

University of Groningen

New Problems for Moral Anti-Realism

Kassenberg, Erik Pieter

DOI:
[10.33612/diss.175044982](https://doi.org/10.33612/diss.175044982)

IMPORTANT NOTE: You are advised to consult the publisher's version (publisher's PDF) if you wish to cite from it. Please check the document version below.

Document Version
Publisher's PDF, also known as Version of record

Publication date:
2021

[Link to publication in University of Groningen/UMCG research database](#)

Citation for published version (APA):

Kassenberg, E. P. (2021). *New Problems for Moral Anti-Realism*. [Thesis fully internal (DIV), University of Groningen]. University of Groningen. <https://doi.org/10.33612/diss.175044982>

Copyright

Other than for strictly personal use, it is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license (like Creative Commons).

The publication may also be distributed here under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license. More information can be found on the University of Groningen website: <https://www.rug.nl/library/open-access/self-archiving-pure/taverne-amendment>.

Take-down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Downloaded from the University of Groningen/UMCG research database (Pure): <http://www.rug.nl/research/portal>. For technical reasons the number of authors shown on this cover page is limited to 10 maximum.



university of
 groningen

New Problems for Moral Anti-Realism

PhD Thesis

to obtain the degree of PhD at the
University of Groningen
on the authority of the
Rector Magnificus Prof. C. Wijmenga
and in accordance with
the decision by the College of Deans.

This thesis will be defended in public on
Thursday 30 September 2021 at 16:15 hours

by

Erik Pieter Kassenberg

born on 3 September 1992
in Roden

Supervisor

Prof. B. Streumer

Co-supervisor

Dr. H.W.A. Evers

Assessment committee

Prof. C. Cowie

Prof. J. Olson

Prof. T. Cuneo

Table of Contents

ACKNOWLEDGEMENTS	III
0. INTRODUCTION	1
0.1 CHAPTER SUMMARIES.....	5
1. DEBUNKING THE ARGUMENT FROM QUEERNESS	9
1.1. INTRODUCTION	9
1.2. THE ARGUMENT FROM QUEERNESS	9
1.3. DEBUNKING QUEERNESS JUDGMENTS.....	13
1.4. OBJECTION: BAIT-AND-SWITCH	17
1.5. OBJECTION: ACCESSIBILITY.....	18
1.6. OBJECTION: DOXASTIC MOVE.....	18
1.6.1. <i>Difference in Epistemic Standards</i>	21
1.6.2. <i>Difference in Accessibility</i>	23
1.7. IMPACT ON OUR MORAL PRACTICE.....	26
1.8. CONCLUSION	27
2. GOING NEGATIVE ON THE MORAL FIXED POINTS	29
2.1. INTRODUCTION	29
2.2. CUNEO AND SHAFER-LANDAU’S ACCOUNT	29
2.3. CRITICISM OF CUNEO AND SHAFER-LANDAU’S ACCOUNT.....	30
2.4. POSITIVE AND NEGATIVE MORAL FIXED POINTS.....	32
2.5. SHIFTING FOCUS	35
2.6. DO THE NEGATIVE MORAL FIXED POINTS SHOW ENOUGH?.....	36
2.7. RETAINING BENEFITS.....	37
2.7.1 <i>Moral Disagreement</i>	37
2.7.2 <i>Remarkable Coincidence</i>	38
2.7.3 <i>The Humean challenge</i>	40
2.8. CONCLUSION	41
3. EXPRESSIVISM AND MORAL IRRELEVANCE	42
3.1. INTRODUCTION	42
3.2. IRRELEVANCE CLAIMS.....	42
3.3. CONCEPTUAL IRRELEVANCE	46
3.4. METAETHICS: REALISM AND EXPRESSIVISM.....	49
3.4.1. <i>Realist Solution</i>	49
3.4.2. <i>First Expressivist Answer: Attitudinal Interpretation</i>	53
3.4.3. <i>Second Expressivist Answer: Supervenience</i>	53
3.4.4. <i>Third Expressivist Answer: Evolutionary Explanation</i>	54
3.4.5. <i>Fourth Expressivist Answer: Illusory Conceptual Truth</i>	58
3.4.6. <i>Fifth Expressivist Answer: Moral Attitudes</i>	59
3.5. CONCLUSION	61
4. RELATIVISM AND CONCEPTUAL MORAL TRUTH	62
4.1. INTRODUCTION	62
4.2. KINDS OF RELATIVISM.....	63
4.2.1 <i>Gilbert Harman: Moral Relativism Defended</i>	63
4.2.2 <i>David B. Wong: Moral Relativity</i>	64
4.2.3 <i>Sharon Street: Constructivism about Reasons</i>	65
4.2.4 <i>Stephen Finlay: Oughts and Ends</i>	65
4.2.5 <i>John MacFarlane: Assessment Sensitivity</i>	67
4.3. AXES OF VARIATION	67
4.3.1 R_{who} : <i>Agent, Speaker or Assessor</i>	68
4.3.2 R_{what} : <i>Standards, Normative Judgments, Ends</i>	71
4.3.3 <i>Content or Truth</i>	71

