General introduction

Prelude

The inimitable sci-fi story *Altruizine, or A True Account of How Bonhomius the Hermetic Hermit Tried to Bring About Universal Happiness, and What Came of It* written by the Polish writer Stanislaw Lem narrates the manifold ways in which good intended attempts to create a better society by disseminating a pioneering psychopharmacological substance can backfire (Lem, 2014 [1965]). In this story, the robot Bonhomius leaves the cave he has spent over sixty years meditating in to help his fellow creatures: "it dawned on me that to spend a life in solitude was wrong, for truly, did all my exceedingly profound thoughts and strivings of the spirit ever keep one rivet from falling, and is it not written that thy first duty is to help thy neighbour and not to tend to thine own salvation" (Lem, 2014 [1965], p. 246).

Bonhomius is accompanied by his mentor Klapaucius, a 'constructor' with exceptional engineering skills. They turn to Gnost, a hesitant supercomputer "able literally to contain the Universe itself within its innumerable memory banks" (Lem, 2014 [1965], p. 256) and ask him to help them find a way to eliminate misery and bestow happiness on everyone. Despite his considerable scepticism – numerous earlier attempts have failed miserably – the supercomputer eventually complies and prints them a recipe for 'Altruizine', a drug that triggers profound empathy[1] by replicating the emotional states of others in one's vicinity. Its workings are detailed below:

> ALTRUIZINE. A metapsychotropic transpitting agent effective for all sentient homoproteinates. The drug duplicates in others, within a radius of fifty yards, whatever sensations, emotions and mental states one may experience. Operates by telepathy, guaranteed however to respect one's privacy of thought. Has no effect on either robots or plants. The sender's feelings are amplified, the original signal being relayed back in turn by its receivers and thereby producing resonance, which is as a result directly proportional to the number of individuals situated in the vicinity. According to its discoverer, ALTRUIZINE will insure the untrammelled reign of Brotherhood, Cooperation and Compassion in any society, since the neighbours of a happy man must share his happiness, and the happier he, the happier perforce they, so it is entirely in their own interest that they wish him nothing but the best. Should he suffer any hurt, they will rush to help at once, so

1 A word about concepts here. In the multidisciplinary field of empathy research, the concept of empathy is used in a variety of ways, to refer to a range of distinct but related (emotional, motivational, and cognitive) psychological phenomena (Batson, 2009; Decety & Cowell, 2014). Arguably, the kind of empathy experienced as a result of Altruizine, resembles what in the literature is termed "emotional contagion": experiencing the feelings of those in your immediate vicinity.

> as to spare themselves the pain induced by his. Neither walls, fences, hedges, nor any other obstacle will weaken the altruizing influence. The drug is water-soluble and may be administered through reservoirs, rivers, wells and the like. Tasteless and odourless. One millimicrogram serves for one hundred thousand individuals. We assume no responsibility for results at variance with the discoverer's claims. Supplied by the Gnost. Computerized representative of the Highest Poss. Lev. Devel. (Lem, 2014 [1965], pp. 267-268)

Bonhomius sets off to a planet inhabited by humans and diffuses the Altruizine drug. Chaos ensues. Overwhelmed by the one-on-one transmission of other peoples' emotions and sensations, people start to roam around in mob-like hordes and to violently attack each other.

> By dawn the Altruizine effect was so strong, that if one nostril itched, the entire neighbourhood for a mile on every side would respond with a shattering salvo of sneezes; those suffering from chronic migraines were abandoned by their families, and doctors and nurses fled in panic when they approached – only a few pale masochists would hang around them, breathing heavily. And then there were the many doubters who slapped or kicked their compatriots, merely to ascertain whether there was any truth to this amazing transmission of feelings everyone spoke of, nor were these compatriots slow in returning the favour, and soon the entire city rang with the sounds of slaps and kicks. (Lem, 2014 [1965], pp. 271-272)

Bonhomius' and Klapaucius' honest attempt to bring about a better society by radically enhancing people's empathy, has failed.

'Introduction to the introduction'

One could easily dismiss this pessimistic and rather gloomy story about the use of a pharmacological agent to increase human empathy as rather farfetched, and an interesting thought experiment at best. Nonetheless, one of the most lively and vehemently debated issues in the bioethical literature of the last decade concerns closely related proposals. Since 2008, the so-called 'moral enhancement debate' asks whether we should actively pursue the development of moral enhancement technologies, and whether it would be permissible – or even obligatory – to put them to use, provided that these interventions would be effective and safe?

Whereas 'traditional methods' of moral betterment (such as upbringing, socialization and education) are arguably as old as humanity itself, the debate on moral bioenhancement focuses on the desirability of methods based on novel or emerging biomedical insights, or the use of biomedical methods. The debate follows a significant rise in fundamental research on the (neuro)biological and genetic underpinnings of human (and animal) morality – moral emotions, moral cognition, moral judgment, moral behaviour. Potential interventions that are being discussed in the literature on moral bioenhancement range from various types of psychopharmaceuticals, neurostimulation, and genetic selection and engineering.

Central topics of debate have been how to understand and define moral enhancement; whether (safe and effective) moral enhancement technologies should be mandated; and the question whether means matter morally, i.e. whether there are inherent ethical differences between biomedical and nonbiomedical means of moral enhancement. Just as earlier human enhancement debates (dealing with the ethical desirability of enhancing beauty, sports, mood, cognition and memory), the debate on moral enhancement is highly polarized. Proponents argue that moral enhancement is our only hope in averting disaster, while opponents fear that moral enhancement would *provoke* disaster.

However, until now, the theoretical debate on moral (bio)enhancement has a strong speculative character and mostly precedes and runs ahead of realistic scientific possibilities. Moreover, to date, the debate risks lacking focus and real world impact, as different commentators fail to agree on how to understand and define moral enhancement, and on what kind of examples would constitute cases of moral enhancement.

This under-examination of potential *practices of moral enhancement* is significant because without specifying intended users, contexts of implementation, and the goals and objectives of developing and applying potential biomedical possibilities for optimizing morally relevant capacities, it is not clear who should be concerned about this debate. At the same time, existing or emerging practices that already contain elements of 'moral enhancement' but as yet lack ethical attention and ethical scrutiny, risk staying out of sight.

In order to address this problem, this thesis focuses on (present and emerging) moral enhancement *practices*, in order to identify ethical issues that are not necessarily part of the current debate on moral enhancement. Rather than distinguishing in a blanket fashion desirable from undesirable moral enhancement *technologies*, this thesis intends to formulate conditions and ethical requirements for ethically justifiable moral enhancement practices.

