization, the resting membrane potential rises rapidly until it reaches the reversal potential for Na+, and then reverts to a more hyperpolarized state before finally rising to the stable resting potential. These changes constitute the rapid influx of Na^+ (rising potential) followed by rapid efflux (lowering potential), and the slower influx of K^+ (the stabilizing increase to resting potential).

This spike in potential is called an action potential. It arises at the axon hilloc, between the base of the axon and the body of the cell, and is transferred by a chain reaction through the axon to the axon terminals. The terminal ends of the axon connect with the dendrites, receptive branches, of another cell, forming what's called a synapse. When the action potential reaches the terminal it signals the release of neurotransmitters into the synapse between pre and postsynaptic neurons. The neurotransmitters initiate ion flow through ion channels in the postsynaptic neuron, which sends signals from the dendrite to the cell body, where the information is processed. If input to the cell from the synapses meets a threshold, it initiates an action potential that repeats the cycle to stimulate the next cell in a neural circuit.

Criterial causation relates to physical and informational criteria for this neuronal firing. In order for a neuron to respond to stimulation in a certain way (i.e. with a spike, or two, or not), the stimulation it receives must meet some physical criteria (such as exceeding a threshold electrical potential before the cell fires from the axon hilloc) or informational criteria (such as what information is mentally realized when the cell is fired). There are two forms of plasticity,

synaptic and intrinsic plasticity, that control how much stimulation a postsynaptic cell is capable of receiving from a presynaptic neuron, and how a cell's intrinsic properties integrates synaptic input and produces output activity.

Now let's discuss how criterial causation can fit into a model of neural events. Tse suggests a three stage model for criterial causation:

> (1) New physical/informational criteria are set in a neuronal circuit on the basis of preceding physical/mental processing at time t1, in part via a mechanism of rapid synaptic resetting that effectively changes the inputs to a postsynaptic neuron. These changes can be driven either volitionally or nonvolitionally, depending on the neural circuitry involved.
>
> (2) At time t2, inherently variable inputs arrive at the postsynaptic neuron,
>
> and (3) at time t3 physical/informational criteria are met or not met, leading to postsynaptic neural firing or not (2.7).

What Tse is suggesting is that, after a neuron participates in the creation of a mental event it adjusts its synaptic and intrinsic criteria for when and how it will respond to future neural stimulation events. He doesn't explicitly refer to intrinsic criterial adjustments because research into intrinsic plasticity is less complete, but there is sufficient evidence to show that intrinsic plasticity fits right in this model. These adjustments can be mediated by volition, or conscious will, or not. After making adjustments, a cell will inevitably receive new stimulation. Depending on whether the new criteria are met, the cell will respond a certain way, maybe it will fire, maybe not. Tse argues that indeterminism can be involved in the first two stages, but not the last, because "intracellular potential either passes the threshold for firing or it does not" (2.7). Tse's model parallels the model I have established as it also defines a period of indeterminacy and the stages after which indeterminism is irrelevant and events occur on a deterministic basis.

Now we have the foundation needed to understand Tse's central thesis: "physically realized mental events can change the physical basis not of present but rather of future mental

events by triggering changes in the physical and informational criteria that must be met by future presynaptic inputs before future postsynaptic neuronal firing occurs". In other words, after a mental event of any magnitude occurs, there is a way for regulatory activity to relay information back to the neurons from the mental event originated, and in so doing can influence the postsynaptic cell's criteria for responsiveness to future stimulation by presynaptic events. This carefully sidesteps the problem of mental causation, which states that a mental event arising from physical substrate cannot modify itself by modifying the physical substrate. But Tse notes that his model certainly describes free will because while "there can obviously never be a self-caused event… criteria can be set up in advance, such that when they are met, an action automatically follows; this is an action that we will have willed to take place by virtue of having set up those particular criteria in advance" (7.5).

A remaining question is how a mental event can influence the change of a neuron's physical and informational criteria. Tse holds that, "assuming amplification of microscopic fluctuations, criterial causation permits downward mental causation and a strong free will because neurons can set up criteria for future action potential release that, once satisfied, lead to nondetermined yet self-selected future actions" (2.6). We will discuss amplification further when we discuss Klemm, but the point here is how this downward mental causation presents itself. The questions that remain are: what is the precise nature of this changing criteria, and how do some mental events cause that change in a way that instills free will in the system?

First, the nature of neuronal criterial changes. We need not dive too far into the details of complex mechanistic criterial changes, a topic that Tse describes in thorough detail in his book (Chapter 5), as it is complex, and not yet fully understood. The strength of the proposal is understood well enough by understanding the following, which is only one example of a type of neural plasticity that can be involved in criterial causation.

There are a subset of receptors that are located in the membrane of a neuron at the synapse called NMDA receptors. NMDA receptors belong to a family of receptors called

ionotropic receptors, which are made up of several subunit proteins. These typically four subunits line up like a bouquet of proteins with a space between them through which ions can pass, like through a channel. In the case of NMDA, the proteins are signalled to open the channel when the neurotransmitter glutamate and the amino acid glycine bind to the surface, and when the neuron in question is currently depolarizing (being excited, or spiking). When open, the channel permits Na^{+} and Ca^{2+} to permeate the membrane, enter the cell and further depolarize the membrane as well as regulate a variety of features, such as the neurotransmitter release subsequent to the cell spiking.

Tse proposes that NMDA receptors can influence future spiking events because they control the precise timing and strength of synaptic input, which mediates spike-timing dependent plasticity. Spike-timing dependent plasticity refers to how the activity and informational coding of a cell will change depending on how it spikes in the wider context of the cells around it and the timing of the introduction of a stimulus. If a spike occurs after an upstream event, it fits in the causal cycle (wherein an object responds to something that came before it) and cooperates with other cells to contribute to the larger response. On the other hand, if the spike time is before an upstream event, then it doesn't fit in the picture and is considered to be unrelated to the phenomenon, and therefore only a competitor for attention, not a contributor. Imagine the scenario of a basketball game: if someone in the crowd cheers before a shot is taken, they would appear out of place, and wouldn't contribute to a group response. It might appear that they are cheering for some other purpose, like the pass preceding the shot, or an email they just received telling them they've been accepted into a competitive graduate program. If they cheered just after a shot, like is normal, then they appear much more coherent with the rest of the crowd.

While this wouldn't happen in the same way with humans, of course, plasticity among neurons, the ability to change connections and individual activity, would result in the exclusion and inhibition of cells that spiked before the stimulus, whereas those that spiked together and at the "right time" would be welcomed as contributors to a pattern of response. This would shape the way cells and the brain as a whole respond moving forward. As Tse writes:

> Over time, such spike timing-dependent weight changes could result in the "sculpting" of neural circuits, in which spikes, and the information realized in them, are causal of later spikes and information, from level to level in an information-processing series of stages. Such circuitry allows the encoding and learning of information about external causation in terms of changes in synaptic weights and the ordering of neuronal firing in a network (5.11).

