

FROM GUILT TO REGRET

The impact of neuroscientific evidence upon our ideas of libertarian free will and moral responsibility

Bachelor Degree Project in Philosophy
15 ECTS
Spring term 2012

Nathalie Hallin

Supervisor: Stefan Berglund
Examiner: Paavo Pylkkänen

**From guilt to regret: The impact of neuroscientific evidence upon our
ideas of libertarian free will and moral responsibility**

Submitted by Nathalie Hallin to the University of Skövde as a final year project towards the degree of B.Sc. in the School of Humanities and Informatics. The project has been supervised by Stefan Berglund.

2012-06-18

I hereby certify that all material in this final year project which is not my own work has been identified and that no work is included for which a degree has already been conferred on me.

Signature: _____

Abstract

In this essay I will investigate some different opinions about how or whether we should change our views concerning moral responsibility if neuroscience shows that contra-causal free will is implausible and also argue that holding people morally responsible is not the best method of making people act morally. A common approach is to argue that we can hold people responsible in a non-absolute sense and many argue that this has practical benefits (i.e. that people would act more morally). I evaluate different alternatives and conclude that there are different beliefs and opinions that cause disagreement in this question. I argue that moral responsibility is implausible without the existence of contra-causal free will (because it is then just a question of luck what actions you will perform) and that it is more practical to shift our focus to other ways of changing people and not blaming them.

Keywords: Moral responsibility, free will, blame

Contents

Abstract	3
1. Introduction	5
1.1 <i>Background</i>	5
1.2 <i>Definitions</i>	7
1.2.1 Free will.....	7
1.2.2 Moral responsibility.....	10
2. Compatibilist views.....	12
3. Incompatibilist views	18
4. Discussion	21
4.1 <i>Summary of the Presented Views</i>	21
4.2 <i>Evaluation of Possible Solutions</i>	24
4.2.1 Free will is not what moral responsibility is founded on	24
4.2.2 Our views are innate and it is more practical to keep them.....	25
4.2.3 We should not hold people who are not in control morally responsible	25
4.2.4 We should not hold people morally responsible	26
5. References	31

1. Introduction

Recent neuroscientific research might undermine the idea of libertarian free will and this might have consequences on our views on moral responsibility. This essay will discuss whether (and in that case how) we should change our views concerning moral responsibility if contra-causal free will does not exist. I will start by presenting some neuroscientific research which gave rise to the discussion about libertarian free will. Then I will define two concepts which are central in this discussion – free will and moral responsibility. I will then discuss answers to this question given by compatibilists and incompatibilists. Finally I will give my own suggestion to how we should change our views on moral responsibility if contra-causal free will does not exist.

1.1 Background

In 1983 Benjamin Libet did a study which suggests that the decision to carry out voluntary acts can be made before reaching consciousness. In his experiment, the participants were told to flex their finger or wrist when they felt like doing so and to remember the exact time, using a very exact clock. The results showed that there was activation in the brain (in the supplementary motor area (SMA)) several milliseconds before the participant thought that he or she began to feel the urge to act (Libet 1983). This shows, Libet argues, that “a spontaneous, freely voluntary act can begin unconsciously, that is, before there is any (at least recallable) subjective awareness that a 'decision' to act has already been initiated cerebrally” (Libet 1983, 623). This study received much critique. Some said that the clock could be distracting and that a conscious decision was too subjective. But many have replicated and refined Libet’s study (K. Smith 2011, 24) and some of these studies will be discussed later in this section. I argue that if decisions can be predicted before the person experiences that he or she decides it consciously, many people would conclude that contra-causal free will is very unlikely. The definition of ‘free will’ will be discussed in the next section.

Soon, Brass, Heinze and Haynes (2008) wanted to examine if there was activation in other parts of the brain, prior to the one found in the SMA. They found activation in frontopolar cortex 7 seconds before the subject’s decision and some later activation in other parts of the brain and if one takes into account the temporal impreciseness of fMRI-scanning, the real time of the activation was up to 10 seconds. This shows that a decision can be shaped by a network of high-

level control areas long before the person becomes consciously aware of it. When it reaches consciousness, the person experiences it as freely chosen (Soon 2008).

According to Churchland (2006), the information we have gained through research is enough to say that it is very improbable that we have a non-physical soul or something which makes it possible for us to make choices independent of causal processes. However, it feels as if we have free will. Churchland gives two explanations to why we think that we have free will. The first is that we cannot access all processes leading to our decisions and therefore it appears to us that they are uncaused. The second is that when our brain decides to perform an action, a copy of the command called 'efference copy' is sent to other regions of the brain and tells us that it was we who performed the action. For instance, when your arm rises in water, an efference copy is not sent which makes me attribute the movement to an external source. In contrast, when you raise your arm in the water an efference copy makes you assume that you chose the action. You may falsely believe that it was uncaused free will which made the arm move when in fact, the difference is that an efference copy is sent in the second case (P. S. Churchland 2006, 5-8).

Mele thinks that it would be a threat to free will if neuroscience could predict a person's decision correctly in every case – before the person was conscious of the decision – just by looking at their brain activity. This may be possible quite soon (K. Smith 2011, 24). Fried, Mukame and Kreiman (2011) used a paradigm similar to the one Libet (1983) used and managed to predict participants' decisions with 80% accuracy and could also predict, with a precision of a few milliseconds, when the decision was going to be carried out (Fried, Mukame och Kreiman 2011).

There is thus a chance that neuroscience will be able to show that contra-causal free will is very improbable and this might make it more commonly believed that contra-causal free will does not exist. This might have consequences for our views on issues related to free will. One such issue is moral responsibility. In case we do not have a contra-causal free will, how should we change our views concerning moral responsibility? Can we still hold people responsible for their actions, and if so, in which way and on what grounds? I will begin by defining the two concepts which are essential to these questions – free will and moral responsibility. Then I will present the different answers that some philosophers, with different opinions and different focus, have given to my questions and finally – in the discussion – I will evaluate the answers and give

my suggestion about how we should change our values concerning moral responsibility if free will does not exist.

1.2 Definitions

1.2.1 Free will

Van Inwagen states that the definition of the “free will thesis” is that people sometimes are in the position when they contemplate a future act that they “simultaneously have both the following abilities: the ability to perform that act and the ability to refrain from performing that act” (van Inwagen 2008, 329). Many people claim that there are different definitions of free will – libertarian free will (genuine, metaphysical free will) and compatibilist free will (a will which does not violate causality) – but van Inwagen argues that there is only one definition. Both compatibilists and incompatibilists mean the same by the word ‘able’, he argues. What it means for both is what the word means in English. The difference is simply that compatibilists think that being able to act or refrain from acting is compatible with determinism while incompatibilists think that it is not. Therefore, they all mean the same by ‘free will’. (van Inwagen 2008, 329-332)

Eddy Nahmias (2011) argues that ordinary people define free will in a different way than the neuroscientific studies. There are two types of control: libertarian powers and compatibilist powers. Libertarian powers require indeterministic gaps where free will can operate (and sometimes also agent causation, i.e. a power of a person to initiate causal processes). Compatibilist powers refer to self-control, rational choice, planning and such things – this can be possible even if determinism is true. The neuroscientific data suggest that free will defined as libertarian powers is implausible, but ordinary people often define free will as compatibilist powers. Even if they understand that it is the libertarian powers which are threatened, some people infer this to be a threat to compatibilist powers, as they falsely think that libertarian powers are required for compatibilist powers (Nahmias 2011, 17-24).