Setting the scene: The new sciences of morality and the debate on moral bioenhancement

The promises of the new sciences of morality

For the last couple of decades, a growing body of research into the so-called "new sciences of morality" studies neurobiological and behavioural genetic models of human morality (Slaby, 2013; Yoder & Decety, 2017). A plethora of scientific disciplines, including social and behavioural neuroscience, biology, evolutionary psychology, experimental psychology, and primatology, study phenomena moral philosophers have examined conceptually for ages.

Primatologists study our evolutionary ancestors, mainly great apes like bonobos and chimpanzees, in order to identify potential continuities between these species and ours. Frans de Waal for example argues that we can find a kind of "proto-morality" in these apes, in the form of empathy, reciprocity, and a sense of fairness that can rightly be called he thinks "the building blocks of morality" (de Waal, 1996, 2009; de Waal & Sherblom, 2018).

In the field of evolutionary and cognitive psychology theorists look upon morality as an evolutionary adaptation and ask how something morality-like might have been advantageous, that is fitness enhancing, for our ancestors (Katz, 2000; Nichols, 2005; Kitcher, 2007). Their research gives way to discussions about whether (and if so, how much of) human morality should be considered 'innate'. Some hypothesize that humans come to this world with an innate capacity for moral judgments already in place (Hauser, 2006). Others go even further and speculate about specific moral judgments being innate, in the form of for example a universal bias against 'unauthorized killing' that forms a necessary precondition for close knit communities to live and work together.

In the early 2000's, in various 'morality labs', the first neuroscientific experiments were set up, aimed at mapping the neural circuits that are involved in moral thinking and moral judging and to find perhaps a specialized moral centre in the human brain (Damasio, 2000; Greene, 2001; Haidt, 2001; Greene & Haidt, 2002). These experiments raised questions about the role played by emotion and cognition in moral judgments, and stirred discussions about which of these two should be considered as a more decisive factor. Later, attention shifted from trying to uncover "domain-specific morality", i.e. identifying brain regions specifically dedicated to morality, to mapping more domain-general processes, such as theory of mind (Young & Dungan, 2012).

In the meantime, a rich landscape of studies on the 'emerging neurosciences of' has appeared, ranging from the emerging neuroscience of justice motivation and justice sensitivity (Decety & Yoder, 2017), empathy (Marsh, 2018), intergroup relations (Cikara & Van Bavel, 2014), social decision-making (Yoder & Decety, 2017), violence (Miczek et al., 2007), to pro-social behaviour (Yamasue et al., 2012), etc. etc. The same goes for (behavioural) genetic studies on for example violence and aggression (Baum, 2013; Asherson & Cormand, 2016), and antisocial personality disorder and psychopathy (Viding et al., 2005; Tracy et al., 2010).

In criminological research a previously improbable interest in biological and social underpinnings of antisocial behavior is visible (Raine, 2013, pp. xii-xiii; Glenn & Raine, 2014). Bio-social criminology stresses the interplay between social and biological factors: "What is important about this new line of research is that it is not a reification of the outdated nature versus nurture debate, but rather an entirely new perspective that highlights the dual influences of genetic and environmental factors in the etiology of crime and delinquency" (Beaver et al., 2015, p. 109).

Moreover, these biological, 'brain-based' explanations of (im)morality[2] have given rise to related debates on the relevance of this growing body of knowledge for a variety of domains (Blank, 2007; Levitt & Manson, 2007; Mendez, 2009; Zik & Roberts, 2015). In the field of 'Neurolaw' for example, the meaning and implications of neuroscience for the law and legal practices is investigated (Meynen, 2014, p. 819; Morse, 2017). Neuroscientific information and techniques are already being used in court (de Kogel & Westgeest, 2015). This research has also stirred debate about the question whether in the future brain imaging could be used to (aid in) accurately predicting future violence and other forms of antisocial behaviour (Poldrack, 2017; Poldrack et al., 2018), and help in tailoring treatment to individual needs.

> Although single genes only contribute to a small proportion of the overall variance in antisocial behaviour, identifying genes that confer risk may aid in the development of treatment methods that could potentially be tailored to specific risk factors of the individual. In addition, it may improve our understanding of the biological pathways that lead to antisocial behaviour. (Glenn et al., 2013, p. 4)

The promises of these new sciences of morality thus go beyond improved understanding of biological underpinnings of morality. As moral (and immoral) behaviour are

2 In *The Anatomy of Violence*, Adrian Raine argues that criminals have "broken brains, brains that are physically different from those of the rest of us" (Raine, 2013, p. 180).

explained in terms of (disordered) brain functions, commentators have also speculated about how insights from neuroscience could benefit conflict resolution and foster peace (Bruneau, 2015) or even ways in which potential neurobiological interventions could be used to correct potential "social or moral dysfunction":

> Understanding the functional anatomy of moral judgment and the anatomical and functional differences between social and antisocial brains is a prerequisite for developing socially and morally acceptable neurosurgical interventions that treat social or moral dysfunction. (de Ridder et al., 2009, p. 156)

In short, the new sciences of morality not only add to and challenge previous predominantly non-biological scientific understandings of human morality (Abend, 2013), but go a step further as well, and suggest novel methods to deal with behavioural problems, such as pharmacological means, neurostimulation, or even neurosurgical interventions as mentioned in the quote above. As Adrian Raine writes in his book *The Anatomy of Violence*: "It's a bitter pill for many criminologists and psychologists to swallow, but medications do work in controlling and regulating aggression in children and adolescents" (Raine, 2013, p. 292). Current techniques still fall short however, and a range of methodological and ethical challenges remain to be met.

There are of course historical predecessors to these biological approaches to moral and immoral behaviour. Felix Schirmann narrates in what ways the moral brain has been an object of scientific study before, and the 19th century ways of "thinking of immoral persons in terms of disordered brains altered what it means to be immoral. Though contested, the immoral persons discussed were conceived of as viable objects for brain science, psychiatry and medicine: their misdoings were observed on a behavioral level, their actions were evaluated with ethical standards and they were punished according to the law. Observation, evaluation and reprimand were part of the social purview, yet the respective experts sought the causes for their misconduct in their neurobiology." (Schirmann, 2013a, p. 43; 2013b) Notorious is the identification of 'human types', among which 'the born criminal', by the Italian physician and anthropologist Cesare Lombroso, whose ideas were influential in the first half of the twentieth century (for example in the eugenics movement). During the second half of the twentieth century, biological explanations for criminal or immoral behaviour were strongly resisted. For instance, in The Netherlands, professor in criminology Wouter Buikhuisen was deterred from doing research on biological features of criminals, next to social factors

(Berkvens, 2009; de Haan, 2009; Oosterhuis, 2014).[3] Now, however, renewed attention for genetic and biological risk factors for criminality has emerged.