Relating this back to NMDA receptors, Tse suggests that changes in NMDA receptors can reweight the interactions between pre and postsynaptic neurons in such a way as to sculpt neuronal circuits to respond to future events (which can be events one second in the future or years later) in a way that can be mandated by top down mental causation (free will) and not just by deterministic responses of a cell to stimulation.

Since NMDA receptors "can play a role in physically realizing so many of the most fundamental properties of information processing in neuronal circuits" (5.25) yet be altered by such small atomic events, they effectively amplify quantum level uncertainty to indeterminism at the level of the cell. This ability enables NMDA receptors to be a basis for the indeterminism that is wielded by higher level processing to produce creative optimization heuristics when deliberating. More on that later. Tse writes conclusively that:

> NMDA receptors may play at least four important roles in the neuronal basis of mental causation: (1) they introduce randomness to the macroscopic domain in the form of spike-timing uncertainty, which may then meet or not meet present neuronal criteria; (2) they may play some role in an "annealing" process that minimizes the chances of getting stuck in local minima of thought or behavior, affording the generation of many possible solutions that can meet a set of criteria; (3) they may realize rapid synaptic resetting that permits neurons to escape the impossibility of self-causation by changing the criteria for firing of neurons on the next information-processing cycle; and (4) they transition neural network behavior

> into a bursting mode that allows for the transfer of bursts as informational packets across a network, which, I argue, is the neural basis of attentional binding (5.51).

For now we will focus in on roles 3 and 4, and how they demonstrate a capacity for the brain to reweight synaptic responsiveness in a way that controls future output. We will go back to roles 1 and 2 when discussing how indeterminism initiates a process of creative problem solving in order to arrive at potentially novel choices in the face of alternative possibilities.

Tse writes that changes in NMDA receptors can be initiated by back-propagating action potentials, which send the signal to dendritic receptors to change. The back-propagation of action potentials can be "altered by network inputs that sculpt which branches are 'active'" and they "could be a graded or gain-modulated process that effectively [adds] or [decreases] resistance to the passage of an action potential up particular dendritic branches" (5.26). Thus, "network inputs" would have control over which and how dendritic branches are able to undergo synaptic reweighting. Tse proposes that network inputs could be the encoded information delivered by bursting activity (repeated neuronal firing) that is optimized in working memory and mediated by endogenous attention.

Endogenous attention and working memory are the core components of Tse's model for top-down mental causation. For Tse, we are able to generate ideas in working memory that are simulations of action and repercussion in order to deliberate on the solution to a problem. Endogenous attention allows for the manipulation of the contents in working memory, to either select (enhance) or deselect (attenuate) specific options. Arriving at optimized solutions we can then direct lower level cognitive mechanisms that will provide input to specific groups of neurons that will undergo synaptic reweighting to accommodate the intent of upper level mental events and thereby sculpt the circuits and networks of our brain into what we choose to be. What gives us the choice, what makes this a model for conscious free will, is the fact that these highest levels of cognitive processing require consciousness and phenomenology to operate.

Consciousness acts as a unifying model for decision-making, much in the same way as

Ramachandran considers it a homunculus for the self. Consciousness is a single framework, inside of which our decision-making can operate. This answers the question about why we have consciousness, why we have phenomenological experience, and why we experience qualia. These are all causal mechanisms, they are all necessary and requisite to higher level processing, which in turn controls future cellular responsiveness to stimuli. Qualia, Tse writes provide "a common, endogenously attendable format [which] is required so that the outputs of different subsystems can be evaluated along a common metric, by a common set of criteria assessing input, which, if satisfied, releases a decision to act" (10.24). In other words, in order for the complexity of the brain to be cohesive and efficient, all levels of processing need to communicate in the same language, which is what Tse proposes is the goal of qualia, or the ways that we experience things, like the indescribable sensation of the color red.

Higher level processing relies on ideational information in the form of representations, which are what we generally think of as thoughts: mental images of basketball games, mental simulations of our actions, the abstract concepts that are debated in a philosophy class. These representations, are built from the cumulative experience of qualia for in one subject, like the combination of sensations during a basketball game. Pivotally, the highest levels of processing, for Tse, require consciousness to function, thereby suggesting a role for conscious, phenomenal experience, in the causal mechanisms of our brain. Further, Tse notes that the qualia we experience aren't simply an uncontrollable stream of processing in lower level unconscious mechanism; he suggests that the contents of experience "can be changed through the manipulations of numerous cognitive operations, such as mental rotation or volitional shifts of attention" (10.48).

By way of concluding the discussion of Tse's many-sided model of free will, I will first allow him to tie it all together:

> Certain operations can only take place over conscious operands, and thoughts and motor acts can follow and build upon or enact the conclusions of such operations,

> such operations over conscious operands can be causal of subsequent thoughts and motor acts. Such mental operations within the domain of experience play a necessary (though not necessarily sufficient) causal role in their mental and motoric consequences and are not mere illusions of volition... It is the deliberative and volitional manipulation of contents of working memory via endogenous attentional operations that offers the primary causal role of consciousness and the primary domain of free will, regardless of whether or not such operations are associated with conscious feelings of willing or agency. Free will is not limited to the kinds of decisions that are reached by tracking or playing things out endogenously in working memory. Free will also includes the creative ability to imagine, regardless of whether any physical act follows from that which has been imagined (10.85).

For Tse, our self is a real and true self. Each of us is a self that is identical with the substrate of our brain, obeying functionalism, but that is capable of engaging in conscious, bottom up indeterministic thought in working memory and that uses attention as top down control over lower level cognitive mechanisms that instigate synaptic reweighting. This reweighting forces cells to respond differently to future events in a way that is mandated by conscious free will.

It is important to recall, before we leave Tse, that his example of NMDA receptors is just one example. There are many more receptors and non-receptor factors involved in synaptic reweighting. Furthermore, there is a whole field of ion channels and associated factors involved in nonsynaptic (intrinsic) plasticity. There are numerous mechanisms that can be involved in adjusting neural criterial causation, many of which have yet to be explored. Interestingly, my own neuroscience research in the optic tectum of Xenopus tadpoles has rendered results that unwittingly highlight the relevance of neural noise in encoding information, which supports indeterministic prerequisite activity to be acted upon by endogenous attention, and which has investigated the role of intrinsic plasticity in maintaining homeostatic, or sustainable, levels of activity in the face of long term stimulation.

Whole Brain Free Will

Tse's discussion on the topic of free will is extensive and covers a great deal of ground. As such, I will be presenting Klemm as more of a supplement to Tse's model that brings us back toward practical reasoning. In our discussion of Tse we have seen a thorough explanation for just about every component of a complete model for free will except how executive control comes about and what it really means at the level of deliberation, representations, and reasoning. How is executive control over endogenous attention and working memory functions free? It is clear how we can enforce a decision, and that these executive processes have a goal, and evolutionarily advantage, but what makes this control indeterministic at such a large scale? How do we have alternative possibilities from quantum uncertainty that reach all the way up to the representations of executive deliberation. Ultimately, the question I seek to answer with Klemm is: how do executive functions constitute a free willing self?