4.3.4 Recap	74
4.4. CONCEPTUAL MORAL TRUTH	74
4.5. THE PROBLEM FOR RELATIVISM	75
4.6. RELATIVIST SOLUTIONS TO THE PROBLEM	76
4.7. RECONCILIATION	76
4.7.1 <i>Specific Solutions</i>	78
4.7.2 <i>General Solutions</i>	82
4.8. REJECTING THE MORAL FIXED POINTS	91
4.9. CONCLUSION	96
5. THE FIRST-ORDER ARGUMENT FOR MORAL ERROR THEORY: THE RETURN OF THE IMPASSE.....	97
5.1. INTRODUCTION	97
5.2. THE FIRST-ORDER ARGUMENT FOR MORAL ERROR THEORY	98
5.3. CONCEPTUAL COMPETENCE	100
5.4. REJECTING THE SECOND ASSUMPTION	101
5.4.1. <i>Assessment</i>	102
5.5. REJECTING THE FIRST ASSUMPTION	105
5.5.1. <i>Assessment: Inconsistency</i>	107
5.5.2. <i>Assessment: How Can We Be Led Astray?</i>	110
5.5.3. <i>Assessment: How Can We Recognize Other Conceptual Truths?</i>	112
5.5.4. <i>Assessment: What About the Axiological Propositions Then?</i>	113
5.5.5. <i>Assessment: What Is Your Evidence?</i>	117
5.6. NOW WHAT?	117
5.6.1. <i>Conservationism</i>	118
5.6.2. <i>Fictionalism</i>	119
5.6.3. <i>Substitutionism</i>	120
5.6.4. <i>Abolitionism</i>	122
5.7. CONCLUSION	123
6. CONCLUSION.....	124
7. NEDERLANDSE SAMENVATTING	129
8. REFERENCES.....	135

Acknowledgements

Every PhD project has its ups and downs. Sometimes the ups outnumber the downs, but not for this PhD project. Or, at least, it is not clear to me if they did. This does not mean that I had a terrible time during these last four years. Nor does it mean that I regret embarking on this project, or that I would have decided differently had I known this four years ago. It does, however, mean that I am deeply grateful for all the people that helped me along the way.

First, I want to thank everyone on the administrative side of the faculty who helped me. I would like to especially thank Marga and Marion for all their help. I also want to thank my fellow PhD candidates. I have sorely missed our interactions over lunch or coffee and tea breaks during the COVID-19 pandemic. I am also thankful for everyone who participated in the metaethics reading group. The discussions we had were either a welcome distraction from my own project, or they were of much use to my own research. I would also like to thank Jonas Olson for his extensive comments on some of my work.

But when it comes to the matter of help regarding the content of my thesis, there are two people that I am most thankful for: Bart Streumer and Daan Evers. Without your feedback, advice, and support I would never have been able to complete this project. I have heard many horrifying stories about PhD supervisors. I am glad that I have only good ones to share. I am especially grateful for your constant reassurance that my health was more important than my PhD project.

Of course I also have to thank my family and friends for their support. First, I want to thank Peter and Renske for their support. I want to thank Maurice for the joy that he brings. I want to thank Lineke and Kris for partaking in many adventures involving dungeons and dragons. I want to thank Bianca for this too, and for many lovely vegan dinners. I also want to thank Kritika for her cooking and her friendship. I hope that we can one day claim to have battered and deep fried all vegetables known to man. And I want to thank my cat. Thank you for reminding me that there are far

more important things than my PhD project, such as: petting your cat, hugging your cat, scratching your cat behind their ears and under their chin. Rest in peace, Mies.

I am also grateful for my family and for all their support. To my brothers Tim and Ben I would like to say that, although as children we were not always on the friendliest of terms, I really treasure the bond that we have now; and to my parents Paul and Marianne, I will always be grateful for the love and support you have given me throughout my life. I could not have wished for better parents.

