Practical and Normative Challenges for the Serious Conceptual Engineer

Wierenga, Niek

Practical and Normative Challenges for the Serious Conceptual Engineer

Niek J. Wierenga

Philosophical Perspectives on Politics and the Economy (MA), Leiden University

5014VTP6: MA Thesis Philosophy 60 EC

Dr. W.F. Kalf

November 2022

Word count: 19323

Contents

1 Introduction

1.1 Three significant problems and political relevance

Historically, one of the primary methods in philosophy has been conceptual analysis. However, in recent years, a new method of *conceptual engineering* has been developed. Instead of identifying what content our concepts *do* have, the new method of conceptual engineering seeks to determine what content our concepts *should* have. More specifically, the conceptual engineer's goal is to revise, improve and implement concepts based on a normative standard to which the conceptual engineer has determined that a concept is defective.

This dissertation will defend that the new philosophical method of conceptual engineering faces three significant problems if the field is understood as the method that aims to intentionally change (at least) the semantic meaning or the psychological structure of concepts that a large group of people uses. These significant problems are, first, the broadly recognised implementation problem. Secondly, the problem that I call the evaluation dilemma. Thirdly, the problem that I call the motivation problem. These problems must be solved before the new field can become what its primary participants want it to be. The political relevance of this dissertation is multifaceted. First, as I will argue in chapter II, collective long-range control is the most likely method for a successful conceptual change. Like climate change and other large-scale normative projects, conceptual engineering could be successful through collective long-range control in which the collective works together over a long time to solve a normative problem. This means that, as with climate change, most individuals have to voluntarily participate for concepts to change through collective long-range control. This is politically relevant because collective large-scale normative decisions have to be made, or the collective has to be forced.

To see the political relevance more clearly, I will provide an example. Yeonmi Park is a North Korean defector who has shared a story about her time in North Korea. Her testimony is that the people in North Korea are forbidden to use the concept of 'love' towards a person

other than the leader of their country. In other words, when individuals in North Korea use the concept of 'love', they automatically refer to 'the Great Leader' (Fridman, 2021, 20:25-22:42). This is an example of conceptual change in a collective and it shows that collective long-range control over conceptual change is at least possible through coercion and oppression. However, we intuitively understand that using coercive power to change the concept of 'love' to only refer to one dictator is morally wrong. The question is, in what way is conceptual change on a large scale desirable and justifiable?

To generalise, coercive conceptual change has already been used in history in the form of conceptual banning. Think about political regimes and religious authorities who banned particular concepts, ideas and stories. The Nazis banned an extensive list of books and beliefs. Other oppressive regimes, such as Mao's and Stalin's, did the same. Every dictator knows that free speech and free use of concepts are only desirable in a free country with free-thinking people. Suppose regimes and authorities control what concepts can be used or how particular concepts are interpreted. In that case, those regimes and authorities have power over how people conceptualise the world. Think about how enslaved people were not allowed to read or study: educated enslaved people were more likely to free themselves. Therefore, conceptual power over people has real-world consequences to real-world people.

Conceptual engineering is a new field that aims to develop a method that can change the concepts *used by large groups of people*. If a successful method of conceptual engineering is developed, the method can be used by regimes or small groups of people to control how people conceptualise the world. I will contend that, by definition, if the conceptual engineer is serious about the project of intentional conceptual change, the engineer aims to acquire the power to change the concepts that large groups of people use. In short, the success or failure of the new method of conceptual engineering has political implications for society.

Moreover, central participants in the literature make a social constructionist argument that justifies conceptual amelioration based on social justice. This is a

moral-political argument. One such example is the social constructionist argument from Haslanger (2000) to ameliorate the concept of 'woman'. Haslanger is one of the primary participants in the literature. This means that her argument has some stake in what the current literature indicates that conceptual engineering should be, which is why I use her example throughout this dissertation. Haslanger argues for ameliorating the concept of 'woman' from the biological concept towards the socially constructed concept of 'woman' in which the structural oppression of women will automatically be recognised. Haslanger's motivation for the amelioration proposal is the *moral-political goal of social justice and equality*.

In short, this dissertation is politically relevant.

1.2 Structure of dissertation

In what follows, I lay out the structure of this dissertation. Chapter I will first discuss the relatively new literature on conceptual engineering and provide some examples. Secondly, I will argue that if the conceptual engineer is serious in his endeavour, conceptual engineering ought to be understood (at least) as changing the semantic meaning or psychological structures of concepts that a large group of people uses. I will also discuss how different metaphysical theories of concepts impact the field of conceptual engineering.

Chapter II will address the implementation problem. I will show that independent of the engineer's metasemantic theory; the literature will need to solve the implementation problem satisfactorily. This is the problem that even if conceptual engineers would objectively determine how a concept ought to be improved, no practical method is known that could implement that change to concepts that will be used by a large group of people. More specifically, I will defend that either the conceptual engineer is not serious about conceptual engineering and his method is trivial and unimportant, or the engineer moves the implementation problem—he does not solve it.

In chapter III, I will address what I call the evaluation dilemma. In short, this is the dilemma that even if the conceptual engineer would have a method to implement

conceptual change on a large scale, it is still being determined how conceptual engineers could find an (objective) agreement for how a concept ought to be ameliorated. More specifically, the dilemma is that the *evaluation* of concepts—specific to conceptual *ethics*—might be complex to the degree that it is practically impossible to do. I will argue that the engineer finds him or herself in a dilemma. Accepting the first horn, the engineer has to confront the all-things-considered problem. Considering all things about a concept will turn out to be practically impossible. In contrast, by accepting the second horn, the engineer only considers a few parameters for evaluating a concept before amelioration. In this case, I will argue that the engineer will likely create more unintended problems than the initial amelioration would solve. These are the first moves on this problem in the literature: more can be said about the problem.

Furthermore, chapter IV will address one specific set of parameters on which the engineer could evaluate a concept. These are the parameters of the conceptual engineer's goals, aims and motivation. Suppose that a conceptual engineer is motivated to earn a lot of money. Is that engineer justified in engineering a concept used by a large group of people for that motivation? Intuitively, this seems to be wrong. Therefore, I will defend two arguments. First, I will show that the necessarily biassed motivation of the conceptual engineer ought to be considered before amelioration. More specifically, because the engineer's motivation is necessarily biassed, the justification for amelioration based on the engineer's subjective motivation is problematic.

The second argument considers how the method of conceptual engineering will provide a small group of conceptual engineers conceptual power over a large group of people. I propose that conceptual engineering owes a justification for constructing such a method that can influence the conceptual schemata of millions of people. Suppose the engineer is serious about his project. In this case, the conceptual engineer aims to intentionally change the semantic meaning or psychological structures of concepts used by most people. Therefore, if the conceptual engineer is successful, the engineer would obtain the power to influence how large groups of people conceptualise themselves, society, and

the world. With such power comes responsibility, and I will argue that the conceptual engineering literature owes a satisfactory justification for why obtaining such a method would benefit humanity and not become another power tool.

In short, before the method of conceptual engineering can become what its central participants hope it will become, the conceptual engineer has to overcome these challenges.

2 Chapter I - The Project of Conceptual Engineering

2.1 Introduction

In this chapter, I will argue that if the field of conceptual engineering is serious, it aims to intentionally change the semantic meaning or psychological structure of concepts that most people will use. This is because all central participants in the literature who argue in favour of a conceptual change aim to either change the semantic meaning or the psychological structure of concepts used by a large group of people. In contrast, based on my definition, I will show that a conceptual engineering proposal becomes trivial or unimportant if the engineer is not serious about the project. In order to argue this, I will first introduce the field of conceptual engineering and show that different participants in the literature define the field differently and have different motivations to participate. Secondly, I will provide examples from the literature and show that conceptual change will impact real-world people. Thirdly, I will discuss different metaphysical views on concepts accepted in the literature and argue for the definition of the *serious* conceptual engineer I will use as the basis for the other arguments in this dissertation.

2.2 Conceptual engineering

Conceptual engineering is a relatively new and fast-moving literature with contributions from many fields and subfields, including linguistics, cognitive science, psychology, history, sociology, ethics, philosophical methodology, metaphysics, and epistemology (Cappelen & Plunkett, 2020, p.15). The field is concerned with assessing and improving concepts. Historically, when philosophers debated concepts, the standard method was conceptual analysis. This method would take a concept C and analyse and define its constraints: the necessary and sufficient conditions for being C. Other philosophers would then challenge the concept through, for example, the method of cases. This method has been the primary philosophical source of evidence for the success of analysis. However, over the last decades, this standard model of conceptual analysis has increasingly been questioned for its ability to provide satisfactory analysis. The lack of satisfaction is because the method has failed to

have generated even a single successful instance of analysis (Nado, 2021, p.1508). Philosophers have found contradictions, paradoxes or irregularities in all proposed philosophical concepts.

Therefore, a new alternative method has increasingly gotten attention: conceptual engineering. Instead of analysing concepts, engineers aim to improve or replace them. In other words, engineers aim to *create* better concepts rather than explore them. However, to what degree are concepts creatable? These are metaphysical questions that are also part of the conceptual engineering project. It divides the literature on conceptual engineering into two parts. The first is case studies focusing on changes to specific concepts and advocating for specific ameliorations. The second is metaphilosophical research on conceptual engineering as a method and its foundational issues. In order to conceptually engineer concepts, it has to be understood *what* is being engineered, *how* it could be done, and whether it would be *ethical* to do so.

Conceptual engineering can be understood as intentionally changing, removing or improving concepts. Engineers take apart, examine, explore, study and try to improve objects. As we will see, some speakers argue that the classic method of philosophical analysis is also a form of conceptual engineering (Jorem, 2021). Others argue that conceptual analysis is the only method that is likely to be practically possible, making the new field of conceptual engineering trivially unimportant (Deutsch, 2020).

Central participants in the conceptual engineering literature provide a broad definition:

> *Conceptual engineering = (i) The assessment of representational devices, (ii)*
> *reflections on and proposal for how to improve representational devices, and*
> *(iii) efforts to implement the proposed improvements (Cappelen & Plunkett,*
> *2020, p.3).*

To show why the field is in its infancy and participants still have fundamental metaphysical and methodological disputes, other participants argue against the idea that conceptual engineering ought to be about representational devices. For example, some argue for

commitment engineering (Löhr, 2021). Commitment engineering is a form of conceptual engineering without engineering representations. Instead, it engineers people's commitments to concepts. Therefore, such a defence of conceptual engineering contradicts Cappelen and Plunkett's purposefully broad definition of conceptual engineering.

All of this shows the infancy of the literature and the degree to which participants have different conceptions of conceptual engineering and what conceptual engineering ought to be. Still, other participants in the literature ask precisely that question. The paper, '*How to conceptually engineer conceptual engineering?*' explores what conceptual engineering *ought* to be (Isaac, 2020). In other words, a large consensus on such issues has yet to be met, making the topic an exciting new frontier.

For example, it is argued that conceptual engineering should include intentionally changing concepts in its broadest sense (Jorem, 2021). This view on conceptual engineering includes, for example, changes in the *speaker*-meaning of concepts. Suppose that the meaning of concept A is changed to mean B in a conversation between two people, this singular act of changing the speaker-meaning of a concept is an example of conceptual engineering. On the other hand, other participants in the literature argue that the real aim of conceptual engineers should be to change (at least in part) the semantic meaning or psychological structure of concepts in a standing language that is spoken by most people (see Deutsch 2020 and Cappelen and Plunkett 2020). Which of these arguments is correct has a defining influence on conceptual engineering and what participants want the field to be.