Klemm writes that "there is no doubt that humans exhibit willful considered behavior. We have, many people say, free will. That is, we supposedly can decide from among several alternatives, and this process is not directly controlled by external or internal imperatives." (168) From here comes the idea of a human avatar. According to Klemm, our consciousness is essentially an avatar, like the characters in video games, which act as proxies for the gamer in the world of the game. However, Klemm writes that "the analogy of computer avatars breaks down when it comes to human consciousness avatars. Computer avatars cannot generate their own agency." (146) He notes that "brain avatars respond to impulses and instructions from unconscious mind as well as ongoing environmental effects. Unlike computer avatars, however, the brain's avatar may act on its own initiative to serve the best interests of the brain and the body" (146).

For Klemm, this avatar is nothing like a miniature person, a homunculus, inside our head that makes free willing decisions for us. The fallacy of the age old 'homuncularist' argument is

clear, if there is some model person inside us that regulates everything we do as a unified agent, then what controls that person? A smaller homunculus? For Klemm, the avatar is spread out in a series of representations of self that centralize internal simulations of action and thought. It is thought that a vital characteristic of human learning is the ability to simulate action in an internal model of the world, kind of like generating a virtual reality, in which we can place ourselves and test out different scenarios to optimize our solution to a problem.

The basis for this avatar is what Klemm calls a "CIP representation of the body" (146), where CIP stands for circuit impulse patterns. According to Klemm, circuit impulse patterns are "simultaneous activity across all neurons in a circuit at a given point in time," which is thought to be able to "have some sort of combinatorial or collective code" (38). Essentially, CIP's are similar to the bursting activity Tse described wherein a collective pattern of activity is capable of encoding more information than the mere some of its parts. As such, CIP's are patterns of activity that form representations, which are a complex combination of associated inputs. For example, granny smith apples. You'd have neurons that represent the color of the apple (green), the shape (spherical), and the size (2-3 inches in diameter). Whenever you encountered an object with those qualities, those neurons that are sensitive to the object's size, or shape, or color would be active in parallel, which would give rise to the representation of the fruit as a whole. Of course, our experience of the world is so much more rich than this, and our brain is equally and indescribably more complex than this example leads us to believe, but the general idea, when we understand it, remains the same. When a current impulse flows through larger patterns of neurons, instead of just a couple neurons, or basic circuits, it is capable of encoding and representing collective information.

It is also important to note that the CIP's do not exist in isolation, they interact with each other on a constant basis, recruiting additional CIP's in order to conduct complex deliberative processes. Back to the apple example, when we encounter this object, we may also recruit the CIP representation of a tennis ball to compare how each fits the scenario. After deliberating, after directing endogenous attention to particular regions of the object to try to identify a stem, or the

braiding pattern on a tennis ball, we can come to the conclusion that this object is an apple and that we can eat it.

Importantly, Klemm notes that CIP's underlie all basic brain functions "including the state of consciousness" (145). Thus, there is a "CIP state of consciousness," which is a network of CIPs that collectively represent the self. This CIP state is "contained within the meta-circuit of the whole brain and is therefore integrated with what goes on in unconscious processing" (145). Having a CIP-level representation of self provides a reference against which all other CIPs can measure their activity. Tse gave us a background for how qualia act as a common language for the brain to generate countless representations; Klemm discusses how CIPs, the architectural/ functional explanation of these representations, refer to a complicated, but unified representation of self in order to be consistent and efficient.

When I speak of common languages throughout the brain, or unified representations of self, I am discussing the cohesiveness of the brain at the cognitive level. Often, it doesn't feel like our brain is so clearly communicative; sometimes it feels backlogged, or stuck, or blank. Likewise, it doesn't feel like our concept of self is so clear either; frequently we feel as though we aren't sure what we like or dislike, what intentions we are driven by, or why we feel terribly sad when we 'should' be happy. There are approximately 20 billion neurons located in the cerebrum, the 'thinking' part of the brain, alone. There are thought to be five times as many neurons in the primal parts of our brain (the cerebellum and brainstem), and ten times as many glial cells as neurons; we know very little about the range of functions of glial cells, which puts in perspective how little knowledge we really have about the brain. Not only is there such a vast number, but they are organized with the efficiency and robustness of the primate family such that the activity of a few neurons is capable of encoding more information than other large brains that are less well organized, like in elephants and dolphins (Herculano-Houzel, 2009). My point is that, as much as examples, metaphors, and simplifications can help us to understand the underlying mechanisms at play, we have to remember and appreciate the fact that these mechanisms are being applied rapidly, very rapidly, at the scale of milliseconds, or even fractions

of milliseconds if we're talking about gene transcription and protein turnover; and all of this in huge quantity, far beyond anything we can really imagine (despite these activities enabling our ability to imagine anything in the first place).

When we discuss cognitive processing in the brain, the fact that all of these activities can combine into the smooth, continuous and sensical movement of my fingers to type these words that you can understand, displays a feat of effective communication and unification that is beyond belief. I hesitate to compare neurons to humans, as they are very different in key respects, but just to try and grasp the sheer numbers, imagine if every human on earth were capable of communicating and executing joint actions to move a giant, planet sized being. This is only one third the number of neurons in the cortex, and less than a thousandth the scale of the number of neural and glial cells that contribute to the thought and action of a single human. It's mind blowing. It only makes sense that neurons are communicating in a single language, and must be very unified in order for such complexity of thought and movement to be achieved. Thus, our brain is unified even if it doesn't feel like it, and often times if we declutter our thoughts and concentrate, we can overcome moments when we blank and return to a clear sense of what we do and do not know. Similarly, when we meditate and deliberate sincerely, we can usually uncover what our thoughts and feelings are if we ever feel confused. Doing so makes us better and stronger people because we exercise conscious control over what our brain does, and sculpt ourselves into people who can overcome our obstacles.

Turning back to Klemm after that moment of reflection, I want to show how he defends his theory. Klemm is sure to note that "skeptics will say that these aspects of the avatar's nature are preprogrammed by past experience or by genetic disposition," which is to say that the CIP structure of the avatar, the anatomical structure of the neurons or the functional characteristics those cells have, is dictated by genetics, not by any conscious sculpting of the self. In response, he argues that "experiential influences are developed and learned and become explicitly incorporated because they are acceptable to the avatar… [and] even much of the genetic influence is actually epigenetic, in which a person's choice of thought, behavior, and experiences

helps to control which genes get expressed" (169). Epigenetic mechanisms, in a nutshell, are modifications to our DNA that up or down regulate (increase or decrease) the expression of specific genes. These modifications, instead of being preprogrammed by our heredity, are accumulated, or removed, as a result of experience. Thus, our choice of what experiences to have exposes us to stimuli that will precipitate some or other epigenetic changes to our gene expression.