Interest in the neuroscience of morality has not been confined to the academic world, but has attracted considerable interest from the general public and from policy makers as well. Popular science books abound, both about the *moral* (Hauser, 2006; Churchland, 2011) and the *immoral* brain (Fallon, 2013; Raine, 2013), and about empathy or lack thereof (Baron-Cohen, 2011). Fascination with psychopathic or criminal genes and brains is reflected in considerable attention in the media (Kahn, 2012; Dennett, 2013; Frentz et al., 2015; Levy, 2015a; Hagerty, 2017), although resistance to the advance of so-called 'brain culture' is visible as well (Dehue, 2011; Bloemink, 2013; Dehue, 2014; de Vrieze, 2017; Koelewijn, 2017).

Also here, attention goes beyond fascination with biological explanations of moral and immoral behaviour, and the question is asked whether we can put this new knowledge about the underpinnings of human moral and immoral behaviour to good use. For example, can we teach people (e.g. our children, or medical doctors) to have more empathy (Boodman, 2015; Krznaric, 2015; Eerkens, 2017), for example by using virtual reality (Polak, 2015; Berdik, 2017)? Or, one step further, by means of a simple pill? Plentiful op-eds about the desirability of so-called 'morality pills' have been published (Kahane, 2011; Singer & Sagan, 2012; Crockett, 2014b; Levy, 2015b; Dubljević, 2017; Mandelbaum, 2017; Rampton, 2017; Tremonti, 2017) and on various online forums people have been discussing the possibility and desirability of developing, propagating, and using such pills (Pellissier, 2012; Aeon, 2015).[4]

3 "by the 1970s it became sociological common sense that fatalism, determinism, reductionism, sexism – a naturalization and legitimation of existing relations of power – would follow inescapably from any engagement with the reality of human biology – as either an ontological question – what were humans really like? – or as an epistemological one – what can biology tell us about the forms of life that humans have made for themselves? Human biology was relevant only in that it provided the preconditions for language, meaning and culture, whose form and content must be accounted for in non-biological terms." (Rose, 2013, p. 10)

4 Fascination with potential ways to improve people morally, including taking a pill, is evident in the television series *The Good Place* as well (Goddard et al., 2016). In one of the first episodes, protagonist Eleanor Shellstrop needs a quick and dirty way to improve herself morally in order to extend her place in The Good Place: "It is time to make me good, partner! How do we do it? Is there a pill I could take, or something I could vape?" But instead of giving her a pill her partner Chidi, a former professor of moral philosophy, decides to teach her philosophy as a means to become a better person.

The bioethical debate on moral enhancement

In the current academic debate on moral bioenhancement it has been suggested that biomedical interventions will aid in solving a number of pressing societal issues, such as crime and violence, or will aid in solving climate change - either by enhancing pro-social tendencies and emotions or by inhibiting anti-social tendencies and emotions. The debate started in 2008, with the publication of two articles advocating moral enhancement in the *Journal of Applied Ethics* (Douglas, 2008; Ingmar Persson & Savulescu, 2008).[5]

In his first paper on moral enhancement, Thomas Douglas main intention is to debunk what he calls the Bioconservative Thesis (Douglas, 2008). It posits that "Even if it were technically possible and legally permissible for people to engage in biomedical enhancement, it would not be morally permissible for them to do so" (Douglas, 2008, pp. 228-229), because, the bioconservative argument goes, although the enhancement might be good for the individual, it might be bad for others. Douglas then presents moral enhancement as a type of enhancement that is not necessarily disadvantageous for others, and can even benefit others (and not necessarily yourself). A good target for moral enhancement according to Douglas, are the "counter-moral emotions":

> I have in mind those emotions which may interfere with all of the putative good motives (moral emotions, reasoning processes, and combinations thereof) and/or which are themselves uncontroversially *bad* motives. Attenuating such emotions would plausibly leave a person with better future motives, taken in sum. (Douglas, 2008, p. 231)

The central question Douglas poses is whether it would be ethically permissible for individuals to morally enhance themselves.

In a book and in a series of articles, Ingmar Persson and Julian Savulescu have argued that there is an urgent need to explore the possibility of using the emerging science of morality to develop means to improve - *enhance* - moral dispositions (Ingmar Persson & Savulescu, 2008, 2012; Savulescu & Persson, 2012). Later they synthesized these ideas in a book: *Unfit for the Future: The need for Moral Enhancement* (Ingmar Persson & Savulescu, 2012). They argue that people's moral capacities are essentially 'flawed', and that there

5 During the course of this research project, the amount of journal articles, books, academic conferences, and media attention for moral enhancement has risen significantly. A systematic literature review we performed at the beginning of this project (up until April 2014) yielded 85 academic articles discussing moral (bio)enhancement (see chapter 2). The amount of publications on moral enhancement now (spring 2018) is nearing 300.

exists a fundamental mismatch between our moral psychology and today's conditions of human life (Ingmar Persson & Savulescu, 2013). In addition there is, they posit, a considerable and growing potential for "a significant improvement of various aspects of human cognition by biomedical means", the risk that a small but potent minority will acquire the power to do great harm, grows (Ingmar Persson & Savulescu, 2008). To counter this risk, they argue, "cognitive enhancement would have to be accompanied by a moral enhancement which extends to *all* of us, since such moral enhancement could reduce malevolence" (Ingmar Persson & Savulescu, 2008, p. 166). An extensive moral enhancement of humankind, in conclusion, is our only hope in averting "ultimate harm."[6]

They ask whether we could, through our knowledge of "the new sciences of behavioural control", strategically influence people's moral dispositions and behaviour (Savulescu & Persson, 2012), and discuss a range of lines of research that they feel could provide new and effective means of influencing human choices, such as priming (affecting choice by unconscious stimuli) and nudging (influencing individual choice by changing the choice architecture of the environment). In addition, they discuss how a series of currently used pharmacological substances may be used to affect moral behaviour as a side effect.