Lastly, I want to thank Laura. I am so happy that I decided I had to make some new friends one random morning during my first semester at university, that you happened to be part of the same seminar later that day, that I happened to overhear you talk about an online video game you were playing at the time, and that I had the courage to ask if you wanted to play it together some time. Many times I have thought that our friendship could not get any stronger and every time I was proven wrong. The value of your friendship is immeasurable. Without your support I could never have finished this thesis. Thank you.

0. Introduction

Recall the last time you had to make a moral decision. Maybe you had to decide whether you should give some money to the charity collector that knocked on your door. Maybe you had to think twice about ordering a new pair of shoes, given that they were likely made in a sweatshop. Maybe you were ordering food and you had to decide whether to order a meat-free dish or not. Or maybe it was something else entirely. When making such moral decisions, have you ever wondered whether there is an objective fact of the matter about which choice is right or wrong, or good or bad? This question sits at the centre of the debate between metaethical realism and anti-realism.

Realists believe there are objective or mind-independent ethical facts; anti-realists do not. The debate between them is old, and there is much to say in favour of either side. Anti-realism, for instance, seems to be a sensible response to the seemingly widespread and sustained moral disagreements we observe. If there are no objective ethical facts, there is nothing outside of us, no objective arbiter that can settle disagreements. Anti-realism also provides a more metaphysically parsimonious picture of the world than some realist theories. The existence of ethical facts is often seen as metaphysically problematic, because ethical facts seem fundamentally different from facts in other fields of inquiry like biology, chemistry, and physics. And given that we do not seem to need ethical facts to explain our moral behaviours, since purely psychological and evolutionary processes plausibly explain them, there seems to be no need to postulate their existence.

But anti-realist theories of course also face a number of objections, some of which are widely discussed in the literature. Here are some examples: first consider expressivism. Expressivists are non-cognitivists. They think that moral claims do not express beliefs but non-cognitive attitudes. That is, cognitivists believe that claims like “giving money to charity is good” express the belief that the act of giving money to charity possesses the moral property of goodness. Expressivists think that this is false. They think that this claim expresses a non-cognitive attitude like, for instance, approval of

giving money to charity. This means that expressivists do not try to explain what makes something good or bad, but what it is to believe that something is good or bad.

One important problem for the expressivist is the Frege-Geach problem (Geach, 1965). This problem takes many shapes, but one of them has to do with the validity of certain arguments. Sometimes we use moral claims in an unasserted way, for example, when our moral argument features a conditional claim. But such unasserted moral claims do not seem to express any non-cognitive attitude. This would mean that seemingly valid arguments are not valid on the expressivist account.

Let me give an example:

- P1.** Giving money to charity is good.
- P2.** If giving money to charity is good, then getting my brother to give money to charity is good.
- C.** Getting my brother to give money to charity is good.

P1 expresses the attitude of, say, approval of giving money to charity, according to the expressivist. For **C** to then follow from **P1** and **P2** the antecedent of **P2** (giving money to charity is good) needs to express the same thing that **P1** expresses. But it does not, because it is not asserted, like in **P1**. That is, **P2** does not in part express an attitude of approval of giving money to charity. This means that even if **P1** and **P2** are true, **C** does not follow on the expressivist account. But that is odd, because the above seems like a perfectly valid argument.

Expressivism also faces the moral attitude problem. This problem asks what non-cognitive attitude is expressed by moral claims specifically. Expressivists have given many possible attitudes that moral claims might express: approval or disapproval (Blackburn, 1984), a plan (Gibbard, 2003), being for blaming or praising (Schroeder, 2008) etc. But surely other normative claims express similar attitudes. That is, if I think that my grandmother's wallpaper is ugly, I seem to express an attitude of disapproval about my grandmother's wallpaper. So the question is: what makes a given expression of, for instance, disapproval uniquely moral and not aesthetic, epistemic, prudential, etc.?

Now consider relativism. Relativists are cognitivists. They think that moral claims express beliefs. But relativists also believe that nothing is simply morally good or bad, or right or wrong. Thus, for example, giving money to charity cannot be good *simpliciter*. It can only be good relative to something. What this something is, is a point of contention, but common suggestions are: moral standards (MacFarlane, 2014), other normative judgments (Street, 2008), and ends (Finlay, 2009). Thus, according to some relativists, charity can only be good if one's moral standards entail that it is good. That is, if one has, for instance, the moral standard that helping the less fortunate is good.