In his primary defense of free will, Klemm proposes that during the process of deliberation our conscious avatar can freely choose which considerations to take when faced with the problem. This process of deliberation, to Klemm, is essentially a complex network of CIP's that are all interacting and recruiting new CIP's in the process of recollecting relevant information in memory for a given problem. He proposes that "when decisions are performed by the conscious mind, the avatar can differentially select what inputs from stimuli and memory to consider, thus ultimately biasing one option over others. If conscious mind is an avatar, acting as an agent with a "mind of its own," what is to keep it from making its own choices? This might be the basis for free will" (195). A mechanism for this action makes sense biologically, which is to say evolutionarily, because, as Klemm points out, "the capacity for conscious action means that humans can rise above unconsciously driven behaviors that are unwise or maladaptive" (203). From free willing deliberation can come imagination, critical thinking, and creative problem solving, all of which generate new solutions that are not limited by what is predetermined by our experience, our heritage, or our prejudices.

Klemm's theory relates directly to the concept of self-generation. He brings in the example of crawfish who "will often move when they don't have to… no external stimulus is required" (173). To explain this, he writes that "there are, however, internally generated stimuli that turn movement initiating systems on. Likely such movement results from random neural activity in the pattern generator circuit, and the option to remain in place or move to a different place was always there" (173). Ultimately he writes that "random things happen in nervous systems. They are not always inevitable," and he suggests that these internally generated stimuli

"may be something as simple as inherent membrane instabilities in certain movement initiating neurons that make the neurons discharge impulses" (173). Tse has explained this randomized "instability" at a cellular level. Given that such events can happen in the nervous system we can see that new, or modified, neural activity can be generated without the influence of events external to the brain. New activity, even at the level of a single cell, can be amplified to influence the activity of many more neurons downstream (London et al., 2010). Thus, small changes could generate new CIP representations, which in their turn would introduce new representations to the process of deliberation. If they were selected and approved by the avatar that is our consciousness, then they would ultimately contribute to the representation of self.

***

Let's compare this to practical reasoning and to my theory of self-generating motivation. The CIP recruitment process is essentially the neural correlate of deliberation and reasoning. A given set of CIP's will represent reasons to act one way or another. During the process of deliberation the CIP's are tested for how well they remain consistent to the representation of self, which is the equivalent of saying that during the process of deliberation our reasons are tested for how well they remain consistent with our identity. As we've seen from the work of Tse and Klemm, during the process of deliberation, which is the activity of executive processes like endogenous attention and working memory, our brain can exercise new patterns of activity. These new patterns are not predetermined by external influence or hardwired programs by virtue of indeterminism that, at the cellular level, generates novel patterns of activity. Novelty at the cellular level can affect the activity of the networks of which they are a part, which can generate or influence entire CIP representations. Given new variants, or completely new CIPs, we can consciously choose one by directing attention to it. When we generate this new activity we are generating new reasons to act a certain way.

As such, I hold that consciousness does have a causal role in the redefining of neural activity, and in the restructuring of network connectivity. Not only is this a position based on

the scientific findings I've described, there are many philosophers who agree with Dennett on many points, but not on the theory that consciousness is strictly an observation. In "Self to Self: Selected Essays," J. David Velleman, who, like Dennett, holds that the self is constituted by a narrative, maintains that consciousness is action guiding, meaning that consciousness has a causal role in our decisions. As he writes, "In Dennett's view, our error about the self is to assume that the protagonist of a human being's autobiography is identical with the author" (206). While he holds that the "'central controller' of a person may indeed be a fiction…" it is still a conscious "reflective representation that feeds back into the person's behavior" (214). As such, Velleman suggests that "subjects are 'living out' their self-conceptions… they are doing things that would accord with what they have described themselves as feeling" (215). This invokes the importance of our conscious construction of identity in guiding subsequent action.

When the right constellation of reasons is found to be sufficiently optimal by our consciousness then it genuinely moves us to act. Selecting some reason over another stabilizes the neural pattern that generated the relevant representations. This physically sculpts our circuits, in the sense of how responsive cells are, and in the sense of how they are physically connected to each other. As such, future activity has an updated framework in which to act. This framework is vital because it forces future unconscious activity to flow through patterns of cellular criteria and anatomy that have been created and endorsed by our consciousness such that they act as an extension of our conscious deliberations, like water flowing through an aqueduct.

As I've said before, new motivations may compel us to act in the same way that we would have given a different constellation of reasons. For example, we may choose not to yell at our friend for their sake, or we may choose not to for our sake, or for the sake of someone else, or some other consideration. Thus, a single action may have multiple, antecedent reasons. I believe it is important to distinguish between which reason compels any given action, between what it is that genuinely motivates us. This motivation will be entirely novel compared to the motivation that would have compelled an automatic decision. An automatic decision would have undergone approval by the avatar for reasons that have been previously established during

former deliberation. I might automatically choose not to yell at someone for my own sake, or, after some deliberation, I may still choose not to yell at someone, and I may still make that choice for my own sake. I hold that the automatic and the deliberative decisions are always different. Any choice that has undergone a process of deliberation will necessarily be novel because it will always have had to verify itself against new information. Thus, every time we are making a conscious choice, not only are we self-generating a motivation, but this motivation is entirely novel and is part of a larger framework of choices that make up who we are and as whom we identify.

Chapter 4
Implications In Daily Life

This project has always been more to me than a purely academic debate over abstract, esoteric terms and conditions. This is not to say that the academic debate isn't valuable, it is, and its success relies on these careful debates over key terms and conditions. But at the heart of this project has been a beating curiosity about daily life. I hold concepts of responsibility and morality to be central in our day to day interactions with people and the world around us. My inspiration to explore free will is inspired by my own questions about what is moral, about what I must take responsibility for, and what I must be responsible for doing. This is not to say that the academic tilt of the past three chapters is a disguise for my true interests. I believe that an important tool for exploring these concepts is to disentangle the vines that have grown around this tree and try to access the core of the issues. To do this, I believe we need to resort to very basic questions about causality and freedom, which we have done. In so doing we have certainly ripped some leaves off the branches, and while some long, tangled vines remain, we can see the trunk more clearly, and I believe that the path ahead is less obscured. In this chapter I'd like to take a step back from the rigor of academic analysis and briefly touch upon some implications of my proposals that will end in casual observations about our day to day activities.

Moral Responsibility

First, how does this discussion of free will impact our morals? As I discussed very early on in the introduction, the fundamental question of morality depends upon responsibility,

or moral culpability. Moral philosophy explores what we should do. For the most part, and I acknowledge that there are exceptions, the question of what we should do in a given scenario depends on what we can be considered responsible for doing. There's no point in holding a position on what we should do if we cannot be considered responsible for the relevant outcome. For example, it would be silly for a moral philosophy to hold that we should not sneeze in the presence of a respected elder. It is silly because we cannot be held responsible for sneezing, it is a biological function that occurs automatically and uncontrollably.