Persson and Savulescu have defended mandatory programs if people fail to morall enhance themselves voluntarily:

> If safe moral enhancements are ever developed, there are strong reasons to believe that their use should be obligatory, like education or fluoride in the water, since those who should take them are least likely to be inclined to use them. (Ingmar Persson & Savulescu, 2008, p. 174; 2011see also: ; Ingmar Persson & Julian Savulescu, 2014b; Ingmar Persson & Savulescu, 2016)

More recently, they have argued that people even have a moral duty to morally enhance themselves (Ingmar Persson & Savulescu, 2017).

The God Machine

In a thought experiment not much different from Lem's futuristic story about *Altruizine* this chapter started with, Julian Savulescu and Ingmar Persson entertain the idea of

6 The concept of Ultimate Harm reminds of research on so-called existential risks (or x-risks); risks that could lead to human extinction or civilizational collapse (Bostrom, 2013), studied for example in the University of Cambridge based Centre for the Study of Existential Risk (https://www.cser.ac.uk/).

an all-powerful, self-learning, self-developing bioquantum supercomputer, named The God Machine that would intervene in human action to prevent "great harm, injustice or other deeply immoral behaviour from occurring" (Savulescu & Persson, 2012, p. 413):

> The God Machine would monitor the thoughts, beliefs, desires and intentions of every human being. It was capable of modifying these within nanoseconds, without the conscious recognition by any human subjects. (Savulescu & Persson, 2012, pp. 412-413)

Without the person involved being consciously aware, the God Machine would 'change his mind', for example by altering his intention to murder an innocent person. The machine is designed to ensure maximum individual freedom, Savulescu and Persson maintain: people's choices to do the morally right thing remain fully autonomous. It is only when a person is about to act in a clearly immoral way that the machine intervenes. Yet, the person in question would never be aware if the intervention, and as such she would still *experience* a life of complete freedom.

Would this God Machine be a bad idea, Savulescu and Persson ask, and, why would it be? The God Machine now makes impossible what in 'the old days' was prohibited by law, on pain of punishment. People are not free to murder now. Surely this small infringement of a person's autonomous choice is preferable to physical incarceration, or worse?

The primary purpose of this thought experiment has been to counter an objection by one of moral enhancement most outspoken critics, philosopher and bioethicist John Harris. He reasons that moral bioenhancement forms a fundamental threat to human freedom and autonomy, in particular to what he has called "the freedom to fall", the ability to choose based on reason and argument, and the ability to make an *ethical choice* (Harris, 2011, 2013a, 2014):

> Ethics is for those occasions on which altruism fails; or for those people who fail to think and feel, and/or who are not disposed to do as they should! Moral reasoning is needed to identify the appropriate objects for sympathy, empathy and the sort of generalized love that is the conclusion of a moral argument and which is often expressed as 'love thy neighbour as thyself'. We will, I believe, always need to use moral reasoning to act as a guide to our emotions and as a way of checking that we are having appropriate feelings in appropriate circumstances and for appropriate objects. (Harris, 2013a, p. 171)

Michael Hauskeller's critique on the God Machine focuses on the aspect of external manipulation and control the God Machine impersonates.

> We treasure human freedom, not the metaphysical one that very likely is a mere illusion, but the social, relational one. It defines what we are, our very humanity. From that perspective even the God machine poses a threat because the problem would still be that we are being controlled by a machine that has been designed with the explicit goal of controlling us. Whether we are actually aware of the manipulation or not is not relevant, or is not the main problem. The main problem is the fact that we are being controlled by, per definition, somebody else's will, and the fact that the purpose of our being controlled is that we conform to certain moral standards does not make it any better, but rather the contrary, because instead of being ends in ourselves, we thereby become means to the end of morality, which ultimately makes morality itself pointless because we no longer have a good reason to treat each other as ends if we are *in fact* nothing but means. (Hauskeller, 2017, p. 375)

The difficulty of defining moral enhancement

From the beginning, the debate on moral enhancement has been characterized by profound disagreements about the way the concept of moral enhancement should be defined. An exchange between neuroscientist Molly Crockett and ethicists John Harris and Sarah Chan can serve as an example (Crockett et al., 2010a, 2010b; Harris & Chan, 2010; Chan & Harris, 2011). According to Crockett et al., enhancing the neurotransmitter serotonin in the brain influences moral judgment and moral/pro-social behaviour, as their study shows that it diminishes the willingness to harm another person; it increases *harm aversion:* "our subjects were less likely to endorse personal moral harms and less likely to punish unfairness in the UG [Ultimatum Game]" (Crockett et al., 2010b, p. 17436). John Harris and Sarah Chan present an opposing view, and argue that serotonin seems to preclude and cloud moral judgment, rather than improving it.[7]

7 Harris Wiseman makes a similar point, mentioning both hesitations about the safety and efficacy of using SSRI's as a moral enhancement intervention, and concerns about whether SSRI's are "worthy of the name" of moral enhancement: "It is hard to imagine a serious moral outlook that can equate chemically induced pacification of behavior with a genuine moral improvement of the individual in question. Add to this that SSRIs cannot be made to work in a contextual fashion, as they are "all-or-nothing," brute agents that require periods of building up and tapering off, and that the only justifiable "populations" they might be effective in treating, assuming no abreaction occurs, are those with more serious mental health issues, in which case, it is more likely that stronger antipsychotics or sedatives will be used rather than SSRIs." (Wiseman, 2014b, p. 27).

> It is not clear that 'enhancing aversion to personally harming others' is something that would promote either moral behaviour in the public at large or indeed, harm reduction more generally, unless it was capable of much more nuanced effects than seems evident from the reported research. Thus, if serotonin affects moral behaviour, it does so adversely by impairing moral judgment, subjugating it to emotional instinct. We should be wary of assertions claiming that serotonin has a role in moral judgment; the opposite seems to be the case. It may enhance aversion to violence, but it does not enhance moral behaviour; it can increase, rather than diminish, harm to others and bypass the use of moral reasoning. (Harris & Chan, 2010, p. E183)

Underlying these disputes often lie profound philosophical battles about what constitutes morality. Harris believes that true moral growth can be achieved through cognitive means, by employing reason; he is a rationalist (Baertschi, 2014). Directly manipulating moral emotions, as advocated by for example Douglas (Douglas, 2008, 2013), would according to Harris lead to *moral decline*, not moral growth.