At the beginning of this dissertation, I will use conceptual engineering in its broadest sense. I will define the term more clearly as the dissertation progresses. Chapter III will switch from conceptual engineering to conceptual ethics. This is a subfield of conceptual engineering. Conceptual ethics is concerned explicitly with concepts' (normative) evaluation. Which concepts should be used and why? Based on what evaluation standards can this be determined and why? And how *ought* a concept be used? The 'ethics' in conceptual ethics is

meant to be very broad. It is not meant to privilege moral-political norms. Instead, it is meant to include evaluation from any relevant discipline (Cappelen & Plunkett, 2020, p.4).

2.3 Goals of conceptual engineering

Conceptual engineers are also divided about *why* concepts need to be engineered. What is the motivation of the conceptual engineer? One answer is to improve concepts based on a particular motivation. However, in order to improve concepts, it has to be determined that they have to be defective in a significant way. Why would concepts be defective? Nietzsche contributed to the idea that concepts are defective and should be improved. In Will to Power, he wrote, "philosophers have accepted concepts as if they were already perfect". Nietzsche was famously sceptical of this. He believed "what is needed above all is an absolute scepticism toward all inherited concepts" (Nietzsche 1901/1968: 220–1, section 409). Conceptual engineering, in other words, is the project that is sceptical towards concepts, tries to take them apart and aims to improve them while constructing them back together. Another answer is that it might be desirable to be sceptical towards the sceptics of concepts. In the last chapter, I will consider precisely this: *should* the conceptual engineer seek a method that, if successful, will have power over how large groups of people conceptualise the world? In other words, if the field becomes successful, a small group of engineers will have conceptual power over large groups of people. I will ask whether this is something that humanity should want.

2.4 Examples of central participants

I will shortly discuss several central participants in the literature to show what these central participants want the new field of conceptual engineering to be. For example, one central participant is Haslanger (2000), whom I discussed shortly. Whether the concept of 'woman' would still concertedly be used as a biological adult human female has other real-world consequences than if individuals would use Haslanger's ameliorated concept of 'woman'. Summarised, Haslanger's ameliorated concept of 'woman' is a member of a social class who is structurally socially subordinated to men based on their presumed or perceived female

biological role in reproduction (Haslanger, 2000, pp.42-43). At least, Haslanger's idea is that if her conceptual amelioration of the concept of 'women' is collectively accepted, social reality would become equal and just. Social change is one of the reasons why the project of amelioration could be desirable: to bring about real-world change based on collective normative goals. Haslanger aims to change the semantic meaning and psychological structure of the concept of 'woman'.

Another participant in the conceptual engineering literature, Joshua Habgood-Coote, argues that the concept of 'fake news'—a widely used contemporary concept—should not be used because the concept suffers from three defects (Brown, 2019). Brown (2019), in contrast, argues that the concept should be used carefully. This is because it allows us to speak about an issue that she believes needs to be addressed. Without the concept she defends, our communication would lack a concept that expresses something valuable. This is an argument about how a concept should be used psychologically.

Another amelioration proposal aims to improve the concept of 'truth' (Scharp, 2021). This argument is based on the finding in consistency theory that there are twelve logical platitudes for the concept of 'truth' which are mainly inconsistent with each other. Based on this inconsistency, there is not one property of being true. Scharp is a conceptual engineer who is motivated to find one such property of being true.

As we shall see in the remaining part of this chapter, it becomes relevant that these central participants in the literature aim to change the semantic meaning or psychological structure of these terms in society and that this, therefore, should be the goal of the serious conceptual engineer.

2.5 Concepts

It is philosophically unclear what metaphysical reality concepts have. As described by most psychologists of concepts, are they mental representations that guide people's categorisations? Or, as the philosophical conceptual analysts would insist, are concepts the meaning of a term that 'expresses' a concept? It is concepts that conceptual engineers

purport to engineer, so what are concepts *exactly?* The fact that this is unclear shows the multiplicity of interpretations about conceptual engineering and what function participants believe the field ought to have. There are disagreements on multiple dimensions in the literature: the metaphysical, practical, theoretical, and methodological.

I have discussed the general literature on conceptual engineering and shown that the field is in its infancy. Almost all metaphilosophical and methodological aspects of conceptual engineering are still being debated. That is why I will argue for an interpretation of what a successful conceptual engineering project would be. Different participants in the literature provide different answers to this question, yielding different theories and practical tools for what the conceptual engineer ought to do. However, the primary participants in the conceptual engineering literature who aim to conceptually engineer specific concepts can be characterised by either semanticism or psychologism. Semanticism is the approach to conceptual engineering that argues that the linguistic meaning of words and concepts in a language must be changed for the engineer to succeed. In other words, the engineer aims to change the semantic meaning (sometimes called the *standing* meaning of a concept) of words in a standing language. The second approach is psychologism. Engineers engaged in this engineering form aim to change the psychological structures that explain our cognitive and linguistic behaviour. To do conceptual engineering on this view is to change how people classify their experiences and under what circumstances they use certain linguistic expressions. In short:

> *Semanticism: To engage in conceptual engineering is to advocate and implement changes in what our words mean (Koch 2021b, p.1960).*
>
> *Psychologism: To engage in conceptual engineering is to advocate and implement changes in how people use words, classify objects around them, and draw inferences about those objects (Koch 2021b, p.1961).*

Based on the primary contemporary participants in the conceptual engineering literature, I argue that if the conceptual engineer is serious in his project, the engineer aims to at least change the semantic meaning or psychological structure of concepts. In contrast, other

participants in the literature defend that conceptual engineering could also, for example, include changing the speaker-meaning of concepts or adopting a new language altogether (Jorem, 2021). The speaker-meaning of a concept is the meaning that an individual speaker gives to a concept. This latter view makes conceptual engineering trivially unimportant.

2.6 The serious conceptual engineer

Thus, some participants in the literature defend that changing the *speaker's* meaning and referent of concepts is enough for a successful case of conceptual engineering (Jorem, 2021; for a counterargument, see Deutsch, 2020). In this case, a speaker can use a particular concept and mean or refer to something different than the particular semantic meaning of that concept in a language. For example, a woman could stipulate that, in the following conversation, when she uses the concept 'plate', she shall mean the concept 'cup'. The stipulation can be coherent and understandable for both participants in such a conversation. This is important because people are more or less free to use the concepts that they use to speaker-mean and speaker-refer to whatever they want any concept to mean. However, they are not free to change the semantic meanings of concepts in their language (Deutsch, 2020, p.3939).

A person can intentionally use a concept and mean something that is not generally in the extension of that concept. In the example of plates and cups, if the woman talks about plates in that conversation, she speaks about cups, which is evident to everyone who witnessed her explicit stipulation. However, the concepts of 'plate' and 'cup' still mean what they have always meant in English. The meaning will not change because of an individual's stipulation. The woman changed the *speaker's* meaning and referent of the concept, not the *semantic* meaning and referent. It is unclear how such semantic change can come about.

For clarity, *stipulative introduction* is something that philosophers have always done. Instead of changing an already established semantic concept, they introduce one. Such introduction has been a common practice and is practically possible because no semantic meaning of the

term has yet been established. Nevertheless, if a semantic meaning is already established, it is unclear how a speaker can change its semantic meaning.

However, most participants in the contemporary conceptual engineering literature accept that a successful project of conceptual engineering is, at least, the intentional engineering of the semantic meaning or psychological structures of concepts that are used by a large number of people (Deutsch, 2020; Jorem, 2020; Nado, 2021; Shields, 2021; Simion, 2018; and Cappelen, 2018). This is important because I base the arguments in this dissertation on the following definition.

> *The Serious Conceptual Engineer: the conceptual engineer who aims to conceptually ameliorate—at least in part—the semantic meaning or psychological structures of concepts used by most people.*

This is because, as I have shown, engineering other aspects of concepts—for example, the speaker-meaning of terms—reduces the new literature on conceptual engineering to something *trivial and unimportant* because the method would be "neither a new nor neglected philosophical method" (Deutch, 2020, p.3935). Moreover, as I have shown earlier in this chapter, the central participants in the literature who have proposed contemporary conceptual ameliorations are all motivated to change at least the semantic meaning or psychological structures of concepts.

This means that based on the serious definition of conceptual engineering, for example, commitment engineering is also a viable form of engineering for the serious conceptual engineer (Löhr, 2021). This is because commitment engineering changes the psychological structures of people surrounding a concept. If such a psychological change is implemented by most people in a standing language, based on the definition of the serious conceptual engineer, the project is successful.

Finally, a successful case of conceptual change is a change to a concept used by most people. 'Most people' means a majority of people in a standing language. If most people in a standing language adopt a conceptual change, then the semantic meaning of that concept has changed. This is because most speakers of that language have adopted a new meaning

for a particular concept. In other words, they mean a different concept when they use the same term.

2.7 Conclusion

To conclude, based on the central participants in the conceptual engineering literature, this dissertation accepts that to be serious about the conceptual engineering project, the conceptual engineer is at least motivated and capable of intentionally changing the semantic meaning or psychological structure of a concept that most people will use. In the chapter that follows, I will argue that if the engineer is serious about his project, the engineer finds him or herself an implementation problem.

3 Chapter II - The Implementation Problem

3.1 Introduction

This chapter will argue that even if conceptual engineering is objectively possible, the engineer does not have a method to *implement* conceptual change. As I will show, the only method likely to work is "collective long-range control" (Koch, 2021a, p.329). I am borrowing from how other contemporary large-scale normative projects are attempted to be solved. Think about reducing plastics, carbon dioxide, or the meat industry. I will conclude the chapter by saying that collective long-range control of conceptual change would be possible through coercion or conceptual activism and I will show why these methods are not justifiable. Instead, in a truthful and free society, conceptual conflict becomes a vital problem for the conceptual engineer and this leaves the implementation problem unsolved.

In order to make this argument, first, I will discuss how the implementation problem is understood in the literature. Second, I will argue against Jorem's arguments in which he aims to solve the implementation problem. I contend that his four arguments make the project of conceptual engineering trivial and unimportant. Thirdly, I will discuss the externalist challenge because the implementation problem could depend on the concept's metaphysical reality. I will conclude that the implementation problem is independent of the engineer's metaphysical theory of concepts. To show this, I will contrast the social constructionist view of concepts with the externalist view. Following that, I will argue for one solution to the implementation problem that does not make the field of conceptual engineering trivial and unimportant: collective long-range control. Similar to how contemporary society aims to solve other large-scale normative issues, this method is most likely to be practically feasible. However, even if this method is practically possible, the method is likely to suffer from either top-down coercion or conceptual conflict. Finally, I will show that this is undesirable by discussing the method of conceptual activism that aims to create conceptual change by ignoring the problem of conceptual conflict.

3.2 The implementation problem in the literature

The implementation problem is widely recognised (Burgess & Plunkett, 2013a; Cappelen, 2020; Deutch,2020; Gibbons, 2022; Jorem, 2021; and Koch, 2021a). Burgess and Plunket (2013a) were the first to recognise the implementation problem. They are two of the leading participants in the field, together with Cappelen, who also recognises the severity of the problem and concludes that he does not know how, " nonetheless we will and should keep trying" to make conceptual engineering work (Cappelen, 2018, pp. 72-75).

In short, the implementation problem is the problem that even if the conceptual engineer would establish that a concept ought to mean B but did mean A, the engineer does not know the steps the engineer has to take to implement the change from A to B. In other words, the implementation problem is constituted by the following question: *how* would it be possible to implement better concepts that most people will use?