How does free will come into play here? Returning to my introduction, I noted that questions of responsibility are essentially questions of control. If we have control over the outcome of a situation, then we can be held responsible for said outcome. Taking one step lower, the question of control is a question of free will, in my opinion. According to my definition, free will is the capacity to control our cognitive mechanisms such that activity in our brains is not a strictly automatic process. What does it mean for "us" to control our brain? It seems to me that consciousness is the most frequent metric for when "we" have control over our brain, as it seems wrong to blame someone for something they have done without consciously supporting it. I therefore suggested that consciousness is the seat of the self, and in order for our self to be considered responsible, morally culpable, for our actions, our consciousness must be capable of influencing the way we deliberate and decide on a course of action.

What we have seen in the chapters leading up to this is a brief background of the contemporary positions on free will, a proposal for how free will could be manifest in practical reasoning, and several perspectives on the scientific plausibility of this proposal. I have hopefully convinced you that the concept of conscious control over the physical substrate of our brain is not divorced from reality. Based on what we currently know, we do not require magic, or radical paradigm shifts to substantiate how conscious processes (the "us," the concept of self, the feeling of me) absolutely can exercise top down control over the lower cognitive mechanisms in our brain, which mediate the complex calculations required to design actions that will attend to a situation in the most optimal fashion. We have seen that the processes making up consciousness

are capable of generating new solutions to a problem in order to find an optimal solution, and then can change the criteria, or the rules, that lower processes in the brain follow to execute the relevant actions.

Returning to our moral approach to the world, the way we define and substantiate free will doesn't seem to impact what activities we define as either moral or immoral; however, I suggest that my model for free will can help us to determine under what circumstances an individual's adherence to our moral system is enough to be considered moral (or immoral), or if it is amoral. According to most moral philosophies an action is amoral, neither moral nor immoral, if it doesn't enter the moral plane, which is to say, if it is accidental, uncontrollable, or unintended. This makes intuitive sense. You may be frustrated by someone who stumbles and drops a box on your foot, but you wouldn't consider them immoral as it was an accident. Likewise, you couldn't possibly blame a driver who has a seizure for causing an accident with another car as the situation was out of their control. Finally, you wouldn't consider someone immoral if, when trying to comfort you, they hurt your feelings more; they were just trying to help. Likewise, in the opposite direction, we wouldn't attribute good results to someone who caused them accidentally, unintentionally (as a result of some other action), or in circumstances that were out of control. This new model of free will provides a framework for recognizing when an action is legitimately unintentional, or uncontrollable, and when it was under the agent's control. Of course, a great deal of empirical studies will need to follow up to flush out the mechanisms, and improve our ability to detect those mechanisms in action, but this is an important step in the philosophical approach.

Some philosophers have addressed the blameworthiness of psychopaths who are capable of deliberation and rational thought, thereby placing their actions in the plane of moral consideration, but who are completely blind to moral codes by no fault of their own, which would seem to render them outside of the moral plane (Talbert, 2008; Watson, 2013). I would suggest that further empirical studies subsequent to my model of free will could help establish if cases of psychopathy, and other mental disorders, warrant blameworthiness.

Some moral philosophies do not require intention, consciousness, or control for an action to enter the moral plane. This reminds us of the relativism of morality, which lends strength to my argument that the perception of morality, responsibility, and personal culpability relies heavily on someone's upbringing. Against Wegner and other illusionists, if we are raised believing that free will is an illusion, it would impact the way we conceive of morality. In a prime example, one study has found that belief in free will directly causes increased feelings of gratitude, as believing in our free will and the free will of others influences how much we consider acts to be sincerely motivated (Mackenzie et al., 2014).

It would be wrong of me not to also acknowledge the ways in which moral relativism opposes my model for free will. According to my model, free will is found in an agent that can be strictly understood to be an individual human. This is because I personally consider the individual person to be the right combination of complex processes that make up a self-sufficient and free willing body, and yet be small enough to represent the basic unit of free will. Many non-eurocentric/non-Western cultures combine the concept of individual free will with the concept of communal free will. For example, there are traditional African philosophies, such as those of the Akan people in current day Ghana and Ivory Coast, where the community is considered the fundamental, basic unit while individuals were considered not to be self-sufficient agents. In contemporary African moral philosophy there is debate about whether those traditional values provide a moral framework that can effectively resolve contemporary conflicts, or if an individualistic moral philosophy more accurately defines where free will is bound, and therefore how to best allocate moral responsibility (Oyowe, 2013).

Community based moral systems where the individual is not considered the basic unit stand in odd contrast to my assumption that the individual is the seat of free will and moral responsibility. There is something to be said for the fact that, without a community, there is no concept of responsibility, punishment and reward, legal culpability, and even the reactive attitudes that P.F. Strawson argues are more important in the schema of moral responsibility than the causal background of an action (Strawson, 1962). At least, so we can expect; there aren't

enough cases of people living in absolute exclusion from communities who can report their experiences. As my thesis does not seek to adjudicate on moral relativism vs objectivism, I will leave this question to someone else.

My model of free will uses the in-depth, combined perspectives of several leading neuroscientists (each of whose proposals stands on the shoulders of extensive scientific literature, so it isn't really just the position of a few neuroscientists) to provide a substantive empirical background to my philosophical model. Moving forward I suggest that further empirical studies can pick apart the mechanisms involved in the version of free will that is roughly outlined here. In so doing we could establish ways to distinguish between acts that are freely willed and those that are not. I do not suggest that my model will discover new moral codes, (i.e. we won't suddenly find out that, yes, we do have control over when we drop heavy objects), but rather that we might learn to distinguish between when we have control over, intention toward, or awareness of the outcome in a given situation, and when we don't. This could help us to elucidate whether an action in some grey area does warrant consideration in the moral plane. For example, it is hard to say whether I should be blamed for being standoffish when it can be argued that my abrasive qualities are entirely unintentional and unconscious. On the one hand, I am not actively motivated to be standoffish, I just come off that way (in this hypothetical world). On the other hand, had I undergone careful deliberation, maybe I could have changed my conscious motivations in a way that impacted my future unconscious activity, thereby preventing my appearing abrasive, even when I wasn't thinking about it. When we know more empirically about free will, we could better determine if I have control over my standoffishness, and should therefore be considered in the moral plane, or not, in which case my trespasses can slip below the radar as an accident.

Legal Repercussions

Despite clarifying some questions of blameworthiness, the problem of how to assign punishment or reward on the basis of actual culpability is still thoroughly complex. As such, how

does this proposal change our system of law and order? In "Incognito: The Secret Lives of the Brain," David Eagleman addresses the wonder of the subconscious brain as it relates to almost every component of our lives. In particular, Eagleman raises the question of blameworthiness for subconscious actions and its impact in our current legal system.