> Although both emotion and reasoning thus affect moral decisions, of the two, it must be reasoning that pulls in the direction of morality. A moral agent is not just someone who performs actions with moral consequences, she is a person who cares about doing the right thing. Such a person must have a way of deciding whether what her emotions prompt her to do, what strikes her as the right thing to do, really is the right thing to do. She will need to think things through, identify the relevant principles she accepts, the values she holds and the moral objectives she believes are right, and apply them to the present circumstances; and to do this she must use moral reasoning. (Chan & Harris, 2011, p. 130)

Taxonomy of definitions of moral enhancement, and their normative implications

In the following quotes, John Shook calls attention to the wide variety of definitions given:

> Too many discussions are proceeding as if both the meaning and the possibility of moral enhancement were already widely understood and agreed upon. (…) Asking such questions, and offering answers, depend on assigning some sense or another to "moral enhancement." However, clear and precise definitions of "moral enhancement" are not to be found; what has been called "moral" enhancement ranges from feeling empathic concern to increasing personal responsibility all the way to heightening respect for global fairness. (Shook, 2012, p. 3)

And:

> anyone using the term 'moral enhancement' as if everyone knows what is meant must either be simplifying matters to the point of negligence, or trying to speak only to those already in local moral consensus. (Shook, 2012, p. 4)

We identified different definitions of moral (bio)enhancement that have been given in the literature, and developed a taxonomy of different definitions and uses of the concept of moral (bio)enhancement as well as their normative implications (Raus et al., 2014).[8]

Some commentators provide a specific definition of moral enhancement. Thomas Douglas for example defines moral enhancement as:

> interventions that will expectably leave an individual with more moral (*viz.* morally better) motives or behaviour than she would otherwise have had. (Douglas, 2008, p. 229)

Other commentators do not offer such an explicit definition, and for the purpose of this taxonomy we therefore had to reconstruct the way the concept was used by the different authors.

Focus on the intervention or on the individual

The first distinction concerns different interpretations of the question *what is moral enhancement?* Whereas some definitions focus on the intended effect of a particular *intervention* or technology (i.e., "interventions that are intended to improve our moral capacities such as our capacities for sympathy and fairness" (DeGrazia, 2014, p. 361)), other definitions focus on questions about when an individual can be considered to have been morally enhanced, or what it means to say that an *individual* is morally enhanced. James Hughes for example argues that moral enhancement goes beyond "the jacking up of virtue with neurochemicals", but that it should be conceived more broadly, in terms of "taking conscious control of our lives to build the kind of character we want to have" (Hughes, 2011).

8 This paragraph draws on a paper by Kasper Raus et al., of which I am a co-author: Raus, K., Focquaert, F., Schermer, M. H. N., Specker, J., & Sterckx, S. (2014). On defining moral enhancement: A clarificatory taxonomy. *Neuroethics, 7*(3), 263-273. When discussing the different definitions and uses of moral enhancement in this paragraph, 'we' refers to the authors of this paper.

Broad versus more specific interventional means

Another relevant distinction we found concerns broad versus more specific definitions. Some authors prefer a more restrictive definition of moral enhancement that focuses on biomedical and genetic means only. These authors also tend to prefer the concept of moral *bio*-enhancement or *genetic virtue* (Walker, 2009, 2010). Others include any practice that brings about – or intends to bring about – positive changes in a person's moral capacities, including practices such as (moral) education and (talk) therapy. Because genetic and biomedical interventions (for example pharmaceutically or by means of (deep) brain stimulation) are generally considered more invasive and more precarious than non-biomedical ones, it might be expected that using "the concept of moral enhancement for *all* interventions or only for those interventions that are (most) invasive, can serve to bias the normative debate towards or away from a conclusion of permissibility or desirability" (Raus et al., 2014, p. 265).

Enhancing individuals versus enhancing humanity

A further way in which definitions of moral enhancement differ, is based on the question who is considered the target of moral enhancement: individuals, or groups of persons? Thomas Douglas, especially in his earlier work on moral enhancement, explicitly focuses on individual persons, as illustrated by the previously mentioned definition. Ingmar Persson and Julian Savulescu on the other hand, have advocated enhancing "the moral character of humanity" (Ingmar Persson & Savulescu, 2008), and have defended mandatory moral enhancement *programs*, aimed specifically at those most in need of moral improvement. As we argued in the taxonomy, these definitions have profound normative implications. For justifying moral enhancement applied solely to individuals, justifications related to autonomy and personal benefit might suffice. However, in justifying moral enhancement on a societal level one would need to turn to other justifications such as issues of justice or of achieving a common good. Moreover, concerns about altering human nature would then come into view as well.

Moral treatment versus moral enhancement

A further distinction concerns the distinction between moral treatment and moral enhancement. Whereas some definitions understand moral enhancement to be any form of moral improvement, regardless of whether a person's moral functioning before enhancement could be considered below average, average, or above average; other definitions clearly damarcate bringing people to an average or 'normal' level of moral functioning from enhancing beyond the average level. Improving towards average would be considered moral treatment, improving beyond average would be

considered enhancement.[9] Yet, defining moral normalcy raises questions. As there is no objective way of determining what falls within the range of 'normal' moral behaviour or functioning of moral capacities, every classification is, necessarily, a normative one.

John Shook and James Giordano have proposed an alternative classification between "moral (re)habilitation" (instilling or restoring some degree of moral capacity and responsibility in someone); "moral normalization" (an improvement of already-existing moral capacity towards society's standard of good moral conduct); and "surpassing enhancement" (improvement above regular requirements of common morality) (Shook & Giordano, 2016b, 2016a, 2017).

Intended versus effective interventions

A further difference concerns definitions that focus on either the intention to enhance versus interventions that actually have a positive effect. John Harris for example understands moral enhancement interventions to be those interventions that have an actual effect: "I do not define enhancements in terms of the intention or the motivation of those who produce them but rather in terms of their effect" (Harris, 2014, p. 372). On the other hand, the definition by Douglas given above, and the definition by David include *intended or expected but failed* interventions (Douglas, 2013; DeGrazia, 2014):

> Interventions that are intended to improve our moral capacities such as our capacities for sympathy and fairness. (DeGrazia, 2014, p. 361)

For them, what is important in determining whether an intervention is to be considered a moral enhancement is the *intention* with which the intervention is used, not its actual effect. This distinction between intended and actual effect is relevant for morally justifying moral enhancement. Because if an intervention can only be labelled a moral enhancement if it has a positive effect, then one important reason to oppose moral enhancement is eliminated. All problematic cases where no enhancement is reached or a person is left *less* moral than before the intervention, would not be considered cases of moral enhancement at all. A consequence of this way of justifying moral enhancement is that it would avoid debate concerning potential risks and side-effects of interventions.