When engaged in changing the semantic meaning and reference of concepts, engineers are either ignorant of how to implement conceptual change or, to the degree that conceptual engineering is already being practised, the literature on conceptual engineering has been trivialised. In the latter case, conceptual engineering is nothing else than conceptual stipulation or the old method of philosophical analysis. As I have shown previously, the former is the intentional introduction of concepts to influence an audience. The latter is simply the old method of describing concepts more precisely. If either counted as a successful instance of conceptual engineering, it would make the new literature on conceptual engineering trivially unimportant.

Semantic change is not *impossible*, but there is an implementation problem for intentional semantic change to a concept used by most people. It still needs to be discovered how to change the semantic meanings of concepts. It needs to be more methodologically clarified precisely how many, how often, or how long speakers must persist in using intentional stipulations in order for the semantic meaning of a concept to change (Deutsch, 2020, p.3940). Even if conceptual engineers have determined that the meaning of a concept ought to have a different semantic meaning and referent, it is unknown how to

implement that amelioration. This is because if a concept stands in a language, that concept is used by most people. Hence, a semantic or psychological change to a concept has to be changed by a large number of people. All those people—or most people in a standing language—need to change the meaning of a concept they use. How change can be brought about is unclear. Moreover, if the engineer is serious about conceptual change, the engineer aims to change concepts' semantic meaning or psychological structure. Changing the speaker-meaning of a concept is not enough.

Thus, two points have been established: first, conceptual engineering should be about the semantic meaning and reference of concepts, and two, to the degree that conceptual engineers aim to change the semantic meaning of concepts intentionally, they owe an explanation of how this can practically be done. In what follows, I will refute the best attempt to solve the implementation problem, contended by Jorem (2021), and argue that if the conceptual engineer is serious in his project, Jorem's four defences against the implementation problem fail.

3.3 Solving the implementation problem by broadening its definition

Some participants defend that conceptual engineering does not have to be the improvement of the semantic meaning or reference of concepts (Jorem, 2021). Instead, it can be contended that a successful amelioration can involve implementation methods different from changing concepts' semantic meaning or psychological structure. In what follows, I will lay out these arguments more closely and argue why they fail if the conceptual engineer is serious about the project of conceptual change. In that case, the project needs to include the intentional semantic change of concepts. Instead, the arguments that I refute make the project from the conceptual engineer trivially unimportant.

Jorem explicitly accepts that changing concepts' semantic or standing meanings is impossible from the individual perspective (2021, p.193). Only what Deutsch has called the speaker-meaning of a concept can be changed by an individual or a small group of people. Instead, according to Jorem, it is only under the control of collectives to change the semantic

meaning of a concept. However, this is the implementation problem: it seems beyond practical feasibility to change the semantic meaning of a concept for a large collective of people. At least, no such practical method has yet been established. The *semantic* meaning of a concept is also called the *standing* meaning of a concept in a language. The idea of the *standing* meaning of a concept is that there is something about the concept that is stable across speakers, times, and places. If the meaning of concepts could instead be easily manipulated, the stability of human communication would suffer. However, based on his definition of the implementation problem, Jorem argues that the problem is insignificant insofar as one of his four proposed implementation methods stands.

Jorem accepts that it is likely *not metaphysically impossible* to change the semantic meaning of concepts (2021, p.194). That is why Jorem argues that his argument works both for externalism and internalism. If the standing meaning of a concept were only determined based on internalistic factors, it would still be *practically* near impossible to change large groups of people's intentional states. These are the internal states of mind in which something can be represented. In order to, for example, change the meaning of an English concept by changing people's internal states, it would mean that conceptual engineers would have to engineer the mental states of around 1.5 billion English language speakers in the world. The only viable method that could work is collective long-range control. I will explain what this is later in this chapter.

In contrast, if only two people spoke a language, the concepts in that language could easily be changed. In this language, the speaker-meanings of concepts are also the standing meaning because stipulative agreement could change the speaker-meaning of concepts into the standing meaning of the language with only one agreement. If two people were to be in a conversation in a language only those two spoke, then through one single agreement, it is easy to see that they can change any standing meaning of any concept to what they, as speakers, agree at the moment. Therefore, if such a project is successful, the implementation problem does not stand for minor languages or small isolated groups. The question then becomes, how large could such groups become before the implementation of

conceptual change becomes a practical problem? Still, if the conceptual engineer is serious about the project of conceptual change, I have argued that the engineer has to change concepts *used by most people*. Therefore, this argument does not solve the implementation problem.

The second proposal is meaning modulation, a theory taken from Peter Ludrow (Jorem, 2021, p.199). Ludrow defends that the standing meaning of concepts often needs to be fixed. Instead, the standing meaning of concepts is more dynamic than most people commonly accept. Following this, Ludrow argues that if a speaker changes the extension of a concept and the hearer in the conversation accepts this change, then the word's meaning is modulated. If this happens to multiple concepts in a conversation, the pair speak what Ludrow calls a micro-language (Jorem, 2021, p.199).

This argument fails to solve the implementation problem if the conceptual engineer is serious in his project and aims to change concepts that most people speak in a language. The changes to the semantic meaning of concepts of Jorem's proposal are conversation-by-conversation based. It is not an attempt to change the semantic meaning of concepts that *most people will use*.

Finally, if the engineer is serious about the project of conceptual change, Jorem's third and fourth proposals to escape the implementation problem need to be more convincing. The third argument that Jorem provides is a change in speaker-meaning. I have already refuted this argument and shown that it makes the project of conceptual change trivially unimportant. Nevertheless, I will provide an example from Haslanger's amelioration proposal. Her project of improving the concept of woman for the cause of social justice seems to fail if her amelioration project is understood as an attempt to change the semantic meaning of the concept of 'woman' through stipulation. Haslanger's ameliorated concept of woman is:

> *S is a woman iif S is systematically subordinated along some dimension*
> *(economic, political, legal, social, etc.), and S is "marked" as a target for this*

treatment by observed or imagined bodily features presumed to be evidence

of a female's biological role in reproduction (Haslanger, 2000, p.39).

Suppose that it is true that the semantic extension of the concept 'woman' in the English language would still include people who are not subordinated based on perceived biological features indicating a female role in reproduction. If this is the case, Haslanger's ameliorative project will not have changed that this is true in the standing English language. In other words, Haslanger can *speaker*-refer to subordinated women when she uses the concept of 'woman', but she can *speaker*-mean anything with the concept of 'woman'. It is, therefore, questionable whether this change is a conceptual improvement (Deutch, 2020, p.3943). The point is that even if Haslanger convinces most people to speaker-mean something, it is unclear when and if the semantic meaning of the concept would also change. If conceptual engineering would simply be the change of *speaker meanings*, then the exciting terminology of *conceptual engineering* would not entail what its visionaries would like it to be. Again, stipulative introduction is something philosophers have always done and are still doing. If this is all conceptual engineers can do, the field will be trivial.

Jorem's fourth option is to start speaking a different language altogether and, in this way, be able to implement new semantic meanings for all concepts. However, this will only result in a new implementation problem: bringing a collective of people to learn and speak a new language. Therefore, this is also not a viable option.

Thus, these proposals do not solve the implementation problem if the conceptual engineer is serious about his project. As I have shown, Jorem defines the implementation problem so that his definition is unclear about the concepts' properties that the conceptual engineer would be engineering. However, most participants in the literature disagree. Here, conceptual engineering is understood as engineering (at least in part) the semantic meaning and reference or the psychological structures of concepts that are used by a large number of people (e.g. Burgess & Plunkett, 2013a; Cappelen & Plunkett, 2021; Deutsch, 2020; Koch, 2021b). If conceptual engineering instead would be about engineering, meaning modulation, then the project of conceptual engineering cannot be what the literature hopes

conceptual engineering would become. This means that the implementation problem stands if the engineer is serious about what we want the conceptual engineer to accomplish.

3.4 The externalist challenge

This section will explore whether a metaphysical (metasemantic) answer about concepts could solve the implementation problem. More specifically, I will explore the externalist challenge to conceptual change. This is the idea that if semantic externalism is true, conceptual change is impossible because semantic meaning exists external to humans and is, therefore, beyond human control to change. I will argue that the implementation problem is independent of the metasemantics of concepts: the implementation problem stands whether semantic internalism or semantic externalism is true. However, I will contend that semantic externalism's truth makes conceptual change practically harder to implement.

Burgess and Plunkett (2013a) have been the first to recognise the implementation problem in the conceptual engineering literature. They made the realisation in relation to externalist metasemantics. They recognised the externalist challenge considering semantic externalism—the view that semantic properties of a concept depend on some property in an external environment. Conceptual engineering can be nothing else than the increasingly precise discovery of those external properties. If that is true, then "deliberate conceptual change might be unnecessary" (Burgess & Plunkett, 2013a, p.1098). Instead, the philosophical project would already be conceptual engineering, but this engineering would be nothing else than discovering those external properties. This would make the conceptual engineering project seem different than it implies because *engineering* implies the conscious creation of something new. Instead, it would only be possible to discover what already existed in an external environment, and the term conceptual *engineering* would be conceptually inaccurate.

This is the metaphysical discussion about what concepts are or what (natural) properties they have. Moreover, how does this affect the engineer's ability to implement conceptual change? For example, to show how this is considered in the literature, Burgess

and Plunkett (2013a, p. 1093) write that the best methodology forward might be to let first-order metaphysics lead the way in the conceptual engineering debate. In this way, the methodology of conceptual engineering could be based on the knowledge of what relationship concepts have with reality. If the engineer understands what metaphysical reality concepts have, the engineer might be able to solve the implementation problem.

The notion of *carving nature at its joints* is central in this debate (See Simion, 2018; Burgess & Plunkett, 2013a; Cappelen, 2018; Haslanger, 2012; and Plunkett, 2020). This is the metaphysical notion that nature, or objective reality, is constituted of fundamental niches to which our concepts refer. Carving at the joints means that there is a precise niche in a body where the carving is most effortless for natural reasons: it is way harder to carve through the bones. A practical example could be given with the concept of 'dog'. There is a natural reason why the concept of 'dog' exists. A concept that would not, for example, carve nature at its joints is the concept of 'a dog born on Tuesday'. This is also a valid concept. Why would this concept be more valid, or less valid, than the concept of 'dog'? Here the notion of *carving nature at its joints* comes in. The concept of 'dog' carves nature at its joints because it can be said to have a *natural* extension. The extension of a concept is the class of entities to which the expression can be applied.

In contrast, the addition of 'born on Tuesday' is a non-natural extension because it reduces the concept of 'dog' to a concept that is only useful in specific (non-natural) cases. Therefore, the concept of 'dog born on Tuesday' does not carve nature at its joints. The notion of carving nature at its joints assumes an externalist metasemantic theory of concepts.

If such an externalist metasemantic theory of concepts is true, it is understandably harder to implement conceptual change. If, for example, concepts carve nature at its joints, then concepts seem to be fixed to these joints, and we can only change semantic change by changing those joints in nature. Externalism does imply that the semantic shift of concepts does require a corresponding shift in the external environment. The real problem, therefore,

remains implementation. Even if externalism does not threaten semantic change, it is still unclear how a change would be implementable.

Thus, the externalist challenge states that if an externalist metasemantic theory of concepts is true, semantic change to concepts seems impossible. However, I will consider arguments that solve the externalist challenge (Koch, 2021a). It is argued that externalism constitutes a problem for implementing conceptual change because it is not *easy* to implement semantic change. However, in the end, I will argue that conceptual engineers have 'collective-long-range control' over the semantic meanings of concepts.