Wegner uses the term *conscious* will, but he seems to mean that what he says should apply for what others call free will, since he quotes and refers to other texts discussing free will. However, he does not explicitly write about how the two terms relate to each other. He says that conscious will has two meanings; (1) the feeling we have when we experience that we knowingly

give rise to an action or thought and (2) a force our minds have, which is the causal link between our minds and an action. The first meaning is essential for saying that anyone practices conscious will – no matter how much the person’s behaviour indicates that he or she acts with his or her conscious will, it still would not make sense to say that it is the case, if the person experiences that he or she did not act with his or her conscious will. However, there are cases when people experience the feeling of conscious will, when they in fact were not the ones who performed an action (e.g. a person suffering from OCD who experiences that his or her compulsive actions keeps horrible events from happening) and there are also cases when people do not experience any control, even though their actions are the cause of an event (e.g. a hypnotised person who is told that his or her arm gets heavy and is drawn toward the ground) (Wegner 2002, 3-9).

The second meaning is often inferred from the first meaning, Wegner argues. The feeling of consciously causing an event can lead one to suppose that one actually does cause an event – but to feel like one is causing something is not the same as the force of causing, says Wegner. The force of will can take different forms; it can be a long-lasting property of a person or small bits of force which come sporadically to produce individual acts. The will can be strong or weak. The traditional way of viewing the force of will is as something contra-causal and thus no actions can be predicted and it is meaningless to study the will scientifically. Wegner argues that the conscious will is an illusion and by that he means that the experience of will is an illusion of the force of will (Wegner 2002, 12-15).

According to Farthing (2002), it has been said that free will refers to the concept of volition – that we are able to choose from different alternatives of actions and that it is the opposite of determinism. But Farthing agrees with Hebb when he claims that free will is “control of behavior by the thought process” (Farthing 2002, 40) and that it doesn’t require indeterminism and violation of scientific law. In fact, it is incompatible with indeterminism – which actually means randomness. We are free to decide our actions within the limits of determinism (Farthing 2002, 40).

Nagel (2005) argues that free will cannot exist because of moral luck. Ultimately, everything is about luck – even who you are and what actions you carry out: “If one cannot be responsible for consequences of one’s acts due to factors beyond one’s control, or for antecedents of one’s acts that are properties of temperament not subject to one’s will, or for the circumstances

that pose one's moral choices, then how can one be responsible even for the stripped-down acts of the will itself, if *they* are the product of antecedent circumstances outside of the will's control?" (Nagel 2005, 209) Nagel also gives an explanation of why we act as if free will did exist. We cannot view ourselves simply as a part of the world – we have a line between ourselves and things which are not ourselves – acts that we “cause” as opposed to things that happen to us. We impose this on other people as well – we assume that it is for them as it is for us. (Nagel 2005, 209-212).

It seems that there is some disagreement about how free will should be defined and this must be kept in mind when examining the answers to my initial question. However, two definitions seem to be dominant. The first is that free will requires genuine, contra-causal, metaphysical (i.e. objective and ontologically real) agency, which the incompatibilists seem to agree about. The second is what Nahmias calls compatibilist powers, which compatibilists seem to agree about. I do not think that people generally define free will as compatibilist powers and falsely believe that contra-causal free will is required for that. According to my experience, people generally define free will as contra-causal and then fail to see the problems with this concept (which will be discussed later). One example is that people generally accept that a person did not act freely if he or she had a mental illness or brain damage which made their action the only possible alternative. If neuroscience will be able to predict decisions perfectly (especially more advanced decisions than the one in Libet's study), this will make people question their belief in contra-causal free will and eventually, this belief will seem less likely by many people.

Even though I am already convinced that contra-causal free will is not possible, I think it is useful to show with neuroscientific studies that contra-causal free will is improbable. People generally respect and trust scientific findings and even if findings suggesting that contra-causal free will does not exist would be considered undesirable by many people, I think that they eventually would accept them. Although contra-causal free will implies that there were not any causal antecedents, it also implies that this free will would causally affect the future. Thus one cannot argue that this free will cannot be detected in studies investigating causal connections. One could argue that the contra-causal free will is unconscious, but most people would not think that it is they who are choosing freely if this happens unconsciously. I argue that neuroscience will affect people's beliefs about contra-causal free will and when more people will believe that

we lack a contra-causal free will, this will have effects on how people view issues related to free will – such as moral responsibility.

The neuroscientific studies do not undermine compatibilist free will and furthermore, people do not generally define free will in this sense. For these reasons, I will define free will as contra-causal free will. The question is thus what the consequences on moral responsibility are if free will defined as genuine agency does not exist. I believe that this type of free will is incoherent, because it seems impossible to imagine what *contra-causal* mean if it would not mean *chance* (see the first paragraph in the section *Compatibilist views*). However, not everyone is convinced by that argument and it is therefore meaningful to use future neuroscientific studies to show that this kind of free will is implausible.

Some might argue that quantum physics opens up a possibility for contra-causal free will, since it seems that some events at the quantum level are not determined. Churchland (2006) presents two arguments against this idea. Firstly, we would not consider uncaused actions free, since that would mean that they are random. This argument was originally presented by Hume and will be discussed later. Secondly, classical physics seems to be able to explain the behaviour of neurons and how they interact with each other. If uncaused events exist in neurons, they are at a very low level and are not related to the events which result in choices (P. S. Churchland 2006, 8). I believe that the first of these arguments is enough to dismiss the idea of contra-causal free will based on quantum physics. The second argument requires much knowledge about how physics works and it is therefore difficult for me to evaluate it. However, if the claims this argument makes are true, this argument is also enough to dismiss contra-causal free will based on quantum physics.

1.2.2 Moral responsibility

Van Inwagen argues that moral responsibility requires the ability to act differently from the alternative we choose. He states that the existence of this ability is obvious: “It is, however, undeniable that people do not always behave as they ought. (...) Denying that agents are ever able to do otherwise is therefore simply not an option” (van Inwagen 2008, 340) (van Inwagen 2008, 339-340). Mackie (1990) does not define the term, but he holds that the question whether we should hold people morally responsible is a moral question and moral questions do not have absolute answers, according to him (Mackie 1990, 208). Mele (2008) states that an agent is

responsible when he deserves credit for a good intentional action or blame for a bad intentional action (Mele 2008, 263).