A prominent problem for relativism concerns its ability to account for moral disagreement. This might be somewhat surprising given that I just mentioned that widespread moral disagreement seems to be best explained by anti-realism. However, it seems that relativists cannot explain why these moral disagreements are proper disagreements, or so the objection goes. Let me explain: imagine that I believe that giving money to charity is bad, and you believe that giving money to charity is good. We clearly disagree. But according to some relativists' understanding of morality we do not. According to them, I believe that giving money to charity is bad according to my moral standards M_{author} , while you think giving money to charity is good according to your moral standards M_{reader} . Both these claims can be true. We both can even believe that both of these claims are true. For M_{author} and M_{reader} might simply be different moral standards. M_{reader} might entail the truth of your claim, while M_{author} might entail the truth of mine. So, what seems like a moral disagreement where at most one of us can be right turns out to be nothing of the sort, according to some forms of relativism. On the relativist account, what seems like disagreement is actually a case of us talking past each other. But that conflicts with the strong intuition that we do actually disagree.

Lastly consider the moral error theory. The moral error theory consists of two claims: the conceptual claim states that when we make a moral judgment we ascribe an objectively categorical or irreducible moral property to an act. The ontological claim states that such properties are never instantiated. Together these claims entail that all our (positive) moral claims are uniformly false.

One problem for the error theory is that it might overgeneralise and thereby undermine itself (Bedke, 2010; Cuneo, 2007). The idea is that if the error theorists' account is correct for our moral judgments, it might also apply to other kinds of normative judgment, like epistemic judgments. This would be so provided epistemic judgments and properties are sufficiently like the error theorist's conception of moral judgments and properties. Critics argue that they are sufficiently alike. This is a problem for moral error theorists. For if the epistemic error theory were true, all positive epistemic judgments would be false. This includes the judgment that there is a reason to believe the epistemic and moral error theory, meaning that the moral error theory would undermine itself.

Another objection to the moral error theory focuses on its implications. If the error theory is correct, then there are no moral truths. But, as error theorists claim, our moral practice relies on these truths. So, the truth of the error theory would entail that we can no longer morally praise or blame people or that we cannot use morality to guide our own decision-making. Given the extent to which these moral practices are integrated into our daily lives, the truth of the error theory would have enormous consequences. Because of this, moral error theorists have given several accounts of how our moral practice might be shielded from the truth of the error theory. But critics say this casts doubts on the truth of the error theory (Wright, 1992). For if regardless of the error theory's clear implications for our moral practice, this practice can be maintained, then that seems to suggest the error theorist's analysis might be off. That is, if accepting the error theory makes no difference to our moral practice, then surely that gives us reason to doubt that is true.

Anti-realists have of course responded to each of these objections. And critics have in turn responded to these responses, to which anti-realists have again responded, and on it goes. This thesis will not be a response to a response to a response. The aim of this thesis is to develop novel problems for these anti-realist theories. That is, the aim is to point out new issues with anti-realism, which requires the anti-realist to say something new in order to solve them. By developing these new problems this thesis does not strive to settle the debate between realism and anti-realism in favour of realism. That is, I do not live under the illusion that the new problems for anti-realism that this

thesis puts forward make anti-realism ripe for the metaethical dump. The scope of this dissertation is more modest. It is to put pressure on anti-realists by developing new problems for anti-realism; problems that require new solutions.

The new problems that this thesis develops are not all unrelated. That is, most of these new problems are derived from a concern for a specific intuition: the intuition that there are moral claims that must be true. Claims such as: clasp one's hands three times an hour in itself cannot be morally good (Foot, 2002). Or, it is *pro tanto* wrong to engage in the recreational slaughter of a fellow person (Cuneo and Shafer-Landau, 2014). Some authors think that the intuition that these claims must be true is explained by the fact that these claims are conceptually true. That is, they believe it simply follows from the concepts in these claims that they are true. In that sense they are more complex moral versions of claims like "a bachelor is an unmarried man". Chapters 2 to 5 of this thesis, which is a collection of self-standing papers, deal with the point of contact between this idea about conceptual moral truths and its underlying intuition and the various anti-realist theories. Chapter 1 gives an objection to the most prominent argument in favour of moral error theory on other grounds.