When confronted with the question of free will, Eagleman is quick to support an eliminitive compatibilist position much like that of Wegner. He holds that the brain strictly follows laws of determinism while simultaneously agreeing that such a state isn't compatible with free will, thereby concluding that free will, as we think of it, does not exist beyond illusion. At one point, Eagleman acknowledges the argument about quantum indeterminacy, but he draws the conclusion that: "a system that is probabilistic and unpredictable is every bit as unsatisfying as a system that is deterministic, because in both cases there is no choice. It's either coin flips or billiard balls but neither case equates to freedom in the sense that we'd desire to have it" (168). Eagleman devotes exactly a paragraph in his book to acknowledging indeterminism, and concludes that indeterminism cannot provide the answer. As I have shown, when we consider the non-probabilistic ends of indeterminism, it defies Eagleman's terse treatment.

As we have seen, the structure of the brain as it has evolved is not only capable of amplifying quantum level indeterminacy but also of wielding it in a way that is not probabilistic after a point. This suggests that indeterminacy in the brain is neither, necessarily, superseded by the larger magnitudes of structure, and nor is it strictly probabilistic, or random, and therefore incompatible with free will. Klemm supported this view. He predicts the conclusion that Eagleman draws: "the doctrine of determinism states that every fact in the universe is guided entirely by physical law. Quantum mechanics also posits that nature is random until observed or measured. Neither facet provides much room for free will in human decisions. Fatalism and destiny governed by probabilities are corollaries" (195). However, Klemm goes on to say that "the equations of quantum mechanics do not predict what will happen, but rather the probability of what will happen." This opposition to Eagleman-esque conclusions about quantum mechanics supports my theory that while starting out as quantum uncertainty, indeterminism in the brain

need not be entirely random. Of course, Klemm writes further that "human behavior is much more complex and subject to numerous variables have nothing to do with quantum mechanics. Free will is greatly removed from the activity of subatomic particles." This would seem to detract from a theory proposing the import of quantum mechanics, but I read this to mean that quantum uncertainty does not govern free will, while not ruling out the fact that it is a factor.

Indeed, he does support the idea that "a more practical criterion for free will is that of 'could have done otherwise'" (195). A 'could have done otherwise' clause invokes a very similar argument to that of the libertarians who suggest that alternative possibilities provide a gamut of choices from which a free willing agent can opt for their preferred action. Recall that Klemm supports the idea that "when decisions are performed by the conscious mind, the avatar can differentially select what inputs from stimuli and memory to consider, thus ultimately biasing one option over others" (195). The concept of an avatar choosing what considerations to select when choosing from an array of options is directly comparable to my definition of free will as the act of self-generating a motivation wherein, during the process of deliberation, our consciousness defines the rules for calculating the pros and cons of each choice.

Aside from the fact that we disagree on what free will is, or can be, and whether or not it exists beyond illusion, I also find Eagleman's address of blameworthiness and law to be incomplete. I truly admire his work, let me be clear. Incognito is a fascinating book and I agree with a great deal of it, almost all of it, except for the admittedly fundamental premise that free will is an illusion. I extend this same belief and respect to Wegner. I believe that the unconscious brain is a vast part of our lives, and I agree that, to a large extent, we are deterministic creatures. However, I disagree that all of Eagleman's and Wegner's evidence means there is not a larger, guiding capability in humans, and perhaps other animals (it's too soon to tell), that maintains and is responsible for a majority of our unconscious action. Discussing Eagleman's theories of blameworthiness and law will be an example of how I can agree with most of his conclusions, while disagreeing with the premise and finding the conclusions to be incomplete.

After concluding that humans are strictly determined, Eagleman takes the stance that our thoughts, decisions, and actions are the consequence of our external influence, our biological makeup, and any relevant pathologies. He notes that "responsibility for your actions parallels volitional control," which is to say that we would need to have free will and need to be able to control what choices we make if we are to be held responsible for those actions. He notes that free will is an open scientific question but goes on to say that "the answer to the question of free will doesn't matter– At least not for the purposes of social policy…" (170). Eagleman develops a theory he calls the principle of sufficient automatism. He claims that, "taking into account the steering power of our genetics, childhood experiences, environmental toxins, hormones, neurotransmitters, and neural circuitry, enough of our decisions are beyond our explicit control" which, to him, means that "free will may exist – but if it does it has very little room in which to operate… the principal states that the answer to the free-will question simply does not matter" (170). Reasserting Klemm's objection that most genetic expression is epigenetic, and therefore modified by experiences we choose, and Tse's argument that automatic neural mechanisms can be criterially adjusted by consciousness, much of what Eagleman claims is beyond explicit control is actually well within the control of a free willing agent.

The question of what is possible without consciousness is an old one. A classic formulation for calling into question the essential nature of conscious intervention in decision-making is the zombie argument. Imagine a creature, the zombie, with every capacity of a normal human except conscious experience. The claim of the zombie problem is that we could imagine such a thing, that conscious experience is a non-causal component of us that can be removed from the equation without disturbing functionality. For Tse, the question doesn't work. Either humans and zombies are completely different, or they are exactly the same down to sharing conscious experience. The two would be different because we've removed "'precompiled' information that endogenous attentional and planning circuitry can access and operate on," which makes it so that the zombies would have "nothing to manipulate in working memory for purposes of planning or internal visualization" (10.83). It could be argued that zombies 'could'

be built with the ability to direct attention, but that the attention would not be manipulating consciously experienced qualia. Tse objects that this is impossible, the point of endogenous attention is that it manipulates consciously experienced qualia. There is no "endogenously 'inaccessible qualia'" (10.83) that can be removed to make this hypothetical zombie. Thus, we cannot debunk the importance of consciousness in decision making using the zombie problem because its presuppositions are incoherent. There is evidence for the fact that endogenous attention is required to perform some tasks in the efficient way that we free willing agents do (Lee, 1999; Carrasco, 2004; Lu, 2004; Chong, 2005; and Busse, 2008, to name only a few examples).

Back to Eagleman, he asserts that assigning blame and our legal system is currently based on what biological factors we are capable of detecting with our technology. The natural conclusion would be, as Eagleman acknowledges, that, "the bottom line of the argument is that a criminal should always be treated as incapable of having acted otherwise. The criminal activity itself should be taken as evidence of brain abnormality, regardless of whether currently measurable problems can be pinpointed" (177). As such, Eagleman argues that blame for a crime or legal culpability, "appears to be the wrong question to ask. Here's the right question: what do we do, moving forward, with an accused criminal?" While I disagree with the premise that blameworthiness, moral responsibility, or legal culpability, is irrelevant because we do not have free will, I agree with the conclusion that the criminal justice system should focus more on what to do with criminals moving forward rather than to throwing them in jail without first trying to help them. As he writes, "a forward thinking legal system will parlay biological understanding into customized rehabilitation…" (180). Agreeing with Eagleman on this point is easy enough. Of course, I believe in mercy, compassion, understanding, and rehabilitation as opposed to the indubitably unfair and cruel criminal justice system in place. I also agree with what he proposes we can do to address this question of what to do moving forward.