9 As Dorothee Horstkotter and colleagues argue: "if there is a health problem, medical treatment is the reasonable reaction, while enhancement, either moral or otherwise, does not arise." (Horstkötter, Berghmans, & de Wert, 2012)

Capacities-oriented versus behaviour-oriented interventions

Some commentators in the moral enhancement debate label a certain intervention a moral enhancement depending on its (real or intended) effect on a person's *behaviour*. Others see moral enhancements as interventions that target or are intended to target a person's *capacities of moral reflection*. In the taxonomy, we have referred to this difference as the difference between a behaviour-oriented and a capacities-oriented intervention.

If moral enhancement interventions can target a person's behaviour or their capacities of moral reflection, any intervention can have one of four results. It can result in a person (1) reflecting the same, and acting the same; (2) reflecting the same, but acting differently; (3) reflecting differently, but acting the same; and (4) reflecting differently, and acting differently. Definitions of moral enhancement differ as to the question which of these should be considered cases of moral enhancement. In identifying which interventions can be labelled moral enhancements, some commentators rely heavily or solely on the intervention's achieved (or intended) effect on a person's behaviour [e.g. (Douglas, 2013)]. For those commentators, moral enhancement occurs only when a person, because of an intervention, *acts* differently. According to them, if a person reflects differently, but acts the same because of an intervention, this should not be considered a case of moral enhancement. A range of other authors do not think that the requirement for successful behaviour change suffices to label an intervention a moral enhancement, and instead categorise them as a form of behaviour control (e.g. Harris) or moral therapy (e.g. Agar):

> I take moral enhancement to involve enhancing our ability to think ethically (...), not manipulating the probability of some reacting in ways that *others* deem ethical. (Harris, 2014, p. 373)

> While the manipulation of moral emotions might change the behaviour of an individual, it does not provide any content, for example, norms or values to guide one's behavioural response. (Jotterand, 2011, p. 6)

> (i) forcing agents to act rightly, (ii) preventing agents from acting wrongly, and (iii) making it harder for moral agents to act wrongly fail to constitute genuine moral enhancement. (Simkulet, 2012, p. 17)

For these commentators, as well as for those who use similar lines of argumentation (e.g. (Lev, 2012; Baertschi, 2014)), the criterion for whether an intervention constitutes a moral enhancement is not (only) behaviour, but rather whether the intervention (also) affects or improves an individual's capacities for moral reflection. In the quote

below, Harris appears to embrace the most ambitious and restrictive definition of moral enhancement; acting differently and reflecting differently:

> It seems to me that moral enhancement, properly so called, must not only make the doing of good or right actions more probable and the doing of bad ones less likely, but must also include the understanding of what constitutes right and wrong action. (Harris, 2014, p. 172)

Active involvement versus passive receiving

A final distinction is that between a process of moral enhancement in which the individual is actively involved, versus a process in which the person is a passive recipient. This distinction does not concern the *target* of a moral enhancement, but rather *the way in which the enhancement is achieved*. A moral enhancement by way of active involvement would then be an enhancement requiring conscious mental processes in the subject as a means to achieve its result. Moral education would be a classical example. In contrast, moral enhancement involving passive receiving would either be enhancement in an immediate way where no active involvement was possible (e.g. a pill with immediate effect) or enhancement by way of a process that required no deliberate involvement of the recipient (e.g. classical conditioning).

This distinction is relevant since, for some authors, the goal of moral enhancement is for individuals to become more virtuous, and often these authors follow Aristotle in claiming that becoming virtuous is always a conscious and deliberate process, where the way to becoming virtuous is just as essential as the result (being virtuous). Chris Zarpantine talks about "the thorny and arduous path of moral progress" (Zarpentine, 2013, p. 141), while Jotterand states:

> Virtue is a behavioral habit under the supervision of reason that can be taught and learned. The control and manipulation of moral emotions by technological means reduce the human mind to neurochemical processes and threaten the very essence of moral agency, that is, autonomy. (Jotterand, 2011, p. 7)

The distinction is used to argue for the idea that the means by which moral enhancement is achieved matter morally (Focquaert & Schermer, 2015). Interventions where the individual is passive (more likely in the case of interventions that work directly on the

brain), Focquaert and Schermer argue, are more likely to compromise autonomy and disrupt identity than interventions in which the individual is an active participant.[10]

Conclusion: Definitions are not normatively neutral

As discussed above, definitions of moral enhancement vary widely – Harris Wiseman speaks of the "tremendous plurality the term conceals" (Wiseman, 2017, p. 398) – and have important normative implications. The chosen definition conveys which practices an author considers more or less ethically desirably. Moreover, the concept of moral enhancement sometimes appears to assume that improving a person's morality is unproblematically good – why else would it be called *moral* enhancement? Finally, as the sciences of morality are evolving, it is not yet clear what will be possible in the future, what kind of moral enhancement technologies will be developed, and how people will react to and interact with them.

Therefore, to a certain extent, this thesis takes the moral enhancement debate itself as the main unit of investigation. It asks in what ways the concept of moral enhancement is used, what kind of reasons participants in this debate have given for or against its desirability, what kind of interventions they have proposed, and what kind of problems they think it will solve. To be able to analyse these aspects, it is unhelpful to have a pre-defined notion of what 'moral enhancement' is or should be.

In this thesis the concept of moral enhancement is used to refer to attempts to influence, change, or optimize moral capacities, and limit immoral capacities, with the aim of changing a person's moral emotions, moral reasoning, or moral behavior. This includes instances where someone (e.g. a practitioner or policy maker) discusses or implements interventions that have an effect on a person's moral capacities, but that person does not necessarily discuss those interventions in explicitly *moral* terms. Finally, this working definition is not limited to biomedical interventions but also includes nonbiomedical, traditional means.

This thesis: Commitment to practices

With the ambition of counterbalancing the highly polarized debate and offering a more nuanced vision of what moral enhancement entails, this thesis focuses on (present and emerging) moral enhancement *practices*. The ultimate goal of doing so is to

10 Whether means matter morally has been a topic of intense debate (see chapter 2 and 6). See also Reichlin (Reichlin, 2017), and Bublitz (Bublitz & Merkel, 2014; Bublitz, 2015, 2016).

identify ethical issues that are not necessarily part of the current debate on moral enhancement.