0.1 Chapter summaries

Chapter 1: Debunking the Argument from Queerness

This first chapter addresses the most prominent metaphysical argument in favour of the error theory: the argument from queerness. This argument claims that it is most reasonable to believe that there are no ethical facts, because ethical facts are 1) metaphysically queer, and 2) explanatorily redundant (Olson, 2014). Criticism of the argument from queerness generally focuses on refuting the reasons why error theorists take ethical facts to be metaphysically queer. This chapter, however, argues that even if we grant the error theorist that ethical facts are metaphysically queer and

explanatorily redundant, the error theorist's conclusion does not follow. The argument from queerness actually leads to a more modest conclusion, namely that it is most reasonable to suspend judgment on whether ethical facts exist. In other words, the argument from queerness cannot establish a moral error theory, but can at most establish that metaethical scepticism is true. This, I conclude, is cause for optimism, because moral error theory is more of a threat to our moral practice than metaethical scepticism.

Chapter 2: Going Negative on the Moral Fixed Points

In chapter 2, we turn to the sub-theme of this thesis: the point of contact between the idea of conceptual moral truths and anti-realism. In this second chapter I discuss the most recent defence of conceptual moral truths, Cuneo and Shafer-Landau's (2014) theory of the moral fixed points, and the criticism that is levelled against it. The objections against this view are generally motivated by what some see as a problematic conclusion of the view: the conclusion that error theorists must be conceptually deficient (Evers and Streumer, 2016). This conclusion follows from this view, because if there are conceptual moral truths, then someone who denies that there are moral truths must fail to properly grasp our moral concepts. That is, they must be conceptually deficient in some way.

I think that this criticism of Cuneo and Shafer-Landau seems sound. But in this chapter I explain that people like Cuneo and Shafer-Landau can avoid this criticism by distinguishing between *positive* and *negative* conceptual moral truths, and by focusing on the latter rather than the former. I further argue that this move does not create any new explanatory burdens and retains the benefits of Cuneo and Shafer-Landau's original account.

Chapter 3: Expressivism and Moral Irrelevance

Chapter 3 focuses on a particular species of negative conceptual moral truths: conceptual moral irrelevance claims. Irrelevance claims state that a certain descriptive feature is not morally relevant. For example, one might claim, when debating a white supremacist, that the mere colour of someone's skin is morally irrelevant. Some of these irrelevance claims seem conceptually true. That is, they seem to bear several marks of conceptual truths: they seem *a priori* knowable, denying them baffles conceptually competent people, etc. This means that there seem to be conceptual limits to what is morally relevant. Metaethical theories need to explain why there seem to be these conceptual limits.

In this chapter I argue that when it comes to explaining this realism fares better than expressivism. The fact that realists can appeal to the existence of moral properties helps them to explain these conceptual limits. But this also means that expressivists cannot give the same explanation, because they do not believe that there are moral properties out there in the world. Expressivists will therefore have to come up with their own explanation. This chapter goes over several potential expressivist explanations of why there are these conceptual moral irrelevance claims and finds all of them wanting.

Chapter 4: Relativism and Conceptual Moral Truth

Were there to be conceptual moral truths, this would not only be a problem for theories that deny that there are any moral truths. It would also be a problem for any metaethical theories that claim that there are no objective moral truths. For conceptual moral truths are true in an objective way. Their truth does not depend on your moral standards, your ends, your other normative judgments etc. Relativists, however, claim that the truth of moral judgments does depend on those things. That is, relativists believe that moral claims can at most be true relative to some moral standards, ends, other normative judgments, etc.

In this fourth chapter, I discuss two general approaches the relativist might adopt when faced with the idea of conceptual moral truths: reconciliation and rejection. As the name suggests, the reconciliatory approach tries to resolve the tension between relativism and conceptual moral truths by trying to bring them together in one theory. That is, the aim is to have a view that captures conceptual moral truths similar to those that others have defended, but which remains relativist for all other moral truths. The rejection approach is simply to deny that there are any conceptual moral truths. On its face, the rejection approach might seem like a better fit. However, as it turns out, this rejection approach comes with serious costs. And, perhaps surprisingly, the reconciliatory approach might be a viable alternative.