More specifically, Koch discusses three distinct responses to the externalist challenge. His first argument is the possibility that the semantic types that the conceptual engineer wants to change are not influenced by semantic externalism. For example, the engineer would not want to change proper names or natural kind concepts semantically. However, if this were true, concepts such as 'truth', 'knowledge', 'belief' and 'woman' could not be engineered[1]. And it is precisely these concepts that engineers wish to ameliorate in the literature. In other words, this would render much of conceptual engineering redundant.

The second argument is constructed by Cappelen (2018). As I noted previously, he argues that externalism does pose a serious threat to meaning control in conceptual engineering but that "nonetheless, we keep trying; there's a sense in which we can't give up." (pp.72-75). This argument needs to be stronger because for the project of conceptual engineering to bear any fruit, as I have argued, an intentional method of semantic meaning control is needed. In other words, according to Cappelen (2018, pp.72-75), to what end should the engineer keep trying? Simply keep trying, as it is not specific to the aims of the

---

[1] I am aware that many contemporary academics do not agree with the claim that the concept of 'woman' is a natural kind. In the next section, I am going to discuss one such view. However, I base this claim on four large-scale and well-cited social science studies that empirically prove that social constructionism about gender is wrong (Bihagen, 2000; Costa, 2001; Schmitt, 2009; Schmitt, 2017). These studies show, contrary to the predictions of the social role model of gender, that in cultures where traditional sex roles are *minimised* gender differences *maximise*. These findings prove that gender has a biological basis and is not fundamentally socially constructed. Therefore, empirical fact is that gender is at its basis a natural kind. This does not mean that there is no social role *element* to gender at all, or that philosophers cannot dispute the validity or acceptability of the empirical sciences. This is precisely the question that I believe to be at the root of the conceptual engineering literature: to what degree is the world socially constructed? And if it is, to what degree can we engineer the world to whatever we want it to be? Is the sky the limit? Or is nature the limit?

engineer. If it is not defined what the engineer *ought* to accomplish for a successful amelioration, how can he become successful, to begin with? Or does Cappelen contend with the engineer changing, for example, the *speaker*-meaning of a concept? In this case, conceptual engineering is not something new or exciting.

3.5 Social constructionism

Contrary to semantic externalism is the metasemantic internalist theory of social constructionism. Cantalamessa (2021) and Haslanger (2000; 2021) assume that we, as a collective, have collective-long-range control over semantic or psychological change in the same way we aim to control climate change. They are already in the process of trying to implement their conceptual amelioration projects practically. For example, in her paper, *Gender and Race: (what) are they? And (what) do we want them to be?,* Haslanger (2000) defends a social constructionist defence of gender. She argues that gender is, at least for the most part, socially constructed.

Moreover, to the degree that gender is socially constructed, it needs to be ameliorated. This argument is based on a fundamentally different metaphysical theory than semantic externalism. To say that something is socially constructed implies that it depends on contingent aspects of our social being. If cooperation in society did not exist, we would not have built or believed it. Alternatively, if our society had been different, we would have built or believed differently (Boghossian, 2021, p.1). This includes our conceptual schemes. This is because it is necessarily true that our concepts would not have existed if society did not exist. This is an internalist perspective. It does not necessarily claim that concepts do not have *some* relation to an external environment; this is indeterminate. However, concepts have at least a socially constructed *element* and the degree to which a concept is believed to be socially constructed depends on the speaker. It is about this element of a concept that both Haslanger and Cantalamessa construct their arguments.

More specifically, Haslanger argues that if the collective ameliorated the socially constructed aspect of the concept of 'woman' to her ameliorated concept of 'woman',

society would become more socially just. The question might be asked to what degree a concept is socially constructed. This is the open metaphysical question that the conceptual engineer would like to have answered. This is because, to the degree that concepts are socially constructed, it would be easier to ameliorate that part of the concept compared to the part of the concept that has an externalist relationship with an environment. It is easier to change the social part of ourselves than an external reality over which human beings do not have direct control.

However, as Deutsch points out, the engineer still faces an implementation problem, even to the degree that concepts are socially constructed. Instead of reality that would have to shift according to the semantic meaning in semantic externalism, the conceptual engineer of socially constructed concepts still has to change the internal mental states of large groups of people. Therefore, it could turn out that socially constructed concepts could be metaphysically ameliorated in contrast to concepts with external reality. However, the engineer would still not have a practical method to implement such change.

Moreover, the social constructionist of concepts also has to deal with the problem of conceptual conflict. This means that the social constructionist also has to negotiate with other engineers who are likely to disagree. This means that it does not matter whether semantic externalism about concepts was true or to the degree that concepts are socially constructed; the conceptual engineer has an unsolved implementation problem.

3.6 Collective long-range control

Thus, social constructionism to concepts still faces the implementation problem. This is in line with Koch's final argument; he defends that externalism does not preclude the meaning control of concepts. Instead, he defends an argument that I will defend to be most likely as a method that will enable the conceptual engineer to change concepts. Independent of concepts' metaphysical reality. This argument is in two parts. First, he argues that we collectively assume that we have collective-long-range control when motivated to engage in normative projects. Think about reducing plastics, CO_2 and the consumption of meat.

Similarly, Koch defends that we can influence the standing meaning of semantic concepts. Unlike Koch's other two suggestions, this argument has merit. We, as a collective, assume that we have collective-long-range control over large normative projects. Think again about reducing plastics, CO2, or eating meat. If enough people would participate in collective action, collective long-range control of normative projects is possible. A similar strategy could be valid for concepts' semantic meaning or psychological change. Moreover, this method works for most people because it is *collective* control.

The second part of Koch's argument combines the theory with a *causal theory of reference.* This is the theory that a name's reference becomes fixed through an original act of naming. If the name is used later in history, it successfully references the original "baptism" (Kripke, 1980, p. 96). Simply put, facts that happened in the past ground the references in our language. Koch argues that such a causal theory of reference does allow for reference change in some semantic externalist theories. This means that even though semantic meanings and references are to some degree externally fixed, the meaning and reference can change through collective long-range control. Simply put, the idea is that if we start to use concepts "*as if it had new reference",* it will change the body of information that is the causal source used for a concept. If enough people keep doing that for an extended period, collective long-range control is possible (Koch, 2021a, p.343).

Of all methods, this method is most likely to be practically feasible. This is because, first, some collective long-range control over large normative projects is possible. For example, climate change is an example of how large-scale normative projects can be solved to some degree. Think about the ozone depletion that humanity has been aware of since 1970. This was a large-scale climate problem that has been fixed. In 1989 a ban came into effect to stop ozone-depleting chemicals, stabilising the ozone layer by the mid-1990s and starting the ozone recovery by the 2000s. This means that large-scale normative projects are possible.

Second, if the engineer is serious about his project, the engineer aims to change the concepts used by most people. This would mean that if the collective of a standing language

would work to change the semantic meaning of a concept in a standing language, there is no direct reason to assume that this would not be possible. Remember that the implementation problem is only the practical problem of implementing conceptual change. Moreover, as Koch shows, if a causal theory of reference version of semantic externalism were true, the project of the conceptual engineer would be saved from the implementation problem. However, in the next section, I will show that even if collective-long-range control over semantic or psychological change is possible, the engineer would still have to agree with all agents involved. This is the next problem that the engineer faces.

3.7 Conceptual conflict

Suppose that semantic externalism is accurate and collective long-range control over semantic meaning and reference is possible. The engineer would still face *conceptual conflict* (Gibbons, 2022)*.* This is the degree to which agents disagree with the proposals from conceptual engineers and therefore try to block their amelioration proposals. In other words, the 'collective' in collective long-range control must involve most agents in implementing the change. Those agents cannot internally disagree because this would reduce the collective to multiple conflicting groups. The problem would be conceptual agreement. To see why this is the case, philosophers nearly always conflict with conceptual proposals. It is rare that philosophers—or even two random individuals—agree on conceptual schemes. Even if concepts would carve nature at its joints, it is easy to see why conceptual agreement over a large collective is unlikely. In contemporary society, there is conflict about fundamental scientific truths. The question would be: how would it be practically possible to make a collective of people speaking a particular language agree on anything? Even what some would call well-established empirical facts, such as *a woman is a biological human female*, are questioned by a substantial group of people. If even such well-established empirical facts are questioned and not accepted, on what other bases could a collective agree on anything? It is unlikely that they can. This means that conceptual agreement in a collective is also improbable.

In conclusion, Gibbons' solution moves the implementation problem to a different place; he does not solve it. This means that even if collective long-range control over conceptual change is possible, I cannot see how a collective can be formed that would agree on that conceptual change without the use of force. This leaves the implementation problem unsolved, and the conceptual engineer has yet to construct a practical method that would enable conceptual change if the engineer is serious about the project.

3.8 The continuity problem

The continuity problem is the problem that implementing conceptual change could be impossible because a change in the meaning of a concept will result in a different concept (see Nado, 2021; Knoll, 2020; Jackman, 2020a; Jackman, 2020b; and McPherson, 2021). As I have already described, if Haslanger's ameliorative of the concept of 'woman'—from the biological concept to the socially constructed concept—would be successful, one possible view is that she changed the meaning of the concept of 'woman' that most speakers in society use. One of the problems is that Haslanger's argument to change the concept of 'woman' will not affect the *need* for the biological concept of 'woman' that would still be functional. Biologists and archaeologists will still want to have a concept that refers to the contemporary biological concept of 'woman' because it is highly *functional* in those fields and their respective literature. To the degree that the old concept was functional for some purpose, it will need to be replaced.

Not only replacement is needed. This argument shows that if a concept's semantic meaning or psychological structure is ameliorated to a particular meaning or structure, it is simultaneously ameliorated away from another meaning or structure. This means that the amelioration of concepts simply changes the concept to a new concept. The niche that the old concept referred to remains. To the degree that niche was practically valuable to a group of people, that niche will need to be represented again. This makes conceptual conflict more difficult to solve: it needs to be clarified what exactly is being ameliorated in conceptual change. To see this more closely, I will consider social constructionism about concepts.

3.9 Conceptual activism

If collective long-range control over concepts is possible, but conceptual conflict is the only obstacle for the conceptual engineer, then the method of conceptual activism could be a solution. *Conceptual activism* is a method proposed by Cantalamessa (2021). This is the activity of aiming to implement conceptual change by making claims that are believed to be "trivial, overstated, incomplete, or simply false, as a means of pragmatically subverting the received use of a concept and to encourage their audience to lower their degree of confidence in their current concept or conceptual framework" (2021, p.49). The final aim of the conceptual activist is: "to get audiences to react in a certain way in order to get them to reflect on some aspect of word use in order to instigate conceptual revision and, by extension, social reform" (2021, p.54).

In other words, conceptual activism does not care about finding an objective ground that could justify conceptual amelioration. Instead, it's a method that uses any means necessary to instigate the social reform that a particular conceptual engineer desires. The internal concepts that people use become means towards a hypothetical good. To change such concepts, deceit is justified.

This, I defend, is unjustified because of the following reason. If a person is willing to use instrumental untruths, it cannot be known whether that person is truthful about his or her stated motivations. If a speaker is willing to use intentional untruths while claiming to desire social reform and justice, on what basis can that speaker, and readers of that speaker, believe in his or her initial motivation? They cannot. That is why Cantalamessa cannot be trusted with her stated motivation to improve the lives of disabled people. I would argue that Cantalamessa cannot trust herself on this stated motivation. This is because if she is willing to use instrumental lies to gain control over the concepts by which people—most notably, herself—conceptualise the world, it is unclear, even for herself, about what claims she is truthful and about what claims she is not. In other words, Cantalamessa, her own conceptualisation of the world, is infested by untruths and no person should adopt or believe her motivation for any conceptual amelioration. This means that, in short, the

method of conceptual activism is unjustified because the willingness to use instrumental untruths for the sake of a moral good cannot be justified. I will consider this problem more deeply in chapter IV of this dissertation.