According to Levy and McKenna (2009), free will is often seen as a necessary condition for moral responsibility (Levy and McKenna 2009, 98, 115). The principle of alternate possibilities (PAP) states that “a person is morally responsible for what he has done only if he could have done otherwise” (Frankfurt 1969, 829) Frankfurt was the first to define PAP, but he does not agree with it. He argues that one can be morally responsible even if one could not do otherwise, if the reason one acts in one way is not the same reason one cannot act otherwise. If you think you have two actions to choose from and you chose one of them, you are still responsible for that decision if you in fact could not have chosen the other alternative (if something else would have prevented you from carrying out that decision) (Frankfurt 1969, 829-830, 832). There has been much debate over whether Frankfurt’s argument is right (Levy and McKenna 2009, 98-99). I think that he assumes that compatibilism is true in his argument. I do not agree with Frankfurt when he states that a person could have done otherwise when he chooses between two alternatives and consequently, the rest of the argument does not make sense. The argument tries to appeal to our intuitions about who are responsible, but our unexamined intuitions (i.e. intuitions which are not thought through) are not very reliable.

Proponents of the *direct argument*, for instance van Inwagen (1984), think that since no one is responsible for the past, the laws of nature and the truth of determinism, no one can be responsible for anything. Widerker (2002) thinks that people can be at least partly responsible, even if determinism is true. Another similar argument or view is the *ultimacy argument* – that you must be the cause of yourself to be responsible, which is impossible. You cannot be held responsible for the way you are, since you did not decide how you are and that is what determines your actions. Clarke (2005, 13-24) questions the assumption that you have to be responsible for the *reasons* for your actions. Vargas (2004, 229) says that maybe that is the folk psychology definition of moral responsibility, but it contains a conceptual error and therefore a new definition of the term is needed. Vihvelin (2007, 310-311) holds that we can influence ourselves in modest ways and in that way become responsible.

One view of moral responsibility is that one is not responsible for actions performed as a result of manipulation. However, Pereboom (2001) says that natural causes are no different from

manipulation by a person and consequently, no one is ever responsible. Lack of knowledge or control is often assumed to excuse a person of responsibility. However, some hold that control is not a necessary condition for moral responsibility – all which is needed is that an action is appropriately attributable to a person, which means that the action expresses the person's attitudes and values. Levy (2005) objects to this view, since it does not separate bad agents and blameworthy agents – an agent can perform a bad action but not deserve blame for it. However, not all agree that this distinction should be made. Angela Smith (2008) argues that a bad agent who does not deserve blame is reduced to a passive victim instead of a moral agent (Levy and McKenna 2009, 115-118).

There is much disagreement about what is needed for holding a person responsible. However, all views I have mentioned seem to agree about Mele's elementary definition – that someone is responsible if he deserves blame for a bad intentional action or credit for a good intentional action. The disagreement seems to be about *when* a person deserves blame or credit – if indeterminism is needed, if there is an objective answer or if we need to come to a practical agreement. The disagreement does not, according to me, lie in the definition of moral responsibility – whether indeterminism is needed for moral responsibility or if there is an objective answer are questions about whether those things follow from Mele's definition. For this reason, and because of the consensus about this definition, I will use Mele's definition of moral responsibility in this essay.

2. Compatibilist views

David Hume argues that an action which is not determined by reasons is not free – if it lacks reasons, it is just the product of chance. We hold people responsible for their actions if they had reasons for performing them and we would not hold a person responsible if the action was random, since it then lacks an appropriate connection to the person. Hume states that “[a]ctions are by their very nature temporary and perishing; and where they proceed not from some cause in the character and disposition of the person, who perform'd them, they infix not themselves upon him, and can neither redound to his honour, if good, nor infamy, if evil” (Hume 1739, 282). People often think that determinism makes moral responsibility impossible – when actually, it is impossible to have moral responsibility without determinism (Hume 1739, 277-282). This is a very good argument against moral responsibility based on libertarian free will. It shows that

libertarian free will is a vague concept and it is hard to say what it would mean if one would try to argue that it is not randomness. The concept is incoherent and that is what Hume shows with this argument. He is correct when he states that we hold people responsible only for actions which have reasons.

However, this does not have to mean that we *should* hold people responsible for actions that have reasons. This does not follow from this argument. Further arguments have to be presented for the position that we should continue to act as we do now. It seems like he suggests that we should follow our intuitions and that our intuitions tells us that we should hold people who act because of reasons morally responsible. However, I do not share these intuitions and furthermore, there is no reason to believe that our unexamined intuitions show what we should do. Although they might show what was evolutionarily adaptive, that does not have to be the same thing as what we should do.

Fischer (2006) thinks that moral responsibility is possible without free will. He presents two freedom-relevant conditions for moral responsibility – two kinds of control: (1) regulative control, which is genuine metaphysical access to alternative possibilities and (2) guidance control, which is a distinctive kind of actual-sequence control, to guide ones behaviour in a certain way. This includes mechanism-ownership and moderate reasons-responsiveness. An agent has mechanism-ownership when he takes responsibility. This means to obtain the belief that one can affect the environment differently depending on one's choices and bodily movements and that one's choices and bodily movements – and their effect on the environment – can get morally charged reactions. This is when an agent makes certain actions his or her own. Moderate reasons-responsiveness means that an agent's mechanism is moderately responsive to reasons in a dispositional way.

He argues that moral responsibility only requires guidance control and not regulative control, which is the only control that disappears with free will. He says that he here provides us with more of intuitive considerations than an argument. An agent becomes responsible by *taking* responsibility (mechanism-ownership) and being appropriately responsible to reasons (moderate reasons-responsiveness) (Fischer 2006). Thus both Hume and Fischer have responsiveness to reasons as a condition for moral responsibility, although Fischer has more requirements. In my opinion, it is obvious that this is based on Fischer's intuitions, as he also admits, and it does not

function well as an argument. I do not think that unexamined intuitions can give a reliable answer to this question. There is no reason for us to believe that our intuitions should point to what is true (even if there is some reason to believe that our intuitions should show what is evolutionarily advantageous, but that is not necessarily the same). However, this is a good formulation of the situations in which people normally ascribe responsibility to other people and their selves. If Fischer wants to argue that his intuition shows what is true about responsibility, he has to present arguments for that position. Right now, it is a position which appears to appeal only to other people's intuitions.

Mackie presents three questions, which he wants to keep separate, when considering moral responsibility: (1) can we make a systematic distinction between intended and unintended actions and between voluntary and non-voluntary actions? (2) How should we ascribe legal and moral responsibility, i.e. credit and blame? (3) When are rewards and punishments appropriate? These questions have two levels: the level of common and legal thought and the deeper, metaphysical level which deals with the issues of free will and determinism.