Similarly, Harris Wiseman has advocated a 'practical-realities first' approach to potential moral bioenhancement interventions, implying that speculation about moral bioenhancement should account for "the specific practical realities to be found on the ground level, which are not at all incidental but the very realities around which the abstractions of the debate must be made to shape themselves (not the other way around)" (Wiseman, 2016, p. 13). A focus on practices can function as a correction to overly wild[11] scenarios, and at the same time call attention to practices where elements of stimulating moral betterment are visible. Much of what is being discussed is not possible yet, if ever (bringing some authors to question whether moral enhancement should be considered science fiction, rather than science fact (Dubljević, 2017; Dubljević & Racine, 2017)).[12]

Arguably, in previous enhancement debates, one of the central ethical questions concerned the desirability of expanding upon current, therapeutic interventions or practices, and to use them "beyond therapy" (Kass, 2003) or to go "beyond what is necessary to sustain or restore good health" (Juengst, 1998, p. 29). The quite extensive literature about the desirability of the non-medical use of prescription drugs meant to treat ADHD (Attention Deficit Hyperactivity Disorder) can serve as an example (Greely et al., 2008; Schermer, 2014). In the case of moral enhancement, it is often far less clear what that current practice is, and in what way the proposals made would impact

11 There is no straightforward answer to the question on how to decide whether a certain scenario should indeed be considered overly wild, or, alternatively, represents an important thought experiment. Arguably, imagination is key for scientific discovery (de Mey, 2006) and ethical analysis alike. At the same time, an important normative function of thought experiments might be to stimulate imagining of "possible usages and desirabilities" (van de Werff et al., 2016, p. 98) in order to expose and elucidate relevant ethical issues. An overly wild or speculative scenario arguably "violates conditions of intelligibility, squanders the scarce and valuable resource of ethical concern, and misleads by casting remote possibilities or philosophical thought-experiments as foresight about likely technical developments. In effect, it deflects consideration from the transformative technologies of the present" (Nordmann, 2007, p. 31, and see chapter 3). Björn Hofmann proposes to introduce a so-called *hype test,* which reads as follows: "*What is special with this technology (e.g., gene editing), compared to existing methods, that makes it succeed in improving human social characteristics to make the world a better place for all?*" (Hofmann, 2018).

12 As Stephan Schleim argues: "notwithstanding the influences of neuro-collaborations and related funding schemes within academia, we scholars indeed collectively overestimate the practical and translational social impact of this research so far. The communicated promises as well as the scholarly and public attention given to these possibilities are, in my view, in no way justified by the scientific possibilities" (Schleim, 2014, p. 3), also see (Crockett, 2014b).

on that practice. Oftentimes, what an author considers relevant practices remains implicit, or is discussed in very general terms only.

For example, one domain that is often implied, but not made explicit, is the field of forensic psychiatry. Participants in the moral enhancement debate have discussed the use of neuro-interventions for offenders or forensic patients who are suffering from various cognitive, motivational and emotional impairments as examples of moral enhancement. As such impairments may involve risk factors for various kinds of immoral behaviour (e.g., sexual crimes, violence, and racism), proponents argue that moral bioenhancement could provide new ways to achieve successful recidivism reduction and rehabilitation (Douglas, 2008; S. Carter, 2016).

However, the diversity of definitions of moral enhancement that are available suggests that not everyone will agree that psychiatric treatments that address neurobiological risk factors for deviant behaviour should be understood as proper instances of moral enhancement (Raus et al., 2014; Reichlin, 2017). Nonetheless, outspoken proponents of moral bioenhancement Persson and Savulescu have argued that a number of psychiatric disorders can be characterized as "moral defects", and therefore, that treating these disorders should indeed be understood as moral enhancement (Savulescu & Persson, 2012).

> The opposite of promoting another's interests is damaging another's interests. Traits which increase harm to others cause immoral behaviour. The paradigm is psychopathic personality disorder, but other personality disorders such as antisocial personality disorders, borderline personality disorder and narcissistic personality disorder can cause great harm to those who come into contact with these individuals. The reduction in these tendencies are thus moral enhancements. (Savulescu & Persson, 2012, p. 410)

Likewise, DeGrazia has characterized the treatment (or prevention) of antisocial personality disorder as a uncontroversial example of moral enhancement (DeGrazia, 2014), and Douglas has discussed institutions of criminal justice as institutions that are arguably "already engaged in a kind of moral enhancement" (Douglas, 2014c, p. 1245).

To conclude, intended users, contexts of implementation, and the goals and objectives of developing and applying potential biomedical possibilities for moral improvement are often not specified. In order to expose and elucidate relevant ethical issues concerning moral enhancement, as well as to formulate ethical requirements, this thesis focuses on potential *practices* of moral enhancement.

The following questions have been guiding my research:

1. What is the debate on moral bioenhancement about?
   a. What do participants in this debate understand moral bioenhancement to be? How do they define and demarcate moral bioenhancement?
   b. What is the rationale proponents give for the need for moral bioenhancement and what are the main ethical arguments given for and against moral bioenhancement?
   c. What ethical issues are most prominent in the debate?
2. What is a sensible way to approach the question of the ethical desirability of moral bioenhancement?
   a. From where should normative analysis start?
   b. For which professional practices and for whom is the debate on moral bioenhancement relevant? Who should care about/ participate in the debate?
   c. Can a focus on potential *practices* of moral enhancement help identify ethical issues that are relevant, but left out or under examined in the debate so far?
3. What are conditions and ethical requirements for ethically justifiable moral enhancement practices?

Methodology

In order to answer these research questions, in this thesis, a variety of research methods is employed. In order to get a better grasp of potential practices of moral enhancement, in this thesis I have employed a mixed methods approach, making use of empirical research methods such as qualitative interviews and vignette studies, as well as of more traditional philosophical methods such as conceptual and argumentative analysis.