Note that this view seems incompatible with conceptual externalism because externalism implies that the semantic shift of concepts requires a corresponding shift in the external environment. So, either the shift in the external environment has in some way to directly follow the intentional untruthful practice of conceptual activism or conceptual activism is incompatible with externalist metaphysics. As Koch described it, it is unlikely that an external environment can shift based on untruths. However, Cantalamessa (2021, p.46) is concerned with "challenging biassed conceptions of disability," which is the socially constructed aspect of society's *biassed conception* of disability. Therefore, Cantalamessa's claim is not about the part of the concept that could have an external relationship to nature.

3.10 Conclusion

Thus, if the conceptual engineer is serious about conceptual change, the engineer's aim should be to intentionally change concepts' semantic meaning or psychological structure. I have assessed Jorem's arguments to solve the implementation problem and shown that if the engineer is serious about the project of conceptual change, Jorem's arguments do not stand. Following this, I have explored the role of metaphysics in the implementation problem. I contended that semantic externalism does make semantic change more difficult but that collective-long-range control over semantic externalist conceptual change is the most likely method that will practically work. However, if the engineer supposes that this method can facilitate intentional semantic change, I have argued that the engineer faces a new problem: conceptual conflict. Following that, I have discussed the internalist metaphysical theory of Social constructionism and shown that this theory also faces the implementation problem. Likely, collective long-range control can work for a concept's socially constructed aspects. Although, conceptual conflict is also a problem here. Finally, I have considered Cantalamessa and argued that her method of conceptual activism cannot

be trusted on its claim to be for social justice and should, therefore, not be considered a useful method. In short, the conceptual engineer, independent of his metaphysical beliefs, faces the problem of conceptual conflict in implementing conceptual change.

Thus, the implementation problem stands**,** and the conceptual engineer lacks a method to implement conceptual change. Cappelen, one of the prominent participants in the literature, writes: "more often than not, what we do [to implement conceptual change] will either have no effect or will have some effect other than what we intended." However, he is firm in that "we cannot give up, and we have to keep trying" (2020, p.75). This is because Cappelen argues that even if we do not know how to implement conceptual change, it might be possible that the proposal of our engineered concepts would be the concepts we *ought* to be using. This is true: even though we do not understand *how* conceptual change can be implemented, it could be the case that we *ought* to find out. For example, Haslanger (2000) could be correct in her proposal that we ought to change the psychological structure of the concept of 'woman' in our society for the virtue of social justice. Therefore, whether the implementation problem stands does not mean that the conceptual engineering project is unjustified or undesired.

That is why I will continue with a related problem that the conceptual engineer faces in chapter III. This is what I will call *the evaluation dilemma.* I will defend that the evaluation dilemma is one of the causes of the problem of conceptual conflict.

4 Chapter III - The Evaluation Dilemma

4.1 Introduction

In this chapter, I will argue that the conceptual engineer faces the evaluation dilemma if the engineer aims to change the semantic and psychological content of concepts that large groups of people use. This is the problem that even if we know how to implement conceptual change, it is unclear how the engineer could come to an objective justification for amelioration if conceptual amelioration is based on the engineer's motivation. If it is determined that a concept ought to be ameliorated, it is necessarily the case that the concept is deemed defective relative to some standard or value. However, as I will show, there are many different standards and values relative to which concepts can be evaluated. Why should one be more important than another? This is based on the motivation of the conceptual engineer. Moreover, engineers who appreciate different standards and values will find themselves in conceptual conflict about how a concept ought to be evaluated. Others have noted the plurality of values by which concepts can be evaluated as "the plurality of parameters" (Burgess & Plunkett, 2013b, p.1104).

I will start by introducing two points that will prove helpful in the discussion. Second, I will consider both horns of the evaluation dilemma. Either the engineer considers *all things* about a concept. This is the first horn, and I will show that this horn turns out to be practically impossible. Alternatively, by taking the second horn, the engineer considers a small set of parameters of a concept. I will show that this will likely lead to unintended consequences for other aspects of the concept. It will result in a conceptual amelioration that will likely produce more problems to the concept than the initial amelioration would solve. Either horn of the dilemma the engineer chooses results in a problem—hence the dilemma. Finally, I will consider the philosophical method of pragmatic functions of concepts (Queloz, 2021; Fisher, 2015) that could resolve this problem because it would render the engineer's motivation unimportant. However, I will show that accepting this view will make

conceptual engineering trivially unimportant precisely because the engineer's motivation becomes unimportant.

4.2 Evaluation of parameters

The first point that will be helpful in the discussion is that the norms to which a concept can be evaluated are likely to be indexed or parameterised (Burgess & Plunkett, 2013b, p.1104). Norms can be, amongst others, the goods, values, (historical) contexts, people, projects, availability of concepts, and goals by which a concept can be evaluated. That concepts are parameterised means that the concept's evaluation by some norm is on a scale instead of binarily defined. The concept of 'truth' could then, for example, be evaluated on the parameter of 'social justice'. To show why the norms of evaluation are on a scale, suppose a concept would be found to be *partly* socially just. It would only need conceptual amelioration for that part; it does not need amelioration for the part of the concept that would be socially just. Thus, norms to which concepts can be evaluated are parameterised, and in what follows, I will call the different norms by which concepts can be evaluated the concept's 'parameters'.

4.3 Conceptual ethics

The second point is that this chapter is about *conceptual ethics* (Cappelen & Plunkett, 2020). This is because the evaluation dilemma concerns how to (ethically) evaluate concepts based on norms that the engineer values. Conceptual engineering is a broader field, including analysing, evaluating and implementing conceptual change. In contrast, conceptual ethics is specific to amelioration: the (ethical) evaluation of concepts. Conceptual ethics can therefore be said to be a subfield of conceptual engineering.

In what follows, I will contend that the conceptual engineer faces an evaluation dilemma.

4.4 Dimensions and parameters

A concept only has to be ameliorated if the concept is found to be defective relative to a norm that the engineer has determined. Therefore, when ameliorating a concept, it has to be determined that the concept is defective based on some set of parameters. However, on what set of parameters can a concept be defective? The answer that will lead up to the engineer's evaluation dilemma is that a concept can be determined to be defective on many sets of parameters. There are roughly five categories that classify the many sets of parameters by which concepts can be evaluated. These are called the different *dimensions* on which a concept can be evaluated. The five rough categories of dimensions are moral-political, epistemological, practical, cognitive, and metaphysical defects (Cappelen & Plunket, 2020, p.10). In respective order, the categories of parameters hinder or promote, for example, social justice, the acquisition of knowledge, good practical functioning, good cognitive functioning, and correspondence to joints in reality. The five dimensions have different parameters in these categories by which a concept can be determined to be defective.

For example, a concept can be ameliorated in its metaphysical dimension by the good of truth. Or, a concept can be ameliorated in its practical dimension for the good of clarity. Or by the moral-political good of honesty (Burgess & Plunkett, 2013b, p.1104). These are three examples of different dimensions and sets of norms on those dimensions by which a concept can be determined to be defective and thus needs to be ameliorated. Another example is the concept of 'truth' itself. In the literature, the concept is evaluated for its epistemological virtues of coherence and consistency (Scharp, 2021). Haslanger (2000) proposes an amelioration of the concept of 'woman' for the moral-political good of social justice. Another example, as I have discussed, is that a concept can be evaluated based on epistemic goods. In this case, the concept could also be evaluated by the parameters of coherence, justified belief, and facilitation. Possible parameters of the metaphysical dimension can be the goods of the naturalness of a concept or its contribution to accurate representation (Cappelen & Plunkett, 2020).

Thus, within the five categories of dimensions, there are different parameters by which concepts can be evaluated. These can be different norms, goods, values, goals, historical context, availability of concepts, and unethical parameters (Burgess & Plunkett, 2013b, pp. 1105-1106). To elaborate, historical contexts are the contexts to which a concept has been subjected in history. Similarly, it could be the case that the availability of concepts can influence whether it should be determined that a concept ought to be ameliorated. Think about Haslanger's (2000) amelioration proposal of the concept of 'woman'. The conceptual change will play a role in discarding the contemporarily used biological concept of 'woman'. It could be the case that this will affect how the contemporary biological concept of 'man' is used. This is because, after amelioration, one of the concepts would represent a biological being and the other a social relation in society. When ameliorating concepts, such changes might have to be considered.

4.5 The evaluation dilemma

In summary, when proposing conceptual change, the goals and aims of the conceptual engineer matter (Shields, 2021), the social values in society could be the most valuable set of parameters (Haslanger, 2000), and the (historical) pragmatic function of the concept (Fisher, 2015), or the standards of evaluation themselves (Burgess & Plunkett, 2013b, p.1107). In short, concepts can be evaluated on multiple dimensions, and many parameters can be considered on those dimensions. I will argue that the engineer only has two options of horns that result in a dilemma when the engineer is motivated to ameliorate a concept. The first horn the engineer can take is evaluating a concept based on *all-things-considered* (Burgess & Plunkett, 2013b, p.1106). This means that the concept is evaluated based on all of the possible dimensions and sets of parameters that affect the concept—at the same time. The second horn the engineer can adopt is to propose evaluation by one set of parameters. This means that a concept is ameliorated on a particular partial defect, and only this singular set of parameters of the concept is considered.

This is the evaluation dilemma that the conceptual engineer faces. To accept the first horn—considering *all things* about a concept—is *practically impossible* because if a large group of people uses a concept, the concept has a massive cloud of (historical) context that has to be considered. In fact, for starters, every time the concept has been used in history has to be considered. Or all the contexts in which a concept has been used. Such a massive project is not practically feasible because considering all things about a concept would be a life-consuming task for an engineer. Even with the help of contemporary artificial intelligence, such a project would take a lifetime.

This leaves the second option for the conceptual engineer to consider one set of parameters of a concept. However, when the engineer adopts this horn, I will show that ameliorating a concept is likely to bring more unintended changes to the concept that are not considered. Therefore, this horn will likely create more problems than the initial amelioration solves. Thus, either horn results in a problem, which is why the conceptual engineer's project of evaluating concepts results in an evaluation dilemma.

Burgess & Plunkett, who are central participants in the literature, have asked the following question (2013b, p.1106): "Should we countenance 'all things considered' conceptual norms, or does the correct conceptual ethics just consist of various relative mandates – for different goods of interest, say?" Haslanger's ameliorative project of the concept of 'woman' can help to clarify again. Do all possible dimensions, parameters and goals from the engineer have to be considered when proposing an amelioration for the concept of 'woman'? Or can an amelioration start by evaluating the concept of 'woman' on, for example, the singular dimension of the moral and political good of social justice?

4.6 First horn: all-things-considered

As I have shown, there are a lot of dimensions and parameters to consider for conceptual amelioration. The first horn of the dilemma that the engineer could take is considering all parameters when evaluating a concept. This is what Burgess & Plunkett call

all-things-considered (2013b, p.1106). In what follows, I will argue why adopting this horn is practically impossible.

As Burgess and Plunkett show, as in other fields, such as broadly ethical or epistemic inquiry, all things can be considered, but examples of specific agents can also be considered. The problem for the conceptual engineer is that if the engineer is serious in his project, he or she aims to intentionally change the semantic meaning or psychological structure of concepts used in a standing language by *most people*. This is important because this means a large number of people. Most of the people in that language, especially if the concept is basic or is commonly used—like the concepts of 'truth', 'woman', and 'disability', in other words, the words that are being proposed for amelioration—have a history with the old concept, and will use the new concept in the future. Therefore, if the engineer considers all things about a concept, he or she *has* to consider all those people's histories and future use of the new concept. The engineer has to consider what effect this will have on those different people and the impact that will have on society. For example, considering all things means considering the likely initial confusion the conceptual change will bring. In other words, if the engineer is serious about considering all things that a conceptual engineer's amelioration will bring about, he or she has to consider all things about a concept. However, as I have shown, there are a lot of aspects and norms to a concept that could—or should—be considered.