Mackie argues that an action is intended when the agent who performs the action is not physically forced. There are cases when the agent is subjected to duress, dangers and temptations, which makes the definition more complicated, but Mackie insists that it is more misleading to say that such actions were not intended. The concept of volition is broader than intention – it is simply to perform (or refrain from performing) an action. These concepts are factual, but the question about moral responsibility is a moral question – and legal responsibility is a legal question. Mackie holds that morality must be created, because there is no correct way to act on a metaphysical level and thus there is no absolute answer to moral questions, such as the question about moral responsibility.

He suggests that we follow what he calls the straight rule, which is what we, according to Mackie, normally feel intuitively is right. The idea is that only intended actions ought to be punished. A moral objectivist might want to accept the straight rule because he feels that it is just. However, Mackie is not a moral objectivist and his reason for adopting the straight rule is that an action which was due to physical compulsion is not an action by the agent at all; a consequence which was due to lack of skill, was not really an action of a conscious agent; an ignorant action is

not either an action of a conscious agent, but rather imposed on the course of action of the agent, since he or she was unaware of it.

Mackie does not explain why people *are* responsible when an action *is* intended. However, he does explain why we in addition to simply being able to say that actions belong to agents, also should blame, praise, reward and punish them. He means that by doing so, we modify an agent's view of possible actions and attach a favourableness or adverseness to the action. This is still a created moral – not an absolute. There are many complications with the straight rule, for instance the case of psychopaths, who do not have the ability to appreciate moral feelings or children, whose ability to take responsibility is not fully developed. But the answers that can be given to those objections are still not absolute – it is something we must discuss. Mackie's answer to my question is that moral responsibility is possible without free will – but moral responsibility is not possible in an absolute sense (Mackie 1990, 203-215; 226).

Dennett (1984) holds that free will does not have to include contra-causality and our usual concerns about what we lose with that kind of free will are exaggerated and are a result of incorrect logic in thought experiments. The kind of free will we ought to strive after is one which allows us to keep our dignity and our responsibility. Similarly to Mackie, he does not focus on a metaphysical basis of moral responsibility, but rather the practical benefits of it. To have a system of responsibility has many social and societal benefits. Dennett mentions deterrence used to reduce crimes. He also says that our distinction between responsibility, diminished responsibility and no responsibility is coherent, empirically real and important, in the sense that it increases our feeling of quality and meaning of life. Even if people do not have metaphysical free will, they still accept being held responsible for their actions and attributing responsibility to people can motivate them to let go of an undesired trait. But what about the responsibility we attribute to ourselves? According to Dennett, we cannot have some before-the-eyes-of-God-responsibility, but some weaker senses of responsibility are still possible. We can hold ourselves legally responsible and socially responsible (in the sense that we follow the moral rules we agree about). This form of responsibility does not make remorse and regret impossible. In fact, the thing we want to achieve with remorse and regret is to motivate a person to change his or her priorities for decisions (to prioritise a more desirable action) – and this is what this kind of responsibility leads

to (Dennett 1984, 153-167). Like Mackie, he wants to use our moral emotions to achieve practical benefits.

I do not think that the distinction between different degrees of responsibility increases my feeling of quality and meaning of life and I do not think that a person gets more motivated to let go of an undesired trait if he or she is blamed (as opposed to realising that other people discourage the action or that performing the action has consequences). I will explain this further in the discussion.

Churchland (2011) holds that morality is an evolutionary adaptation shaped by our caring for kin and kith, our recognition of others' psychological states, our ability to solve social problems and our learnt social practices. A result of how our morality is evolutionarily shaped, we create certain values and if we consider these values, some solutions to social problems are better than others (P. S. Churchland 2011, 8-9). Churchland argues that the foundation of morality can only be natural and she attempts to reveal this foundation by investigating how our brain works and what the adaptive function of morality is. Some have criticised attempts to find a natural foundation for morality by claiming that one cannot derive an 'ought' from an 'is', but Churchland consider this an unreasonable demand. An argument does not have to show that something follows logically from it, in order to be sound. She argues that knowledge about how human beings function can help us to better answer moral questions (P. S. Churchland 2011, 5-8; 204). She states that selective punishment probably is the reason 'cheaters' (i.e. people who take advantage of other people's trust instead of cooperating) are not dominating populations. Children learn what social practices are permitted by feeling pleasure when behaviour is approved and pain when it is disapproved. These emotions are very strong and consequently, the different actions associated with these feelings have strong emotional valence. This reward and punishment system helps us to develop a conscience, but it also hinders moral values to change fast, when this is needed (P. S. Churchland 2011, 81; 130-132).

Churchland (2006) mentions Hume and argues like him that a choice which is not caused by desires, drives, hopes, fears, beliefs, intentions and motives, cannot be called a choice at all. We hold (and should hold) people responsible for actions that *are* caused by these things – not for uncaused, random actions. Contra-causal free will is not needed for holding people responsible. Like Mackie and Dennett, she argues for the practical benefits of holding people responsible for

their actions. We are social animals who benefit from holding people responsible. It helps us to keep our society well-functioning. The practical results of holding people responsible are that we can remove dangerous individuals from society, deter potential criminals and have a formal structure for revenging (in order to prevent people from revenging unjust actions themselves).

Churchland (2006) wants to make a distinction between voluntary and involuntary actions, although she has a different approach than Mackie. She wants to use neuroscientific information (which will be acquired in future studies) to differentiate a brain ‘in control’ and a brain ‘not in control’. To be in control means roughly to be able to inhibit inappropriate impulses, maintain goals, consider and evaluate consequences of a planned action, balance long- and short-term values and to resist being carried away by emotions. There are certain structures in the brain which are important in maintaining control (for instance, orbitofrontal cortex, cingulate cortex and amygdala) and there are six *non-specific neurotransmitter projection systems* which could influence control. Different hormone levels could also influence control. Churchland wants to identify all parameters which could influence the level of control and let each represent a dimension in an n -dimensional parameter space. Since it is only possible to present a parameter space with three dimensions visually, we cannot present this visually. However, it can be useful to think of it simplified as a three-dimensional cube. If a brain is within a certain volume, it is in control – and if it is outside that volume, it is not in control. There are thus many ways to be in control. There is no absolute answer to where the limit goes and we must therefore reason together where to draw the line. This information can then be used in legal contexts and according to Churchland this matches our intuitions about moral responsibility. (P. S. Churchland 2006, 10-15).