In his book *Experiments in ethics* Kwame Anthony Appiah warns against considering anything experimental as alien to philosophy (Appiah, 2008). Philosophy nowadays is mostly defined in terms of what it is not – psychology, physics, anthropology. Appiah argues against this crude and a-historical view on what philosophizing consists of, and consisted of in the past: the cliché of a separate group of 'professional thinkers' that look down on anything not purely theoretical and "confined themselves to the realm of unsullied abstraction" (Appiah, 2008, p. 9). Doing so misrepresents the interests and fascinations of many of the canonical philosophers: "what's novel isn't the experimental turn; what's novel was the turn away from it" (Appiah, 2008, p. 6).

Also (bio)ethics has witnessed what has been called "an empirical turn" (Hope, 1999; Molewijk et al., 2004; Sugarman, 2004; Borry et al., 2005), indicating that increasingly normative analysis is being supplemented with empirical input, as well as the working together of more empirical and more philosophical oriented researchers. A recent systematic review of methodologies used in bioethical research demonstrates a striking heterogeneity of methodologies, both quantitative and qualitative (Davies et al., 2015).

There is an extensive literature discussing the benefits and downsides of integrating empirical research with normative analysis and appropriate ways of doing so. Strong et al. list a range of familiar advantages of incorporating empirical research in bioethics, next to philosophical reasoning (Strong et al., 2010). The strength of empirical research in bioethics, they write, is that it allows for:

> The description of the experience of individuals or populations with respect to morally relevant issues; the description of the attitudes, beliefs, moral opinions, reasoning patterns and decision-making of those involved in a certain practice; the generation of data which can challenge authority, dogma, convention, norms and experience by showing how practice varies; and the identification of moral issues that have escaped the attention of ethicists, but are relevant in a specific context, including those that are not obvious because they are embedded in practice. (Strong et al., 2010, p. 318)

Central challenges concern both the empirical methodology, and the integration of empirical and normative research. With regards to methodology, it can be questioned what kind of empirical information is most useful – the opinions, arguments or expressed moral views of various stakeholders, or perhaps rather their actual moral actions and decisions? And which stakeholders should be included? Moreover, what are the best ways to collect data on stakeholders' moral views: open interviews, questionnaires, focusgroups, or vignettes? As regards the integration of empirical data in moral reasoning a number of models have been proposed, such as "integrated empirical ethics" (Molewijk et al., 2004), "broad reflective equilibrium" (van Thiel & van Delden, 2010), "reflexive balancing" (Ives, 2014), and "symbiotic empirical ethics" (Frith, 2012), just to name a few.

In my opinion, the most challenging issue concerns the question what type of normative claim or conclusion one can generate through conducting or studying empirical research. Davies et al. distinguish two "poles of methodological orientation", one *dialogical*, the other *consultative* (Davies et al., 2015). In dialogical approaches "the ethical analysis and reaching of a normative conclusion is part of the research encounter

itself" (Davies et al., 2015, p. 8), whereas in consultative approaches ethical analysis is "undertaken post-engagement, by the researcher or the research team. (…) stakeholders feed into the ethical analysis, but are not involved in it directly" (Davies et al., 2015, p. 9). The approach to empirical data generated in the context of this thesis is mostly consultative in nature.

A review from 2012 maps empirical research on human enhancement specifically (Schuijff & Munnichs, 2012). Empirical studies on enhancement employ a diversity of methodologies, ranging from more quantitatively to qualitatively oriented, and mostly focus on studying public attitudes and preferences. The authors argue that empirical research, and research on public opinion specifically, can enrich the largely expert-only debate on human enhancement (Dijkstra & Schuijff, 2015). Because, they argue, consulting the public may aid in better understanding what shapes public preferences, and this may help in policy making. Ethical analysis of potential practices of moral enhancement indeed needs to include and account for views of relevant stakeholders such as potential users, potential medical professionals or 'operators', policy makers, and the general public. Indeed; "the question of moral bioenhancement is a paradigmatic instance in which understanding public attitudes is critical" (Reiner, 2013, 2017).

Outline of this thesis

Chapter 2 contains a systematic clustering of ethical arguments that emerge in the moral enhancement debate (based on a systematic search of the literature). Each argument is discussed separately, and the debate as a whole is assessed. It is concluded, first, that there is little discussion on what distinguishes moral bioenhancement from treatment of pathological deficiencies. Furthermore, remarkably little attention has been paid so far to the safety, risks and side-effects of moral enhancement, including the risk of identity changes. Finally, many authors overestimate the scientific as well as the practical feasibility of the interventions they discuss, rendering the debate too speculative.

Chapter 3 discusses a number of contexts or domains in which (future) moral bioenhancement interventions possibly or most likely will be implemented. By looking closely at similar or related existing practices and their relevant ethical frameworks, we identify ethical considerations that are relevant for evaluating potential moral bioenhancement interventions. Domains that are examined are, first, debates on the proper scope of moral education; second, proposals for identifying early risk factors

for antisocial behaviour, and third, the difficult balancing of the best interests of the individual and third party concerns in (forensic) psychiatry.

Chapter 4 and **Chapter 5** report on a series of interviews with forensic practitioners on the prospects, threats, and limitations of integrating neurobiological and behavioural genetic interventions in forensic psychiatric practices. Forensic practitioners were interviewed about their expectations as well as moral views regarding potential applications of current neurobiological and behavioural genetic research aiming to understand (and possibly help prevent, contain, or treat) violent and antisocial behaviour. Potential applications that were included were potential biomedical possibilities to lower aggression, the possible usage of neuroimaging in assessing legal responsibility, and the potential use of biomarkers in assessing risk for future violent and antisocial behaviour. Furthermore, these interviews explored to what extent forensic practitioners consider moral development and moral growth to be a part of their current professional practices and to what extent they think that stimulating moral development is a legitimate objective in the context of forensic psychiatric treatment. In addition, we asked how forensic practitioners balance public safety and risk management concerns with the interests and wellbeing of the individual patient.

Chapter 6 reports on an empirical study aimed at gaining insight into the reasons that the public may have for endorsing or eschewing pharmacological moral enhancement for themselves or for others. The study is driven by three hypotheses: first, the hypothesis that the degree to which members of the public support an empathy-enhancing moral enhancement program depends on whether or not the means employed are pharmacological or non-pharmacological. Second, that people are less supportive of pharmacological moral enhancement of their own children than they are of other people's children. Third, that the degree to which respondents support these programs depends on whether they imagine themselves or someone outside their immediate circle of concern to participate. In other words, we expect that the distinction between self and other is relevant.

Chapter 7 reflects on the main findings of the preceding chapters, as well as their limitations. It revisits the main research question, formulates normative conclusions, and suggestions for future research.