4.7 Practically impossible

Suppose this argument is taken to the extreme, and it is argued that the engineer has to consider *all things* when considering an amelioration. In that case, it is easy to see why this is practically impossible. In chapter I, I discussed that philosophers are still unclear about what metaphysical reality concepts have. Nevertheless, even if philosophers knew the metaphysics of concepts, when a large number of people use those concepts, the concepts will have a large number of interactions with either objective reality, human psychology, the historical semantic use of the concept, or our socially constructed reality. What metaphysical

reality concepts will turn out to have, considering all things about them, is practically impossible.

4.8 Second horn: one set of parameters

I have thought about this problem, and the argument I provide in this section is the first elaboration of the problem that Burgess & Plunkett opened in their discussion about "the plurality of parameters" (2013b, p.1104). There is a lot more to be said about this. The problem that I will consider is that considering few parameters for conceptual amelioration is likely to lead to more unintended problems that cannot be foreseen.

It might be enough to evaluate a concept by considering one, or a small set, of all the possible parameters. This would leave the engineer with one, or a few sets of parameters to evaluate a concept. For example, if we take the example of Haslanger's amelioration proposal of 'woman' again, the concept of woman could be evaluated on the moral-political good of social justice. If this evaluation of the concept of the singular dimension of the good of *social justice* resulted in the realisation that the concept of 'woman' is defective in this dimension, this would lead to the recommendation to conceptually engineer the concept of 'woman'. However, the amelioration of the concept could impact (all) other dimensions that would not have been considered. We likely put the cart before the horse and create more problems in other dimensions that we have not foreseen by focusing only on this singular dimension.

The engineer could also consider what the engineer determines to be the most essential set of parameters. However, this suffers from the same evaluation problem. On what basis can a conceptual engineer claim to have found the most essential set of parameters of a concept? This will suffer from conceptual conflict.

To see why this is the case, suppose that Haslanger's amelioration proposal of 'woman' is successful. This change will affect the historical function of the concept of 'woman'. The new concept, for example, could not coherently refer to historical texts in

which the old concept of 'woman' was used because the concept would not represent that historical concept now that it has been changed.

The problem is that considering few parameters is highly likely to lead to unconsidered problems that cannot be foreseen. For example, after amelioration, biologists and archaeologists will lack a concept to refer to the biological thing that the contemporary concept of women refers to. This will confuse the field, or the fields will be forced to create a new concept to represent the biological thing represented by the old concept. This is an obvious example. But similarly, Haslanger's analysis does not, for example, address the value of *harm done to women who do not accept the idea of structural oppression*. Suppose that because of Haslanger's amelioration proposal, some women will believe they are structurally oppressed even though they are not. Will the conceptual engineer proposing this amelioration consider this possibility?

Haslanger's amelioration proposal might be desirable if social justice is the *only* metric to be considered. However, her amelioration could also harm women on parameters other than social justice. These are some examples of the *unintended consequences* that amelioration of one set of parameters of a concept is likely to have. Therefore, even if the engineer's intention is honest and the good of social justice is the primary motivator, amelioration on one set of parameters is likely to be problematic. This is because it is nearly impossible to know the unintended consequences that conceptual amelioration is likely to cause.

I will provide another example. Consider the possible unintended consequences of Cantalamessa's amelioration proposal for the term 'disability' (2021). The motivation for the amelioration that she proposes is "for social reform" and "to improve people's lives" (2021, p.49). All other parameters to which the concept 'disability' can—and maybe *ought* to—be evaluated are trumped by this motivation. More strikingly, Cantalamessa contends that it is justified to make (scientifically) incomplete and false claims if social reform is attained and people's lives are improved. In other words, she claims that for conceptual amelioration, the goods of (scientific) coherency, truth and honesty *ought* to be trumped if an increase in the

good of social justice can be attained. In this specific example, it is clear to see the problem: this amelioration proposal disregards almost all parameters of the concept in favour of *one* most valued set of parameters. Other possible dimensions of the concept are not considered.

Suppose that society would indeed adopt Cantalamessa's proposal, and most speakers in that language, when speaking about disabilities, would communicate (scientific) untruths and defend incoherent claims about the concept with the motivation to improve social justice. It might be the case that social justice for disabled people would indeed improve. However, if this change were to be implemented, communication and trust between people in that society would also necessarily decline. This is because people would tell each other untruths and deliberate inconsistencies. Moreover, the untruths and incoherencies would be about disabled people, which would necessarily and inadvertently impact those disabled people negatively.

Thus, Cantalamessa's argument might be correct and desirable if only one aspect of a concept was considered for amelioration. However, I have shown why this attitude toward amelioration is problematic. Cantalamessa ought to consider more things about the concept when the aim is to ameliorate it for the total benefit of society. If she does not, her proposal will result in more social harm than social good, all-things-considered, even though her initial motivation was to improve the concept of disability to improve people's lives.

To make matters even more complicated, the problem of conceptual conflict is also relevant to the evaluation dilemma. Even though Gibbons (2022) discussed this problem specifically regarding the conceptual engineer's implementation problem, this problem also holds for conceptual evaluation. Suppose considering-all-things is not practically possible, and considering a smaller set of parameters to a concept is undesirable but necessary. In that case, engineers will conflict on which parameters should be considered before other parameters. What if two speakers construct opposing arguments, one favouring the parameter of social justice, the other of scientific coherency? Which set of parameters ought to be valued more by the conceptual engineer?

In short, my first thoughts on this dilemma leave the conceptual engineer with a dilemma. Considering-all-things about a concept is practically impossible, and considering a few parameters of a concept is highly likely to result in more unintended problems that cannot be foreseen than the initial amelioration will solve.

4.9 Pragmatic function of concepts

A possible solution for this problem could be to only consider the pragmatic functions of concepts as a justification for amelioration (see Queloz 2021 and Fisher 2015) because it would render the subjective motivation of the conceptual engineer unimportant. First, I will consider this view. However, secondly, I will show that the truth of this externalist view would make most of the current conceptual amelioration proposals redundant.

To consider this view, Queloz defends that the conceptual engineer should not necessarily consider a concept in its most abstract form, away from day-to-day human practice. Instead, amelioration projects should start with "men and their conversation" (2021, p.1). This view assumes that concepts satisfy the practical needs of humanity. Empirical research could then result in an understanding of what pragmatic function a concept has fulfilled for humanity and, therefore, what function the concept has had for human survival and flourishing. This research would provide the engineer with a justification for amelioration because if a concept was pragmatically valuable at some point in history, it could be determined that it would not be pragmatically functional for contemporary use by individuals or society.

For example, Fisher's (2015) philosophical tool of *pragmatic, experimental philosophy,* based on teleo-semantic discovery, could be a basis for amelioration. Teleo-semantics is the field that seeks to explain meaning in terms of the function it has for our species. Conceptual representation, in this view, is based on the biological function that the representation has for humanity. A concept can then be said to be *good* if pragmatically living by that concept yields beneficial outcomes. Therefore a justification for amelioration can be based on pragmatic, experimental philosophy. Suppose a concept is experimentally

changed based on one set of its parameters, and the lives of a study group improve relative to a control group. In that case, the engineer has scientific evidence supporting a justification for his amelioration proposal, and the evaluation dilemma is solved.

However, the problem with this view for conceptual engineering is that it assumes that concepts have a function determined by teleo-semantics. Concepts can, therefore, not be ameliorated based on an engineer's subjective values or goals. Instead, they can only be ameliorated based on the outcome of teleo-semantic research. If such a view were accurate, it means that central participants in the conceptual engineering literature make invalid amelioration proposals. For example, Haslanger's conclusion would fail: "rather than worrying, what is gender, really? or, what is race, really? I think we should begin by asking what, if anything, *we want them to be*" (2000, p.52). Based on a functional view of concepts, the engineer could not decide *what they want gender to be* because teleo-semantics would be the only justifiable way to determine what a concept should be.

This results in the fundamental question that the new field of conceptual engineering has to settle: *what ought conceptual engineering be about*? Should conceptual engineering be about the intentional change of concepts based on a subjectively determined goal from the engineer? Or can conceptual change only be justified based on empirical research based on an external reality? As I have argued, the former option suffers both an implementation problem and an evaluation dilemma. Suppose the values and motivations an engineer believes to be most important about a concept trump other aspects of the concept, all-things-considered. In that case, the change to the concept is most likely not an amelioration, all-things-considered, because the change will have unintended consequences for the concept. The latter option, however, reduces the conceptual engineering project to the empirical analysis of the externally determined functions of concepts. This reduces the conceptual engineering project to something trivial and unimportant.

Finally, in line with Cappelen's answer to the implementation problem (2018, p.72), it could be argued that the engineer finds himself in an inscrutable evaluation dilemma but that the engineer simply has to keep trying. What if the engineer would evaluate a concept

based on the engineer's values and motivations and use any means necessary to implement this change in society? The change would have unintended consequences, but ameliorating one set of parameters could be desirable to the degree that they would trump the unintended consequences of the amelioration. This is unlikely but possible and could perhaps be tested empirically.

4.10 Conclusion

This chapter is the first attempt to explore "the plurality of parameters" discussed by Burgess & Plunkett (2013b, p.1104). I have considered that because concepts can be evaluated based on a large set of parameters, the engineer finds him or herself in a dilemma at the evaluation stage of conceptual amelioration. If a concept is baptised, the concept will have a history and a standing meaning relative to most people. Such a baptism creates the dimensions and parameters that should be considered when ameliorating the concept. There are two modes of evaluation possible for the engineer. Either the engineer adopts the horn of considering all things about a concept, or the engineer considers the horn of evaluating a few parameters through the engineer's subjective and biassed motivation. As I have argued, the former horn is practically impossible because of the enormous amount of goods, goals, motivations and values that the engineer ought to consider. The latter horn is problematic precisely because it is selective, subjective and biassed. The conceptual amelioration, even if the engineer's motivation is transparent and pristine, a change by this horn is likely to result in more unintended problems than solutions.

Finally, I have discussed a possible solution to this dilemma by considering whether a pragmatic functional approach to concept amelioration could solve it. However, if the conceptual engineer is serious in his project, the engineer ought to at least change the semantic meaning or psychological structure of concepts used in a standing language by most people. Since a pragmatic functional approach to conceptual amelioration assumes that concepts can only be ameliorated relative to an external reality in which concepts carve

nature at their joints, this approach will render the project of conceptual engineering trivial and meaningless.

5 Chapter IV - The Problem of Motivation

5.1 Introduction

I am going to make two arguments in this chapter. First, the conceptual engineer faces a motivation problem because, as I shall argue, the engineer faces a wrong-kind-of-reason problem (Simion, 2018). Second, I will argue that the new field of conceptual engineering ought to justify the method prior to engaging in the method. This is because I will argue that if the field becomes successful, a small group of conceptual engineers will obtain a method that yields conceptual power over a large group of people. As I have argued in chapter I of this dissertation, if the conceptual engineer is serious in his endeavour, the engineer aims to change (at least) the semantic meaning or psychological structure of concepts *used by most people*. If this is the case, then it is necessarily the case that if the engineer succeeds in this endeavour, a small group of engineers will gain the power to influence how most people will conceptualise the world. This means that a small group of conceptual engineers will be highly powerful and able to influence a large number of people by controlling how most people conceptualise the world. I contend that the literature on conceptual engineering should justify the field by addressing the likelihood that a small group of people will gain conceptual power and control over a large group of people through the success of the method.