In her book from 2011, Churchland does not write specifically about moral responsibility, but rather generally about morality, but it seems very likely that she argues that moral responsibility has an evolutionary basis. She seems to hold that we should punish the people we hold morally responsible (who are in control) and refrain from punishing people who are not in control. Does she mean that we should not do anything about the people who are not in control? Probably not – I would suppose that she would like to give those people treatment. She does not make it clear whether she argues that this is a hypothesis about what would be most practical (i.e. would function best in our society) or about what our intuitions about moral responsibility say

(and if she would then argue that we should follow our intuitions because it has been evolutionarily adaptive). I do not think that it is the most practical alternative if she argues that holding people responsible implies that we should blame them. I will explain why I think so in the discussion of this essay. If she means that we should hold certain people responsible because it has been adaptive, I wonder why this trait is different from other traits which we do not think are preferable anymore. For instance, we do not think it is preferable to eat as much fat and sugar as possible, even though this has been evolutionarily adaptive and similarly, we do not think it is good to kill, abuse or bully other people even though this, too, has sometimes been evolutionarily adaptive when human beings evolved. She would probably say that it is the most practical alternative *because* of the way human beings function – that our brain works in a certain way and that our evolutionary history has shaped us in a certain way and therefore shows what is more practical. I disagree with this idea and will consider this more in the discussion.

3. Incompatibilist views

Van Inwagen presents three principles which state when a person is morally responsible: (1) “A person is morally responsible for failing to perform a given act only if he could have performed that act” (van Inwagen 1984, 180), (2) “A person is morally responsible for a certain event-particular only if he could have prevented it” (van Inwagen 1984, 180) and (3) “A person is morally responsible for a certain state of affairs only if (that state of affairs obtain and) he could have prevented it from obtaining” (van Inwagen 1984, 180). He argues that these principles, together with two premises, entail that moral responsibility requires free will.

The first premise clarifies under what circumstances moral responsibility would not exist: “If (i) no one is morally responsible for having failed to perform any act, *and* (ii) no one is morally responsible for any event, *and* (iii) no one is morally responsible for any state of affairs, then there is no such thing as moral responsibility” (van Inwagen 1984, 181). The second premise clarifies the conditions for free will to exist: “If (i) someone could have performed (i.e. had it within his power to perform) some act he did not in fact perform, *or* (ii) someone could have prevented some event that in fact occurred, *or* (iii) someone could have prevented some state of affairs that in fact obtains, then the free-will thesis is true” (van Inwagen 1984, 182). In this text, he uses the term “could have”, however, in van Inwagen (2008), he states that “able to” is a better term to use.

These premises, together with the three principles, entail that the non-existence of free will would make moral responsibility impossible. If moral responsibility does not exist, we would have to stop making judgements about people's actions like "You'd think a person with her advantages would know better than that" (van Inwagen 1984, 207). Van Inwagen rejects the explanation that you are judging the action but not the person and argues that it is not consistent to say that an action was despicable but the person was not morally responsible. If we found out that there were justifying reasons for the person to perform the despicable action (for instance, that the person had been drugged), we would no longer call the action despicable. He also argues that to state that an action was despicable is to describe it just as objectively as stating that a certain car is dangerous (van Inwagen 1984, 207-209).

I do not think that the terms "could have" and "able to" are perfectly clear. People have different opinions about what it means to be able to. I would interpret it as contra-causal free will and I agree with van Inwagen's conclusion that this means that moral responsibility is not possible without free will. However, I do not agree that this means that you cannot say that an action is despicable. What you cannot do is to claim that a person should have known better (because obviously, he or she did in fact not know better). I would not call an action despicable if the person was drugged. In that example, the person does not run an increased risk of performing similar actions in the future and I think that is what we think is important. The actions we call despicable (or make similar judgements about) are actions which are more strongly connected to the person who performed the action, i.e. when the person runs an increased risk of performing a similar action in the future. These are also the actions we hold people responsible for (at least roughly these actions). But I argue that it is possible to separate these two – to say that the action was despicable but to refrain from holding the person responsible. To say that an action is despicable *is* to say something about the action and not about the person and it is not a very good argument to say that in all cases we make judgements about actions we also normally hold the people who performed the actions responsible and the two are therefore not separable.

Nagel also argues that moral responsibility is impossible without free will. However, his reason for holding this is that without free will, there is nothing left which can be said to belong to a responsible self if we do not decide our actions. Genuine agency and moral judgements are not possible if free will does not exist. Events can be celebrated or deplored, but never praised or

blamed (Nagel 2005, 209-212). He sees this as a problem, although he has no solution: “The problem of moral luck cannot be understood without an account of the internal conception of agency and its special connection with the moral attitudes as opposed to other types of value. I do not have such an account. The degree to which the problem has a solution can be determined only by seeing whether in some degree the incompatibility between this conception and the various ways in which we do not control what we do is only apparent. I have nothing to offer on that topic either.” (Nagel 2005, 211-212)

Wegner argues that there are some contradictions between how we think about moral responsibility and the possibility that conscious will is an illusion, but the changes he proposes are relatively small compared to other deterministic theories (Wegner 2002, 334). He says that our actions do not have to be caused by some conscious will in order to be relevant to morality. All that is needed is that we have conscious thoughts about our actions in order for us and others to predict our future behaviour. Hence he emphasises the practical benefits like Mackie, Dennett and Churchland and like Hume and Fischer, he argues for the importance of an appropriate connection to the agent. Wegner mentions Monahan and agrees with him that people do not have to have knowledge about why they act in order for us to determine that they have done something wrong, say, a crime, what treatment they should get, how big the risk is that the person would commit crimes in the future and if this risk can be reduced. The function of conscious will, Wegner claims, is to be a moral compass and it is necessary for producing feelings of guilt and pride. Even if conscious will is an illusion, it is an important feeling which guides our moral behaviour. Unfortunately, Wegner does not propose any concrete changes to the way we think about moral responsibility now (Wegner 2002, 341-342).

4. Discussion

4.1 Summary of the Presented Views

Hume and Churchland argue that actions which are not caused by reasons are random and this is not what we consider free will. They, and also Fischer, argue that we should only hold people responsible for actions which result from reasons connected to the person. Fischer further argues that moral responsibility requires that one takes responsibility, that one is aware that some actions can get morally charged reactions and that one's actions can affect the environment. Mackie argues that when we should hold people responsible is a moral question and that we should follow the straight rule, which is what we most often feel is intuitively right – that we should be held responsible for intended actions. Mackie and Dennett hold that inducing moral feelings like remorse by blaming can make people change their priorities of actions. Churchland, Dennett and Wegner argue that holding people morally responsible has practical benefits. Churchland wants to make a distinction between people who are in control and people who are not in control. People who are in control should, according to her, be held responsible. Van Inwagen and Nagel hold that libertarian free will seems to make moral responsibility impossible.

I agree with Nagel and van Inwagen when they state that the nonexistence of free will is incompatible with moral responsibility. Nagel's linkage to moral luck explains well why: you cannot choose what genes you are going to get and what events you will experience and get shaped by and if this is what ultimately determine your choices, then you cannot be held responsible for them. Also, if chance is what your choices are based on, you cannot be held responsible for that either. But if we cannot hold people responsible, we have a very serious practical problem. This might tempt some people to search for a different conclusion.