Eagleman proposes that "to help a citizen reintegrate into society, the ethical goal is to change him as little as possible to allow his behavior to come into line with society's needs"

(182). The goal is to change the criminal as little as possible in order to avoid the frightening conclusion of science fiction novels such as Clockwork Orange, and 1984. He goes on to say that this can be done by focusing on impulse control. As he writes, "poor impulse control is a hallmark characteristic of the majority of criminals in the prison system" (182). Thus he concludes that the best rehabilitative strategy for criminals is "to give the frontal lobes [involved in long term planning] practice in squelching the short-term [impulse driven] circuits" (183). Our impulses, like the desire for chocolate, money, snacks, alcohol, television, or sex, are kept at bay by our frontal lobes, which suppress this short term activity when keeping in mind larger concepts like social acceptability, discipline, and future plans. Most of us have a temptation that overwhelms the frontal lobes ability to suppress it, but normally acting on these temptations does not invoke legal ramifications. For those of us with especially poor impulse control, Eagleman suggests that these can make for criminal activity and can be corrected through customized impulse control activities, which he describes in more detail in his book. The long and short of it is that I agree with his proposals as a system for rehabilitation, and as an alternative to the counterproductive criminal justice system we currently have in place.

The mistake, or the missing piece to a complete solution, is the lack of attention to the concept of responsibility. At the level of the application of law, personal concepts of responsibility are irrelevant. The law will determine culpability according to its predefined principles. However, moral decision making comes into play in situations that are not resolved by "blind" justice. One example is that of rehabilitation. The goal is to bring the person to a place where they can better understand what their responsibilities to society are, and how they can exercise more control over their actions in order to fulfill their duties. In this case, I believe that it would be unhealthy, and counterproductive, in addition to inaccurate, to instill the belief that conscious control is an illusion and that our actions are ultimately uncontrollable by anything but the automatic, hardwired components of our brain.

Rehabilitation need not only be a cure for criminal behavior, in "The Brain That Changes Itself," psychiatrist and psychoanalyst Norman Doidge discusses cases of patients

with mental disorders and damage from trauma who have been successfully cured by a variety of techniques aimed at 'rewiring' the brain. As he writes extensively, it is becoming clear in contemporary neuroscience that the analogy of the brain as machine has the fundamental flaw of underestimating neuroplasticity, the inherent malleability of neural connections in our brain. Of particular relation to the topic of free will and consciousness, he even demonstrates how imagination and mental practice can optimize the connectivity in our brain in order to perform specific tasks. In essence, he is describing the ability of our conscious mental mechanisms, such as Klemm's CIPs or Tse's endogenous attention and working memory, to sculpt neural pathways. These changes to the brain would not only influence those actions mandated by our consciousness, they would also influence the unconscious processing of information, which would be forced to proceed through the sculpted framework.

Personal Identity

Other examples can be found in daily life. This brings me to my final comments on free will. Given current neuroscience, I believe that we need not rely on Wegner's concept of morality and responsibility, but can instead turn to the theories of practical philosophers such as Korsgaard. Recall from Chapter 1 that, to Wegner, the strength of our illusion is enough to maintain the phenomenology of free will, which in turn maintains moralistic behavior.

As I have said when previously discussing Wegner, beyond my contention that free will isn't an illusion, I hold that the intentional stance would not be infallible to the belief that free will is an illusion. As such, there's no point in drawing such an influential conclusion when the evidence doesn't exist. Wegner makes a strong case for the role of unconscious action through his use of examples from pathologies. I willingly concede that consciousness is not necessary for action. However, the arguments I have presented on the causal role of consciousness do not require that decision-making and action be contingent on conscious intervention. What Ramachandran, Tse, and Klemm suggest is that consciousness optimizes the process of decision-making in order to drive action that is consistent to a single self-identity. Upholding a concept

of self-identity through conscious influence on decision-making has ultimately led to the optimization of interspecies interactions, which have enabled the construction of cooperative efforts, societies, and, at this point, the miraculously interconnected global community.

The fault in Wegner's argument is in understating what he calls phenomenal free will. Instead, Wegner pushes to the fore the strength of the role of unconscious decision-making. Ultimately, it has been shown that phenomenological qualities of consciousness can sculpt the way our automatic cognitive mechanisms function in a way that essentially makes our unconscious decisions an extension of our conscious decisions. Wegner concedes that an ontological free will exists wherein higher level conscious mechanisms influence those lower down, but he maintains that phenomenology cannot have a causal role based on Libet's studies finding that the experience of will comes after the brain decides on an action.

I have shown that contemporary neuroscience since Wegner supports the role of phenomenology in free will. Phenomenology provides the context for the concept of self. Our phenomenological selves experience qualia that are precompiled by our brain, as Tse suggests, and can consciously direct attention to some or other qualia. Directing attention ultimately shapes which qualia are processed with greater priority. This cognitive explanation parallels the practical reasoning explanation wherein reason is a process of deliberation to select the optimal considerations, or motivations, to hold during decision-making.

To Wegner's credit, I find that his reference to the import of authorship is entirely accurate, and distinctly similar to the way that philosophers of practical reasoning describe morality and reason. Through a lens of practical reasoning we can, lastly, look into how this thesis on free will can impact our day to day activities. Practical reasoning provides an explanation of how the deliberative quality of free will enables the sensation of authorship, and how that motivates moral action. Practical reasoning operates under the assumption that the human capacity of reason constitutes free will, which, given the contents of previous chapters, need no longer be an assumption or an illusion. Reason, which is defined in practical reasoning

as the capacity to reflect and deliberate, is essentially the process of our endogenous attention manipulating lower cognitive functions, favoring some patterns of neuronal activity over others, in order to change the criteria or rules for future neuronal activity.

Recall Korsgaard's key quote on practical reasoning and moral decision making: "the reflective structure of human consciousness requires that you identify yourself with some law or principle, which will govern your choices. It requires you to be a law to yourself. And that is the source of normativity" (103-104). This suggests that the way we identify ourselves is vitally important in our day to day activities, which is why I turn back to this in my final notes on the impact of free will. It is important that we uphold an identity that respects personal culpability instead of taking the easy way out and blaming nature and hardwiring. As science enables us to explain behavior in greater and greater detail, there can be a tendency to translate those results into excuses for our own behavior. For example, we sometimes excuse stubbornness or a tendency to yell when we're upset on the way we were raised, stress, too much coffee, or too little coffee. Sometimes we disregard some unconscious social error like invading someone's personal space, or repeatedly saying something demeaning by appealing to obliviousness.

Of course, there are times when appealing to a lack of control is perfectly legitimate. Something as profound as childhood abuse, for example, can have very powerful lasting effects that are out of the control of the agent. When I say that we need to maintain an identity that appreciates our personal culpability, I am not suggesting that we all have to assume a radically individualistic nature where we consider ourselves responsible for every minuscule action we perform. Ultimately, I want to caution us against over-exaggerating those circumstances we do not have control over. Blaming standoffish behavior on stress or upbringing is a cop out, and it undermines what truth there might be in such a statement for someone exposed to constant domestic abuse, who might have a far more compelling reason to be blameless for such behavior.