Chapter 5: The First-Order Argument for Moral Error Theory: The Return of the Impasse

Chapter 5 focuses on a recent proposal that suggests that *positive* conceptual moral truths might actually help to establish a moral error theory. The resulting error theory allegedly has a dialectical edge over moral error theories supported by way of a more traditional metaphysical argument. This is because metaphysical arguments for moral error theory generally devolve into an impasse with (non-natural) realists on fundamental metaphysical matters. This is what the new first-order argument supposedly avoids. But this chapter shows that this is false: the first-order argument for moral error theory also results in an impasse between error theorists and realists. Thus, *positive* conceptual moral truths do not improve the error theorist's dialectical situation. The chapter even concludes that this new way of supporting the error theory leaves error theorists dialectically worse off.

1. Debunking the Argument from Queerness¹

1.1. Introduction

Moral error theory states that ethical facts do not exist. It thereby threatens to undermine our moral practice, which, according to error theorists, depends on the existence of these facts. Following Mackie (1977), Olson (2014) presents a metaphysical argument from queerness (hereafter “the argument from queerness”, unless explicitly stated otherwise) in favour of this view. This argument is intended to show that it is most reasonable to believe that ethical facts do not exist, because they are metaphysically queer (hereafter “queer”, unless explicitly stated otherwise) and explanatorily redundant (Olson, 2014, p. 84-87). I will grant the error theorist both of these claims, yet argue that the argument from queerness does not warrant the rejection of ethical facts.² The argument from queerness therefore does not deliver what the error theorist promises. At most, it shows that we should not be realists: that is, it shows we should refrain from believing that ethical facts exist. But this does not give us anti-realism, because it does not show that we should deny that ethical facts exist. It merely gives us agnosticism about ethical facts. This weakens the threat that the argument from queerness presents to our moral practice.

First, I will discuss the error theory (section 1.2). Second, I will show that the argument from queerness only supports metaethical scepticism, i.e. suspending judgment on the existence of ethical facts (section 1.3). Third, I will address the most pressing objections to my argument (sections 1.4, 1.5, 1.6). Fourth, I will show that my criticism of the argument from queerness weakens its threat to our moral practice (section 1.7).

1.2. The Argument from Queerness

¹ Published in *Ratio* (2021).

² This does not undermine all error theories. Although moral error theorists generally defend their view with the argument from queerness, there are exceptions (see Streumer, 2017).

Moral error theory consists of a conceptual and an ontological claim: the conceptual claim states that our moral practice commits us to the existence of objective and irreducibly normative ethical facts, and the ontological claim states that these facts do not exist. I will grant the error theorist the truth of the conceptual claim but argue that she cannot establish the ontological claim. This requires a proper understanding of the main argument for the ontological claim: the argument from queerness.

According to Olson (2014, p. 84-87), the argument from queerness is a two-step argument.³ First, the error theorist argues that ethical facts are queer. Second, she establishes the explanatory redundancy of ethical facts.⁴ The combination and satisfactory defence of these steps would produce a convincing argument from queerness (Olson, 2014, p. 148).

The argument from queerness thus has the following structure:

- P1.** If an entity is queer and explanatorily redundant, it is most reasonable to believe that it does not exist.
- P2.** Ethical facts are queer.
- P3.** Ethical facts are explanatorily redundant.
- C.** It is most reasonable to believe that ethical facts do not exist.

If an entity is explanatorily redundant, this means our best explanations do not rely on the existence of that entity. If an entity is queer, this means that it is ontologically suspicious (Olson, 2014, p. 84): its queerness is a reason to believe that it does not exist.

But what is it for an entity to be queer? There are two general interpretations of queerness: metaphysical and epistemological. Shepski interprets Mackie as saying metaphysically queer entities

³ “[T]he argument from queerness has a bipartite structure, comprising at least one queerness argument and debunking explanations of moral beliefs” (Olson, 2014, p. 87). Queerness arguments try to establish that ethical facts are queer and debunking explanations of moral beliefs try to establish that ethical facts are explanatorily redundant.

⁴ We often take the explanatory redundancy of an entity by itself to show that it does not exist. Olson does not think this holds for ethical facts (2014, p. 84-85).

are queer because they are *sui generis* i.e. fundamentally different (2008, p. 377).⁵ Olson tries to be more specific and suggests that several error theorists think ethical facts are metaphysically queer because “moral properties and facts are ontologically fundamental additions to a scientifically based naturalistic worldview” (Olson, 2014, p. 85). Olson himself rejects this emphasis on naturalism, but does not provide an alternative account (2014, p. 86-87). Yet it seems that the error theorist is at least committed to the view that a metaphysically queer entity is ontologically mysterious, peculiar, or odd in some way. Otherwise it is unclear why the metaphysical queerness of an entity would count against its existence. So, to be a bit more precise one might suggest that an entity is metaphysically queer if the existence of that entity would entail a fundamental addition to what we take to exist.⁶