Wegner's suggestion is not very concrete, but he seems to argue for a weak form of moral responsibility, which has a practical benefit. The points he agrees with Monahan about (that criminals should be treated and that one should try to reduce the risk of future crimes by the criminal) suggest that moral responsibility should be more about helping people to act in another way in the future, which I also agree about, but he never states explicitly that this is what he

thinks. If he thinks this, I am not sure if ‘moral responsibility’ is the right term to use, since it means that a person should be blamed or praised.

Both Churchland and Dennett argue for the social and practical benefits of moral responsibility and they seem to have rather similar opinions – except for Churchland’s distinction between a brain in control and a brain not in control, which Dennett does not propose. Churchland and Dennett’s suggestion is compelling, because it seems to have adequate grounds for just keeping our intuitive sense of moral responsibility, even when our initial ground for it – free will – is gone. However, I do not think that their suggestions would lead to the most well-functioning society. I will discuss this more later.

Mackie makes it clear that he realises that moral responsibility is not metaphysically possible. He argues, and I agree, that moral responsibility is a moral question, to which there is no absolute answer. But his suggestion is not the most practical (i.e. what makes our society function best and what makes people feel good), although that is what he tries to argue. Of course it is more practical than letting everyone act as they wish, without any consequence, but there are other alternatives, which I will discuss later.

Fischer does not realise that moral responsibility without free will is not possible in an absolute sense, which I think is obvious because of the fact that you cannot choose what genes you are going to get and what events you will experience and get shaped by and if these are what ultimately determine your choices, then you cannot be held responsible for them.

I will argue that the non-existence of free will entail that we should not hold people responsible. I will first discuss if people are responsible in a metaphysical sense and then if it is more practical to hold people responsible. I have defined moral responsibility as deserving blame for a bad action or praise for a good action, but I think it is necessary to define the term ‘blame’ here. Blame is directed to a person who carried out an action and not to the action itself. The person is at least to some degree guilty for carrying out the action. The person might be asked why he or she chose to carry out the action, but no answer is acceptable (if it is, the person is no longer blamed). If there are fully explanatory reasons for carrying out the action, these are either not accepted or the person is no longer blamed. The way blaming is used today, you cannot be blamed for an action if carrying out this action was your only alternative. One could argue that

this is not what the term refers to, but this is in fact what most people do when they blame a person. I would not be opposed to a view claiming that we should blame people *without* claiming that the person had other alternatives and is guilty. The difference is merely that I choose not to call that blame, since the term currently refers to this. This might seem inconsistent with Hume's argument that we hold people responsible only if they have reasons for acting. However, the problem is that most people think that we use reasons when we decide to act (with our free will) and when our actions seem to lack reasons, people interpret that as a sign that we did not use our free will. Most people do not realise that our actions must be either causally determined or random (or a mix of both). We do hold people responsible for reasons, but we do not realise that these reasons are fully explanatory (i.e. the reasons can explain actions entirely and no mysterious free will has a role).

As Nagel pointed out, if free will does not exist, your actions are a direct result of what genes you have and what events you have experienced and are shaped by. It might also be possible that chance plays a role. Without free will, no one could have acted differently – there are fully explanatory reasons to every action, unless chance exists. However, actions resulting from chance do not open up a possibility for people to act differently. It only entails that different outcomes (i.e. actions) might sometimes be possible and people do not influence which of the different outcomes will occur. If a person could not have acted differently, he or she is excused from responsibility. Therefore, the non-existence of free will entails that no one is ever responsible. However, this does not necessarily entail that we cannot hold people responsible anyway. A well-functioning society might require that we *hold* people responsible despite the fact that they are not actually responsible. There might be practical benefits with holding people responsible, such as preventing people from acting immoral (because they know that they will be blamed) and motivating them to act differently in the future once they experience the discomfort of being blamed. However, there might also be negative consequences with holding people responsible, such as paralysing feelings of guilt and failure to see the opportunity and possibility to change. In the remainder of the essay, I will first discuss what is morally relevant about moral responsibility and then discuss whether these are the practical consequences of holding people responsible and whether the advantages are greater than the disadvantages.

Whether we should hold people morally responsible is a moral question, but I think that people have different feelings about what is morally relevant about moral responsibility. Some might feel that we should be allowed (or even that we have an obligation) to hold people responsible – that it should be our right. I do not agree with this idea. The only morally relevant aspect of moral responsibility to me is that I find it negative if someone who could not have acted differently is held responsible. I do not see anything morally relevant in being allowed (or obliged) to hold other people responsible. I have the same feelings about revenge. In fact, I consider blaming a kind of revenge, because it is negative for a person to be blamed and it might satisfy the vindictiveness of other people. These are my moral feelings and I understand that other people might not agree with me. I also want to point out that I do not think that morality has an intrinsic value; rather, we value it because it helps us to cooperate and makes us feel good. To say that something is “moral” is not a good reason to accept it, unless it helps people collaborate, makes them feel good or has some other benefit. This is a better reason to accept an alternative than just stating that it feels intuitively right. Of course, people have different opinions about the benefits of different alternatives and some of these differences are due to differences in how we experience emotions and which emotions different events provoke.

The disagreement in this question seems to be based in different beliefs and opinions. I will present and evaluate different alternatives and elucidate the beliefs and opinions these alternatives are founded on. There are some alternatives that imply that our views on moral responsibility would not need any changes if genuine agency or free will does not exist and some alternatives which argue that we should change our views.

4.2 Evaluation of Possible Solutions

4.2.1 Free will is not what moral responsibility is founded on

One reason for arguing that we should keep our current view is the belief that this kind of free will is *not* what moral responsibility is founded on. To hold people responsible is a way for us to deter people from actions we find immoral and to change people. Although some people think that free will is needed for holding people responsible, proponents of this position argue that is not the case. To hold people responsible is a practical method that gets the results we want. People who defend this alternative (such as Nahmias (2011)) might argue that many people would act more immoral if we told them that they are not really responsible for their actions and

if we would treat them accordingly. However, opponents might argue that normal people distinguish between people who they think acted with their free will and people who did not (children, mentally ill people etc.) when they attribute responsibility to them. If this is a practical method, should every person who acts immoral be attributed equal responsibility, no matter what caused them and what condition they are in? The proponents might answer that it is not practical to hold some people responsible who did not even have Fischer's guidance control or was out of Churchland's "volume" of control. They seem to argue that people with compatibilist free will respond differently to being held responsible than the people who do not have that control. People who are in control will act better if we hold them responsible than if we do not. However, they seem to think that you are not allowed to give people any kind of sanctions if they cannot be held responsible. This is not necessarily the case. They could argue that holding people responsible is the only or at least the best functioning alternative.