A final contribution from practical reasoning are the concepts of natural goodness and deep happiness (Foot, 2001). The natural good for humans is based on the fact that we are

reflective, rational beings in the same way that Korsgaard described. We are naturally compelled to do as best fits our natural goodness, which means doing what we feel is best for us. This is not to say that all action is motivated toward furthering the self, but rather the general human good, which Foot argues is what we are all compelled toward in order to achieve deep happiness. Deep happiness goes beyond the superficial happiness we might obtain from selfish actions; it is experienced by one who can "not only see his or her good as bound up with goodness of desire and action, but also feel that it is, with sentiments such as pleasure, pride and honor" (98). Our reflective capacity discourages immoral behavior by denying us access to this deep happiness, which is experienced about objects "that are basic in human life, such as home, and family, and work, and friendship" which can all yield a "shared reaction among human beings to things that are very general in human life" (88). Anyone motivated toward selfish ends will be plagued by self-inflicted guilt and the inability to trust others, which prevents any feeling of deep happiness. As rational agents, we apply practical reasoning and are motivated to act by reasons that we can weigh as more or less important to us. Since we are capable of seeing grounds to perform a variety of actions, we choose one based on what I call our self-generated motivations. Our choice of motivation is what is subject to moral scrutiny.

Knowing the power that conscious will can have over molding plastic activities in the brain empowers us to improve ourselves. Through careful deliberation we can sculpt our moral engagement with other people and the world at large. By consciously acknowledging the needs and desires of others in conjunction with ourselves, and by identifying human goodness through deliberation, we can self-generate motivation to pursue that goodness. For the most part we do not skirt responsibility for our actions. Many people will excuse themselves in front of others and suffer guilt alone. I do not find the experience of culpability to be in immediate danger. However, I do believe that commitment to an identity that promotes human goodness through careful deliberation, and the recognition of our power over our actions will make us better co-workers, friends, lovers, and participants in the global human good.

Bibliography

Banks, G.; Short, P.; Martinez, J.; Latchaw, R.; Ratcliff, G.; Boller, F. "Alien Hand Syndrome: Clinical and Postmortem Findings." Archives of Neurology, 46.4 (1989): 456-459.

Bayne, T. "The Phenomenology of Agency." Philosophy Compass, 3.1 (2008): 182-202.

Busse, Laura, Steffen Katzner, and Stefan Treue. "Temporal Dynamics of Neuronal Modulation during Exogenous and Endogenous Shifts of Visual Attention in Macaque Area MT." PNAS 105.42 (2008): 16380–16385.

Carrasco, Marisa, Sam Ling, and Sarah Read. "Attention Alters Appearance." Nature Neuroscience 7.3 (2004): 1–15.

Clark, A. "Mindware: An Introduction to the Philosophy of Cognitive Science." Oxford: Oxford University Press (2001).

Chong, Sang, Duje Tadin, and Randolph Blake. "Endogenous Attention Prolongs Dominance Durations in Binocular Rivalry." Journal of Vision 5 (2005): 1004–1012.

Dennett, D. "Consciousness Explained." Boston, Back Bay Books (1992).

Doidge, N. "The Brain That Changes Itself." New York: Penguin Books (2007).

Eagleman, D. "Incognito: The Secret Lives of the Brain." New York: Vintage Books (2011).

Foot, P. "Natural Goodness." Oxford: Clarendon Press (2003).

Frankfurt, H. "Freedom of the Will and the Concept of the Person." The Journal of Philosophy, 68.1 (1978): 5-20.

Haggard, P. "Human volition: towards a neuroscience of will." Nature Reviews, 9 (2008): 934-946.

Kane, R. "Torn decisions, luck, and libertarian free will: comments on Balaguer's free will as an

open scientific problem." Philosophical Studies,169 (2014): 51-58.

Kant, I. Trans. and Eds. Mary Gregor and Jens Timmermann. "Groundwork of the Metaphysics of Morals." Cambridge, Cambridge University Text.

Klemm, W. R. "Mental Biology: The New Science of How the Brain and Mind Relate." Amherst: Prometheus Books (2014).

Korsgaard, K. "Sources of Normativity." Cambridge: Cambridge University Press (1996).

Lee, D. "Effects of Exogenous and Endogenous Attention on Visually Guided Hand Movements." Cognitive Brain Research 8 (1999): 143–156.

Libet, B. "Do we have free will?" Journal of Consciousness Studies, 6.8 (1999): 47-57.

Lu, Zhong-lin, Seong-taek Jeon, and Barbara Anne Dosher. "Temporal Tuning Characteristics of the Perceptual Template and Endogenous Cuing of Spatial Attention Q." Vision Research 44 (2004): 1333–1350.

Mele, A. "Effective Intentions: The Power of Conscious Will." Oxford: Oxford University Press (2009).

Mele, A. "Free Will and Neuroscience." Philosophic Exchange, 43.1 (2013): 1-17.

Oyowe, O. "Individual and Community in Contemporary African Moral-Political Philosophy." Philosophia Africana, 15.2 (2013): 117-136.

Perlovsky, L. "Free Will and Advances in Cognitive Science." Open Journal of Philosophy, 2.1 (2012): 32-37.

Ramachandran, V.S. "A Brief Tour of Human Consciousness." New York: Pi Press (2004).

Slingerland, E. "Trying Not To Try: Ancient China, Modern Science, and the Power of Spontaneity." New York: Broadway Books (2014).

Smilansky, S. "Free Will and Illusion." Oxford: Oxford University Press (2000).

Smilansky, S. "Free Will, Fundamental Dualism and the Centrality of Illusion." Oxford Handbook

of Free Will. Ed. Robert Kane. Oxford: Oxford University Press (2002), p. 489-505.

Strawson, P. F. “Freedom and Resentment.” British Academy. N.p., 1962. 1–15.

Strawson, G. “Realistic monism: why physicalism entails panpsychism.” Journal of Consciousness Studies, 13.10 (2006): 3-31.

Talbert, M. “Blame and Responsiveness to Moral Reasons: Are Psychopaths Blameworthy?” Pacific Philosophical Quarterly, 89 (2008): 516-535.

Tse, P. U. “The Neural Basis of Free Will.” Cambridge: MIT Press (2013).

Velleman, J. “Self to Self: Selected Essays.” Cambridge: Cambridge University Press (2005).

Watson, G. “Free Agency.” The Journal of Philosophy, 72.8 (1975): 205-220.

Watson, G. “Psychopathy Agency and Prudential Deficits.” Proceedings of the Aristotelian Society, 113.3 (2013): 269-292.

Wegner, D. “The Illusion of Conscious Will.” Cambridge, The MIT Press (2002).

Wegner, D. “Precis of The Illusion of Conscious Will.” Behavioral and Brain Studies, 27 (2004): 1-45.

Widerker, D., McKenna, M. “Moral Responsibility and Alternative Possibilities: Essays on the Importance of Alternative Possibilities.” Farnham, Ashgate Publishing Company (2006).