5.2 The motivation problem - Part I

Different motivations, goals, and aims that the conceptual engineer has, give rise to different norms valued for a proposal of conceptual amelioration (see Shields, 2021; Cappelen, 2018; Cappelen & Plunkett, 2020; and Burgess & Plunkett, 2013). For example, whether the engineer's goal is practical, theoretical, political, semantic, personal, for the collective, or a mixture thereof, affects how the engineer determines that a concept ought to be evaluated and possibly ameliorated.

It could be the case that some motivations or goods that the conceptual engineer holds ought to be considered before other motivations. To see why this is the case, suppose

that an engineer's primary motivation to ameliorate a concept is to make money or win an election. This amelioration, at the same time, increases social exploitation. In this case, the former motivations of the engineer do not seem directly justified. At least, in contemporary society, individuals who choose to make money through a method that increases social exploitation are criticised. To the degree that this criticism is justified, the engineer who proposes amelioration based on the motivation to make money over the motivation to increase the good of social justice is injudicious. In this case, the motivation to increase social justice ought to trump the engineer's personal motivation to make money. In other words, a case can be made that, in many cases, a conceptual engineer's motivations should not be the primary justification for amelioration. The engineer's goals and motivations could still matter, but how much and relative to what is an open question. In other words, (universal) goods could be more "explanatorily fundamental" than an engineer's motivations or goals when it comes to settling the normative questions of evaluation in conceptual ethics (Burgess & Plunkett, 2013b, p.1105).

Central participants in the literature note that "many claims in conceptual engineering appeal to the *aims* an agent has in using a given concept" (Cappelen & Plunkett, 2020, p.10). This means that many of the recent amelioration proposals are based on the subjective aims and motivations of the conceptual engineer. People are naturally tempted to evaluate with specific goods and motivations in mind. All human beings have their fields of interests, goals, aims and motivations. Moreover, when evaluating concepts, engineers are likely to primarily bring considerations from their specialisations and respective areas (Burgess & Plunkett, 2013b, p.1106). For example, as seen in their amelioration projects, Cantalamessa (2021) and Haslanger (2000; 2021) are motivated by social justice. Scharp (2021) is motivated by the desire to create a coherent concept from a concept accused of incoherence.

In contrast, Queloz (2021) and Fisher (2015) are motivated by the idea that concepts carve nature at its joints and that concepts can carve nature closer at the joins, improving concepts by improving their niche in nature. All of these motivations are necessarily biassed.

They blind evaluators to other motivations for which a concept *ought* to be evaluated. This is the first motivation problem that the field of conceptual engineering faces. How *should* the conceptual engineer's biassed goals, aims and motivations be considered and accounted for?

A solution could be the existence of objective normative facts. If these existed, the engineer could use these to objectively determine how a concept ought to be ameliorated. The engineer's goals, aims, and motivations would be out of play. However, the new literature on conceptual engineering essentially assumes that no objective normative facts exist. The truth of objective normative facts that the conceptual engineer should follow will make the argument of central participants in the literature redundant. Think about Haslanger and Cantalamessa. In what follows, I will provide an example based on an objectivist theory of concepts, and as we shall see, this will reduce the conceptual engineering project to something trivial. It could not be what the central participants in the literature want it to be.

5.3 Wrong-kind-of-reasons

To give an example of the problem, what if a concept is semantically perfectly functional and carves nature epistemically at the joints but is determined by a conceptual engineer to have harmful moral consequences? Should the concept be ameliorated to become more moral even if the conceptual change will make the concept semantically faulty? This dilemma is explored by what Simion calls the wrong-kind-of-reasons problem (2018). Intuition says that such an amelioration is not justified.

To begin with, suppose a conceptual engineer would enforce such a conceptual change; the amelioration would necessarily have unintended consequences—in this case, to both the semantic and epistemic dimensions of the concept. The engineer who unidimensionally determines that a concept has harmful moral consequences is necessarily biassed and blind to other dimensions at which the concept serves a function. The question

becomes, based on what motivations can the conceptual engineer be justified to ameliorate a concept? Or, what are good-kind-of-reasons to ameliorate a concept?

As I have already argued, not all motivations are justified as primary motivations. Therefore, some speakers have tried to set limits for the conceptual engineer so that his project does not become trivial and unimportant but is justifiable simultaneously. I will consider one such view to see whether it could solve the motivation problem. In her paper, Simion (2018) constructs a *limiting recipe* that limits the conceptual engineer in his or her project. Simion's goal is to set boundaries for the engineer within which amelioration is justified. This is what she calls the Epistemic Limiting Procedure.

> *The Epistemic Limiting Procedure (ELP): A representational device should be ameliorated if (1) There is all-things-considered reason to do so and (2) The amelioration does not translate into epistemic loss (Simion, 2018, p.923).*

There are two parts to this procedure. The first is that Simion's limiting procedure assumes it is *possible* to consider all things. As I have argued in chapter III, if *considering-all-things* is accepted literally, the limiting recipe is practically impossible to execute. In some sense, considering all things about a concept means considering an infinite amount of consequences that cannot all be overseen by the engineer. Therefore, if my argument in chapter III is correct, this limiting procedure is practically impossible.

The second part of the ELP is based on Simion's claim that concepts are best understood as *representational devices* and have a primary epistemic function: representing the world. In other words, Simion defends that a concept's essential dimension and set of parameters is its epistemic function. This means the ELP assumes that if a concept is ameliorated, its epistemic function should be considered first and foremost. Political or moral reasons can only justifiably be the cause for amelioration if no epistemic loss occurs.

If Simion's argument is correct, and the epistemic function of a concept should be considered first and foremost, this would limit the conceptual engineer to a large degree. Suppose I would fraise Simion's argument differently. In that case, her point is that the ability to communicate with concepts that epistemically represent the world trumps all

other considerations by which a concept can be evaluated. I contend that this argument stands because the ability to communicate coherently through concepts representing the world should be more crucial than socially just communication that is epistemically incoherent with reality. This is because if a concept is determined to be socially just, but also epistemically incoherent to reality, the speaker will be blinded to reality and communication with other speakers will be based on untruths. Some people disagree with this argument. To show why this is the case, Simion writes:

> *"Of course, if the life of millions is at stake, revising a perfectly adequate representational device will be all-things-considered permissible. However, the resulting concept will not be a better concept, in virtue of being less likely to fulfil its central function; therefore, no conceptual improvement will have taken place"* (Simion, 2018, p.924).

I agree with Simion's statement that such an amelioration does not result in a better concept. If concepts do indeed represent nature's joints, this is necessarily true. However, I argue against the justification that if millions of lives are at stake, revising a perfectly adequate representational device is straightforwardly all-things-considered permissible. Once again, this is because the engineer is not omnipotent. Maybe it would be correct that a conceptual change (I will not call it an amelioration because I do not believe such a change to be an improvement) in this case would save millions of lives. However, because human beings are not omnipotent, it could be the case that this conceptual change saves a million people today but that it will be the cause of two million deaths in the future.

This problem is reflected in the famous discussion of normative ethics between deontologists (Kant) and consequentialists (Bentham). Should a moral decision be made based on a universal moral law? Or based on the consequences of pain and happiness? In line with my argument from chapter two that considering all things is practically impossible, I would argue that it is impossible to calculate the actual consequences of any moral decision. How far in the future will the calculation go? No human being is omnipotent, and, in the final analysis, no person can calculate the actual consequences of any action.

Suppose a society that holds perfectly adequate representational devices on which the lives of a million people are at stake results in the deaths of a million people. What would that society be if it could not coherently rely on its conceptual representational devices to represent the world? All other things being the same, such a society would be worse. Who is to say that this initial conceptual change that saves a million people today does not result in a worse situation, all things considered, tomorrow? As my argument against the consequentialists shows, it is impossible to calculate the consequences.

More attempts to construct a limiting recipe that would limit the conceptual engineer so that amelioration would be justified are needed. Simion's argument for her ELP is itself not satisfactory. First, the former premise of ELP is based on the engineer's ability to consider all things. As I have shown, this is practically impossible. And second, recent amelioration proposals in the conceptual engineering literature are based on this view, making those proposals redundant if ELP is correct. However, even though Simion's attempt is not satisfactory, it could be the case that other limiting recipes could limit the conceptual engineer in such a way that the project remains both practically possible and does not become a trivial or unimportant project.

5.4 The motivation problem - Part II

As I have argued, if conceptual engineering is understood as a serious endeavour, it aims to (at least) be a method that can change the semantic meaning or psychological structure of concepts that a large group of people uses. This means, for example, that if some concepts from the standing language of English were successfully engineered, the change would have real-world consequences for a large part of the English-speaking population.

For the sake of argument, let us again assume that the engineer could overcome the implementation problem. The engineer would have a practical philosophical method that gives an individual, or a small group of people, the power to implement semantic meaning or psychological structures of concepts in a standing language that a large group of people speaks. If this is the case, I will argue that there are reasons to be careful with the power to

be able to ameliorate concepts because I will show that it would not be an irreproachable tool. As we have seen, changes in concepts have an impact on real-world people. Suppose humanity would create such a method—and therefore the power—to implement changes to concepts. In that case, this power can also be used not to *ameliorate* concepts but to gain conceptual power over people's lives—even use that power with the motivation to do evil.

To the degree that responsibility comes with power, it is not straightforwardly desirable to create a tool that enables a small group of people to fundamentally change how large groups of people conceptualise the world. To be clear, the issue I am concerned about is not power itself. It is the conceptual power that a small group of conceptual engineers will necessarily have over a large group of people. With such a tool comes the responsibility for all potential conceptual engineers to use that tool *responsibly*. If a successful case of conceptual engineering would indeed be a conceptual change that a large group of people would use, such changes would influence a large group of people and their lives. There is no straightforward reason to suppose that participants in the conceptual engineering literature are occupied with something other than the desire to create better concepts to benefit humanity and the planet. However, if conceptual engineering would be successful, individual engineers would also create a method that has the potential power to work against the well-being of humanity and the planet at a large scale.

I will provide an example to show this more clearly. It could be said that it would have been better if the atomic bomb was never created. The only two bombs that have ever been used on people, Hiroshima and Nagasaki, were thrown with the (probably ethical) intent to end the second world war. Still, a strong case can be made (and is still being made) that it would be better if no atomic bomb had ever existed. This is because with great power comes great responsibility, especially if that power is in the hands of a select few. History shows that humanity has not consistently been mature enough to bear that responsibility.

Think about the contemporary conflict with Russia. Would it not have been better if the leaders of Russia and the West did not have access to an atomic bomb? It is impossible to look into the future precisely because it is impossible to consider all things. However, this

does not mean that with the current state of our humanity, it would be wise to create another large-scale method that can influence millions of people's lives by a select few. As I have argued, if the field of conceptual engineering is serious, the engineer would necessarily gain such power.

Suppose that in the future, a method of conceptual engineering would be accessible to the conceptual engineer and humanity would find itself in the same situation as at the end of world war two. It would not be two bombs that can be thrown onto two cities of a country that refuses to surrender and has been the cause of terrible horror in the world. Instead, in the future, it could be a group of people who have been using a concept in a way that has resulted in terrible horror in the world. Conceptually changing the concept would stop this horror and doom two cities' worth of people. Would this conceptual change be justifiable? Moreover, should humanity put itself in the situation in which such a decision must be made?