4.2.2 Our views are innate and it is more practical to keep them

A similar position is that we should not change our views because we are evolutionarily shaped this way. Human beings hold other human beings morally responsible for their actions because it was evolutionarily adaptive and that we lack free will does not change that. It is very unpractical to try to change the way people function and naturally react and feel. Since there is not a right answer to how we ought to change our views, we should choose the most practical alternative – to continue to hold the views we have naturally. An opponent could point out that we have changed our legal system as well as our values before – why should it be a problem now? The proponents might answer that we would not win much on changing our legal system and that it would be very demanding to change a strong innate trait. But what would the benefits of changing our views on moral responsibility be? It would of course depend on how we would change our views, but some version could result in removing blame and guilt and to make people strive to change their behaviour instead of mourning about who they are. This is the opponents' view, but the proponents of this position might either think that it is impossible to get these results or that they would come with a side effect of immoral actions.

4.2.3 We should not hold people who are not in control morally responsible

An alternative which includes changing our views is Churchland's position. The absence of free will as such does not require us to change our views, for the same reasons as those given by

the proponents of the alternative presented in the previous paragraph. What should change our views is the knowledge we get from neuroscientific research – that our brains directly causes our behaviour and our behaviour is thus dependent on the state of our brains. Existent and future knowledge will make it possible for us to determine whether or not a brain is in control. The reason for making this distinction – and holding the people in control responsible – might be that the behaviour of people in control can be roughly estimated and controlled with different kinds of punishment and that they are believed to have what Fischer calls guidance control. People who are not in control do not act as consciously and they would thus not change their behaviour in the same way if we held them responsible. The opponents could say that although the distinction between brains in control and not in control is useful (because these two groups of people might need different kinds of treatment), that people in control would be ascribed moral responsibility does not make sense. There is nothing beneficial about holding people responsible instead of treating them to act better in the future. We would then change our views on people not in control while keeping our not very optimal views to everyone else.

4.2.4 We should not hold people morally responsible

My position is that there are greater disadvantages than advantages to holding people responsible and we should therefore not hold people responsible. I do not claim that it is the definite truth. Rather, I want to elucidate the assumptions underlying this position and I want to make the reader consider this alternative. Although my experiences point to this alternative, others might have different experiences and all readers might therefore not agree with me. However, I claim that it is possible (and, according to my experiences, probable) that these assumptions are true and it is therefore possible that this position is the most practical. The benefit of holding people responsible is that it sometimes can function as a deterring factor and make people choose moral actions instead of immoral actions, because it is undesirable to be blamed. The disadvantage is that the *person* is the one who is blamed and not the *action* he or she performed and consequently, the person might fail to see the possibility to change his or her behaviour. I will explain this further later in the essay.

I want to point out that I do not want to simply stop holding people responsible without an alternative. Without holding people responsible, it is still possible to say that a person was the one who performed the action, to say that the person acted intentionally, to discourage certain

behaviour, to encourage and argue for alternative courses of action, to punish people as a way of deterring them, to give them treatment against their will, to protect society from them and to simply avoid people who act immorally. Some might say that this is very similar to holding people responsible, but there is one important difference – the focus is on the *action* (and choosing differently in the future) instead of the *person* and the person is not considered guilty.

Today, we are holding people responsible. Will people act more immoral if we stop holding them responsible? I say no – we have plenty of reasons to act morally: we feel good about ourselves, other people appreciate us, it feels good (because of feelings like empathy), others might treat us well in return and we do not have to risk legal consequences. It is not like our responsibility is the only thing that keeps us from acting immoral all the time. The opponents of this alternative do not think that moral responsibility does more harm than good. They might think that being held responsible is an important reason for many of our moral actions. I argue that being held responsible has a relatively small role in our moral actions – other reasons are vastly more important.

Of course we must state that an action belongs to a person to be able to do something about it, but Mackie's distinction between intentional and unintentional is not precise enough. Since no one is responsible, the distinction must be made between people who can be changed to choose a better action next time a similar situation arises and people who cannot. Otherwise, it is not useful to try to change the person. In part, our intuitive feeling of responsibility attribution already follows this principle. We would not ascribe any responsibility to a natural event, since we cannot change nature's mind. However, we would ascribe responsibility to a sane, adult person more than to a mentally ill person (who cannot control his or her behaviour as well) or a child (who has not yet fully developed the ability to act in an appropriate way). This is, as Churchland realises, because it has been evolutionarily advantageous. I argue that evolution has not equipped us with the most practical solution for the society we live in today. I will discuss this later.

We normally think that some persons (for instance some mentally ill persons) cannot choose some of their actions. Now we must accept that we are all like that, in every decision we make, so we should treat all persons as we would treat persons who cannot choose their actions – not blaming them, but instead help them to act in a better way. Some might argue that we refrain

from holding some mentally ill persons responsible because it is their illness which makes them act immoral and it is not a permanent trait of the person. However, all mental illnesses cannot be treated and many people have to live with the illness (and its influence on moral actions) for the rest of their life. Some people develop mental illnesses because they are genetically predisposed to get the illness and others develop an illness because of environmental factors (and some because of both factors) – just like healthy people are influenced by both genes and environment. Perhaps the treatment will differ between mentally ill persons and healthy people (which is why Churchland's distinction is still useful), but other than that, there is no fundamental difference between these two groups regarding moral capacities. If a mentally ill person and a healthy person commit the same crime, you cannot say that the healthy person could have known better – because obviously, he or she did not know better. What is a good way to act is another question.

Our belief in free will leads us to assume that we cannot influence people that much, because they have some metaphysical free will which has more control than the experiences the people are subjected to. The nonexistence of free will means that this factor can be removed from the calculation and consequently, the experiences we expose people to have a greater impact. We have a better chance of changing people and this is not best done by blaming them.

Why is this not best done by blaming people? To blame people seems to be the solution evolution has equipped us with and yes, it may be cost-efficient to simply feel the need for holding people responsible and punish them and it works well in enough cases. However, for us today, with the knowledge we have about how people function (i.e. that we lack a free will and that we can influence people), it is possible to go against our evolutionarily created intuitions and instead of blaming people, use better methods of changing them. The opponents might argue that it is not possible because our innate feelings regarding responsibility attribution are too strong. However, we have done this with other evolutionarily created behaviour – we no longer accept that people abuse or kill each other, despite the fact that this might have been evolutionarily adaptive.

Our values and norms have changed before when we have gained new knowledge and this case is not different, according to me. In fact, the Swedish legal system has abolished punishment based on revenge and only punishes people for other reasons – deterrence, protecting society

from dangerous persons and treatment. However, many Swedish persons still disagree with this system.

In what direction we should change people is a moral question and there is no objective answer. Of course it is possible to argue for various answers, but those arguments would be aiming to change people's opinions about what is important and people's intuitions. Although people change opinions, I argue that we should change people in the direction of what people think is morally right. This can be different in different societies and groups.