One might believe that this does not capture what the error theorist means when she makes a metaphysical queerness judgment. One might think that a queerness judgement is not about what we take to exist but more straightforwardly about *what there is*. That is, one might think that the error theorist believes that an entity is metaphysically queer when its existence would entail a fundamental addition to what there is. But this does not seem right. For if entity E is a fundamental addition to what there is, it follows that E is not part of what there is, and it therefore immediately follows that E does not exist. That is not what the error theorist takes the metaphysical queerness of an entity to show. If they did, it would be unclear why error theorists like Olson believe that the explanatory redundancy of an entity is an integral part of the argument from queerness. Error theorists do not think that the metaphysical queerness of an entity by itself shows that it does not exist. They believe it just shows that the entity is ontologically suspicious (Olson, 2014, p. 84). This supports taking an entity to be metaphysically queer if its existence would entail a fundamental addition to what we take to exist. For although an entity that is not part of what we take to exist is to some extent suspect, it might still exist.

⁵ This of course relativizes the queerness, but queerness has to be relative. An entity cannot be queer *simpliciter*; it can only be queer compared to something else. In case of metaphysical queerness, entities can only seem queer compared to other entities, and interesting queerness judgments compare the queer entity to the entities we believe to exist.

⁶ ‘We’ refers to the community of the one inquiring into the queerness of an entity.

But the phrase “what we take to exist” is still a bit vague. It does not make explicit that, when it comes to metaphysical queerness, the additions that error theorists are concerned about are additions on the fundamental ontological level. In what follows, I will therefore assume that an entity is metaphysically queer if the existence of that entity would entail a fundamental addition to our ontological framework, without this framework necessarily being a naturalistic one.⁷

What about epistemological queerness? A fact is epistemologically queer when the apprehension of this fact would involve a queer faculty (Mackie, 1977, p. 38-39). In other words, a fact is epistemologically queer when knowing this fact entails a fundamental addition to our cognitive capacities.

These two different kinds of queerness lead to different interpretations of the argument from queerness. The metaphysical interpretation of the argument supposedly shows that ethical facts do not exist because they are metaphysically queer and explanatorily redundant. The epistemological interpretation relies on the idea that the apprehension of these facts would involve a queer faculty. If the epistemic argument from queerness is successful, it merely establishes moral scepticism: the view that we cannot have any moral knowledge. It cannot establish that ethical facts do not exist (Olson, 2014, p. 100). Since I am concerned with whether the argument from queerness can establish that the error theory is true, I will focus on the metaphysical argument from queerness.

However, before we move onto my criticism of the metaphysical argument from queerness some things need to be cleared up. There are many ways in which error theorists have defended P2 and P3. As for P2: Olson thinks ethical facts are queer because they are irreducibly normative (2014, p. 116–126), Garner thinks they are queer because they imply a demand-without-a-demander (1990, p. 143), and Mackie, amongst other things, points to how ethical facts inexplicably supervene on natural facts to argue that they are queer (Mackie, 1977, p. 41).⁸ As for P3: some argue that evolution can explain our moral practice without invoking ethical facts (Joyce, 2001, p. 168; Olson,

⁷ Note that I use the term ‘ontological framework’ in a theoretically neutral way and not necessarily in the Carnapian sense of the term.

⁸ Olson calls these queerness arguments, since they are intended to show that ethical facts are queer (2014, 84). They are not to be confused with the argument from queerness, which, as we have seen, is intended to show that it is most reasonable to believe that ethical facts do not exist.

2014, p. 142), others argue that psychological facts can do this job (Harman, 1977, p. 6). It needs to be clear that it does not matter for the purposes of this paper which arguments for P2 and P3 are most convincing. This is for two reasons: one, I will simply assume, for the sake of argument, that ethical facts are queer and explanatorily redundant. Two, as mentioned, my criticism targets the argument from queerness as laid out by Olson (2014, 84-85). This only establishes the general form of the argument from queerness. It does not mention any specific reasons for why ethical facts are queer and explanatorily redundant. My criticism will therefore hold regardless of these specific reasons. This also means that my criticism applies not just to Olson, but to any other error theorist that uses this argument from queerness as well.