I contend that the field of conceptual engineering ought to justify its method that includes taking stock of the responsibility that the conceptual engineer will bear if the engineer could succeed in a conceptual amelioration project. This means that I contend that the field of conceptual engineering ought to consider its necessarily biassed motivations but also ought to account for the power and responsibility that the successful method of conceptual engineering will release in the world. What small group of people will use the power to influence how large groups conceptualise the world? Can humanity bear that responsibility? Even if the engineer's motivation is pure, can the necessary unintended consequences of conceptual change be overseen? These questions are not satisfactorily addressed in the literature.

One of the questions is whether the conceptual engineer ought to be invested in getting the best or correct view of a concept (Shields, 2021). It could be the case, as Shields points out, that a conceptual engineer is not invested in a motivation to ameliorate a concept that will result in the best or correct view of that concept. In the contemporary literature on conceptual engineering, there is an assumption that speakers are invested in

getting at the best or correct view of a concept. This is the inquiry assumption: "the assumption that speakers are motivated by and committed to the goal of arriving at the correct or the best view of the relevant concept" (Shields, 2021, p.15047). This contrasts with the conceptual engineer, who is not invested in getting the best or correct view of a concept. Shields calls this a conceptual dominator: a speaker who uses their position of power to enforce an amelioration that is not based on the engineer's motivation to arrive at the best or correct view of a concept.

Any conceptual engineer can be a conceptual dominator, and all amelioration proposals can have (partial) motivations other than arriving at the best or correct view of a concept. For example, as I have already shown, Cantalamessa contends that claiming something that is semantically or representationally "trivial, overstated, incomplete, or simply false" is justified when the audience ought to change their concepts for the causes of "positive social reform and improving people's lives" (2021, p.49). This means that Cantalamessa is motivated by her claim that concepts are primarily instrumental to social justice, not epistemic coherence. In Cantalamessa's view, intentionally claiming epistemic incoherencies is encouraged if that claim would increase societal justice. This means that Cantalamessa is not invested in getting the concept's best or correct view. Instead, she is invested in her subjective motivation to use any conceptual means necessary to increase the unidimensional goal of social justice. This is a social constructionist perspective; no objective normative facts exist. Therefore Cantalamessa is a conceptual dominator if she uses her position of power to implement her amelioration through her method of conceptual activism that uses untruths. If Cantalamessa, or any conceptual engineer with a comparable unidimensional motivation, would gain the conceptual power to change concepts that a large group of people uses, her amelioration is unjustified.

5.5 Conclusion

I have made two arguments in this chapter. First, the engineer faces a motivation problem because the motivation for amelioration suffers wrong-kind-of-reasons. Not all motivations

based on which a conceptual engineer wants to ameliorate a concept are good motivations. Simion attempts to limit the conceptual engineer with a limiting recipe so that amelioration will be performed for proper-kinds-of-reasons. I have refuted her proposal, but future attempts to set limiting recipes for the conceptual engineer could solve the issue. In this way, it could be the case that a limiting recipe would be found that would provide the engineer with proper-kinds-of-reasons and does not trivialise the field as something unimportant.

Second, suppose the engineer is serious about the project of conceptual change, and the engineer aims to change the semantic meaning or psychological structure of concepts that most people will use. I have argued that the conceptual engineer owes a justification for creating a tool that enables a small group of people to have conceptual power over a large group of people. I have provided the metaphor of the atom bomb: ought a small group of people ever be in the position to obtain the power to influence the lives of millions of people? To the degree that the answer is no to this question, the method of conceptual engineering ought not to be obtained. If this is not satisfactory, the conceptual engineer at least owes a justification for why the engineer ought to obtain a method that enables the conceptual engineer the power to influence the lives of a large number of people, even if the intention for creating the method is purely altruistic.

6 Conclusion

This dissertation argues that if the conceptual engineer is serious about the conceptual engineering project, the engineer aims to (at least) change the semantic meaning or the psychological structure of concepts used by most people in a standing language. Based on this definition of the serious conceptual engineer, in chapter II, I have contended that the engineer faces an implementation problem. This is the problem that even if a concept is considered defective, no known method has successfully been tested to implement that change. The most likely method that could work is collective long-range control over conceptual change. This method already works in other large-scale normative projects, such as climate change related issues. However, conceptual amelioration proposals would suffer from conceptual conflict. This is a disagreement between speakers who have evaluated concepts differently and, therefore, disagree about the correct view of a concept. I have discussed one possible solution to this problem: conceptual activism. However, this method is not desirable because the method is based on deceit and untruths.

Chapter III considers the evaluation dilemma. These are the first thoughts on this problem, and more has to be said about it. In short, this is the dilemma that the conceptual engineer faces if the engineer could implement conceptual change but has to consider based on what parameters a concept ought to be evaluated. A concept has to be determined to be defective relative to some normative standard. However, what normative standard(s) should be used? I defend that the first horn of the dilemma that the engineer can take is considering all things about a concept. I argue that this is practically impossible because, considering all things about a concept would take more than a lifetime.

On the other hand, accepting the second horn of the dilemma will only consider one or a small set of parameters to evaluate a concept. However, my argument is that this will likely lead to more unintended problems than the initial amelioration would solve. Therefore, even if the engineer has a method to implement conceptual change, the engineer finds him or herself in a dilemma on what basis to evaluate any concept for conceptual

amelioration. Again, these are the first thoughts on this problem, and more has to be said about it.

Chapter IV explores the evaluation problem more deeply by explicitly considering the conceptual engineer's motivations, aims, and goals that play a role in evaluating concepts. Here I contend that, by accepting the serious definition of the conceptual engineer, the engineer faces two motivation problems. The first motivation problem is the problem of accounting for the biassed motivations of the conceptual engineer. I argue that the subjective motivations of the engineer are problematic and that such motivations are likely to lead to the wrong-kind-of-reason-problem. This is the problem that the engineer could ameliorate a concept for the wrong reasons.

The second motivation problem is the problem of great power in the hands of a small group of engineers. Based on the serious definition of conceptual engineering, the engineer intentionally aims to gain the power to change the concepts that most people will use in a standing language. Therefore, if the conceptual engineer is successful, one or a small group of engineers has conceptual power over large groups of people. I argue that the conceptual engineer owes an explanation for the motivation to gain this power. Moreover, even if a particular conceptual engineer has an honest motivation, the power could be used by other people who do not have an honest motivation. I argue that the engineer who desires to create a method that could improve concepts owes a justification for why it would be better if such a power were created.

In short, based on the serious definition of the conceptual engineer, the new field of conceptual engineering has both practical and normative challenges to overcome before the method can be what its primary participants want it to be.

7 Bibliography

Bihagen, Erik & Katz-Gerro, Tally. (2000). Culture consumption in Sweden: The stability of gender differences. Poetics. 27. 327-349. 10.1016/S0304-422X(00)00004-8.

Boghossian, P. A. (2001). What is social construction?. TLS.

Burgess, A., & Plunkett, D. (2013a). Conceptual Ethics I: Conceptual Ethics I. Philosophy Compass, 8(12), 1091–1101. https://doi.org/10.1111/phc3.12086

Burgess, A., & Plunkett, D. (2013b). Conceptual Ethics II: Conceptual Ethics II. Philosophy Compass, 8(12), 1102–1110. https://doi.org/10.1111/phc3.12085

Cantalamessa, E. A. (2021). Disability studies, conceptual engineering, and conceptual activism. Inquiry, 64(1–2), 46–75. https://doi.org/10.1080/0020174X.2019.1658630

Cappelen, H. (2018). Fixing language: An essay on conceptual engineering (First edition). Oxford University Press.

Costa, Paul & Terracciano, Antonio & McCrae, Robert. (2001). Gender Differences in Personality Traits Across Cultures: Robust and Surprising Findings. Journal of personality and social psychology. 81. 322-31. 10.1037//0022-3514.81.2.322.

Deutsch, M. (2020). Speaker's reference, stipulation, and a dilemma for conceptual engineers. Philosophical Studies, 177(12), 3935–3957. https://doi.org/10.1007/s11098-020-01416-z

Fisher, J. C. (2015). Pragmatic experimental philosophy. Philosophical Psychology, 28(3), 412–433. https://doi.org/10.1080/09515089.2013.870546

Fridman, L. (2021). Yeonmi Park: North Korea | Lex Fridman Podcast #196. [online] www.youtube.com. Available at: https://www.youtube.com/watch?v=usDqSEKDVsA&t=474s [Accessed 24 Oct. 2022].

Gibbons, A. F. (2022). Rational conceptual conflict and the implementation problem. Inquiry, 1–27. https://doi.org/10.1080/0020174X.2022.2037455

Haslanger, S. (2000). Gender and Race: (What) Are They? (What) Do We Want Them To Be? Noûs, 34(1), 31–55. https://doi.org/10.1111/0029-4624.00201

Haslanger, S. A. (2012). Resisting reality: Social construction and social critique. Oxford University Press.

Isaac, M. G. (2020). How to conceptually engineer conceptual engineering?. Inquiry, 1-24.

Jackman, H. (2020a). Temporal externalism, conceptual continuity, meaning, and use. Inquiry, 63(9-10), 959-973.

Jackman, H. (2020b). Construction and continuity: conceptual engineering without conceptual change. Inquiry, 63(9-10), 909-918.

Jorem, S. (2021). Conceptual engineering and the implementation problem. Inquiry, 64(1–2), 186–211. https://doi.org/10.1080/0020174X.2020.1809514

Knoll, V. (2020). Verbal disputes and topic continuity. Inquiry, 1-22.

Koch, S. (2021a). The externalist challenge to conceptual engineering. Synthese, 198(1), 327–348. https://doi.org/10.1007/s11229-018-02007-6

Koch, S. (2021b). Engineering what? On concepts in conceptual engineering. Synthese, 199(1–2), 1955–1975. https://doi.org/10.1007/s11229-020-02868-w

Kripke, S. A. (1980). Naming and necessity. Oxford: Basil Blackwell Ltd.

Löhr, G. (2021). Commitment engineering: Conceptual engineering without representations. Synthese, 199(5–6), 13035–13052. https://doi.org/10.1007/s11229-021-03365-4

McPherson, T., & Plunkett, D. (2021). Topic Continuity in Conceptual Engineering and Beyond. Inquiry, 1-27.

Nietzsche, F. W., (1901/1968). The Will to Power. Trans. W. Kaufmann. New York: Random House

Scharp, K. (2021). Conceptual engineering for truth: Aletheic properties and new aletheic concepts. Synthese, 198(S2), 647–688. https://doi.org/10.1007/s11229-019-02491-4

Schmitt, David & Realo, Anu & Voracek, Martin & Allik, Jüri. (2009). "Why Can't a Man Be More Like a Woman? Sex Differences in Big Five Personality Traits Across 55 Cultures," (vol 94, pg 168, 2008). Journal of personality and social psychology. 96. 118. 10.1037/a0014651.

Schmitt, D. P., Long, A. E., McPhearson, A., O'Brien, K., Remmert, B., & Shah, S. H. (2017). Personality and gender differences in global perspective. International Journal of Psychology, 52, 45-56.

Shields, M. (2021). Conceptual domination. Synthese, 199(5–6), 15043–15067. https://doi.org/10.1007/s11229-021-03454-4

Simion, M. (2018). The 'should' in conceptual engineering. Inquiry, 61(8), 914–928. https://doi.org/10.1080/0020174X.2017.1392894

Queloz, M. (2021). Why We Came to Think as We Do. In M. Queloz, The Practical Origins of Ideas (pp. 1–19). Oxford University Press. https://doi.org/10.1093/oso/9780198868705.003.0001