A result which is often desired from holding a person responsible is that the person should feel guilt. To feel guilt is to feel like one is at least to some extent a bad person. The emotion's connection to the action one feels guilty for, is that the action made one a worse person or that the action is just a symptom which shows that one is (to some extent) a bad person. To feel regret is different, since regret is directed toward the action and not the person who performed the action. It is obviously easier to decide to act differently in the future than to decide to change the kind of person one is. To say that the problem is the person suggests that the action is more strongly connected to the person. Hume was not wrong when he stated that actions proceed from some cause in the character, but that cause does not have to be a permanent trait (if it is, then trying to change the person is useless). The goal is to make the person act differently in the future. All that is needed to reach that goal is to change the reason for performing the action (by making the person regret the action) and thereby make the person act differently in the future. To make a person feel guilty is to make the person believe (or just feel) that the action has a strong connection to him or her and thus makes it seem harder for the person to change. I argue that this discourages the person to change and instead makes it more likely that he or she reluctantly accepts the kind of person he or she is. In contrast, making a person regret an action motivates him or her to act differently in the future. I think that guilt often is thought of as a very intense emotion, while regret is thought of as a mild emotion. However, it is possible to feel intense regret as well – in moral contexts often accompanied by empathy for people who suffered as a result of the action.

Even if a person does not feel guilt when held responsible, the focus is still on the person instead of the action and therefore, it will be less likely that the person sees the opportunity to

decide to act differently in the future. Some people might also be more reluctant to accept that their behaviour is averse and that they should act differently in the future if they are blamed.

To sum up my position, there are both advantages and disadvantages with holding people responsible and I argue that the disadvantages are larger. The alternative is to stop blaming the *person* and focus more on the *action* and thereby induce feelings of regret instead of feelings of guilt. It seems to me that a person's genes and environment are the only things which determine what action the person is going to perform – no mysterious free will has something to say. Therefore, other peoples' reactions and other consequences of actions have a big role in influencing the person's future behaviour. To discourage certain behaviour, to encourage and argue for alternative courses of action and to inflict legal consequences are some ways of changing people. There are many reasons to act morally – being held responsible is only one not very important reason out of many reasons. In a metaphysical sense, no one is morally responsible and the most practical alternative (i.e. the one which would make our society function best and make people feel good) is not to hold people responsible. Therefore, I argue that we should stop holding people responsible and instead focus on better ways to try to change people.

5. References

- Churchland, Patricia S. *Braintrust*. Woodstock: Princeton University Press, 2011.
- Churchland, Patricia Smith. "Moral decision-making and the brain." In *Neuroethics*, by Judy Illes, 3-16. New York: Oxford University Press, 2006.
- Clarke, Randolph. "On an Argument for the Impossibility of Moral Responsibility." In *Free Will and Moral Responsibility*, by Peter A. French and Howard K. Wettstein, 13-24. Oxford: Blackwell Publishing, 2005.
- Dennett, Daniel. *Elbow Room*. Cambridge: MIT Press, 1984.
- Farthing, G. William. *The Psychology of Consciousness*. New Jersey: Prentice Hall, 2002.
- Fischer, John Martin. "The Free Will Revolution (Continued)." *The Journal of Ethics* 10 (2006): 315-345.
- Frankfurt, Harry. "Alternate Possibilities and Moral Responsibility." *The Journal of Philosophy* 66 (1969): 829-839.
- Fried, Itzhak, Roy Mukame, and Gabriel Kreiman. "Internally Generated Preactivation of Single Neurons in Human Medial Frontal Cortex Predicts Volition." *Neuron* 69 (2011): 548-562.
- Hume, David. *A Treatise of human nature*. Kitchener: Batoche, 1999, 1739.
- Levy, Neil. "The Good, the Bad and the Blameworthy." *Journal of Ethics and Social Philosophy* 1 (2005): 1-16. Cited in Levy, Neil, and Michael McKenna. "Recent Work on Free Will and Moral Responsibility." *Philosophy Compass* 4 (2009): 96-133.
- Levy, Neil, and Michael McKenna. "Recent Work on Free Will and Moral Responsibility." *Philosophy Compass* 4 (2009): 96-133.
- Libet, Benjamin. "Time of Conscious Intention to Act in Relation to Onset of Cerebral Activity (Readiness Potential)." *Brain* 106 (1983): 623-642.
- Mackie, John L. *Ethics*. London: Penguin Books, 1990.
- Mele, Alfred R. "Manipulation, Compatibilism, and Moral Responsibility." *The Journal of Ethics* 12 (2008): 263-286.
- Nagel, Thomas. "Moral Luck." In *Basic Readings*, by Nigel Warburton, 202-213. Oxon: Routledge, 2005.
- Nahmias, Eddy. "Why 'Willusionism' Leads to 'Bad Results': Comments on Baumeister, Crescioni, and Alquist." *Neuroethics* 4 (2011): 17-24.

Pereboom, Derk. *Living without Free Will*. Cambridge: Cambridge UP., 2001. Cited in Levy, Neil, and Michael McKenna. "Recent Work on Free Will and Moral Responsibility." *Philosophy Compass* 4 (2009): 96-133.

Smith, Angela. "Control, Responsibility, and Moral Assessment." *Philosophical Studies* 138 (2008): 367–392. Cited in Levy, Neil, and Michael McKenna. "Recent Work on Free Will and Moral Responsibility." *Philosophy Compass* 4 (2009): 96-133.

Smith, Kerri. "Taking Aim at Free Will." *Nature* 477 (2011): 23-25.

Soon, C. S., Brass, M., Heinze, H.-J. & Haynes, J.-D. "Unconscious determinants of free decisions in the human brain." *Nature Neuroscience* 11 (2008): 543-545.

van Inwagen, Peter. *An Essay on Free Will*. Oxford: Oxford University Press, 1984.

van Inwagen, Peter. "How to Think About the Problem of Free Will." *Journal of Ethics* 12 (2008): 327-341.

Vargas, Manuel. "Responsibility and the Aims of Theory: Strawson and Revisionism." *Pacific Philosophical Quarterly* 85 (2004): 218-241.

Wegner, Daniel M. *The Illusion of Conscious Will*. London: The MIT Press, 2002.

Widerker, David. "Farewell to the Direct Argument." *Journal of Philosophy* 99 (2002): 316-324. Cited in Levy, Neil, and Michael McKenna. "Recent Work on Free Will and Moral Responsibility." *Philosophy Compass* 4 (2009): 96-133.

Vihvelin, Kadri. "Compatibilism, Incompatibilism, and Impossibilism." In *Contemporary Debates in Metaphysics*, by John Hawthorne, Theodore Sider and Dean Zimmerman, 303-318. Malden: Blackwell, 2007.