To recap, I grant error theorists three of their core claims: (a) the conceptual claim, i.e. the claim that our moral discourse commits us to the existence of objective and irreducibly normative ethical facts, (b) ethical facts are queer, and (c) ethical facts are explanatorily redundant. I will now argue that even if we grant the error theorist (a)-(c), the argument from queerness fails to establish that ethical facts do not exist.

1.3. Debunking Queerness Judgments

Error theorists believe that the queerness of an entity is a reason to believe it does not exist. But this is not necessarily the case.

Consider the following example: imagine a community of people who are, and always have been, completely cut off from the rest of humanity. And further imagine that because of a genetic defect every member of this community is born blind. Light, which we cannot access or detect through smell, sound, taste, or touch, will be alien to the people of this community. Their inability to access light presumably makes it difficult for them to place light within their ontological framework.

For their ontological framework only encompasses the kinds of entities they can epistemologically access, and light is distinct from any kind of entity they can access.⁹

The example illustrates how the epistemic inaccessibility of an entity can ensure its exclusion from our ontological framework. This means that, were members of the blind community to somehow discover light, they would have to make an addition to their ontological framework. But would they have to make a *fundamental* addition to their ontological framework? Here we should, as Mackie does (1977, p. 38), recognize the relation between the degree of difference between the way one would access the entity in question and one's other cognitive capacities, and the likelihood of that entity being fundamentally different from the entities one does believe to exist. The greater this difference, the greater the likelihood that the entity in question is fundamentally different from the entities comprising our ontological framework. The question is: is sight different enough from the blind people's other cognitive capacities to entail the judgment that light is queer?

Assume, for purposes of elucidation, that it is.¹⁰ This means that light would be metaphysically queer to the people of the blind community. This in turn means that inaccessibility of an entity can make an entity metaphysically queer, in the sense that its existence would be a fundamental addition to our ontological framework. But the inaccessibility of an entity need not be a reason to believe it does not exist. For this inaccessibility might simply stem from a limitation of one's cognitive capacities, as is obviously the case in the blind community example. Recognizing this relationship between our cognitive capacities and our ontological framework shows that the queerness of an entity is not always a reason to believe it does not exist. That is, an accurate queerness judgment does not necessarily provide a reason to deny the existence of the entity in question.¹¹ However, this is only the case when the queerness judgment concerns an epistemically

⁹ Of course they can potentially epistemically access light if someone who has epistemic access to light tells them about it. Or if for some reason members of this community develop the ability to see. This, however, is not possible without violating the stipulations of the thought experiment. So, as it stands, we can claim that the people of the blind community cannot epistemically access light.

¹⁰ It is not a problem if this assumption turns out to be false. An analogy can be illuminating without being perfect. As long as ethical facts are inaccessible and the way to access them is utterly different from our other cognitive capacities, my argument works.

¹¹ It might seem that Olson already accounts for this possibility by discussing examples of seemingly queer entities that are

inaccessible entity, and when the cognitive capacity needed to access the entity is sufficiently different from one's other cognitive capacities.

Interestingly, the error theorist's epistemological argument from queerness ticks both boxes for ethical facts (Mackie, 1977, p. 38-39). As Mackie says, "if we were aware of [ethical facts], it would have to be by some special faculty of moral perception or intuition, utterly different from our ordinary ways of knowing everything else" (1977, p. 38). This means we might judge ethical facts to be metaphysically queer just because if they existed we could not access them and the way to access them would be very different from our other cognitive capacities. The judgment that ethical facts are queer might therefore stem from the limitations of our cognitive capacities. But queerness from this source is not a reason to deny the existence of ethical facts.

Yet, as we've seen, the metaphysical argument from queerness relies on the queerness of an entity to be a reason to deny its existence. More specifically, my criticism undermines P1: the premise that says that if an entity is queer and explanatorily redundant, it is most reasonable to believe it does not exist.¹² To establish the consequent of this premise, the queerness of an entity must be a reason to deny its existence. But the queerness of ethical facts is not necessarily a reason to deny their existence, so the consequent cannot be established in the case of ethical facts. Thus, the argument from queerness fails to show that it is most reasonable to believe that ethical facts do not exist. And therefore fails to establish the error theorist's ontological claim: the claim that ethical facts do not exist.

But although its force is diminished, the argument from queerness is not completely impotent. The queerness and explanatory redundancy of ethical facts are still reasons to refrain from

not ontologically suspicious, e.g. neutrinos, aardvarks, and impressionist paintings (2014, p. 87). These entities might be "*prima facie* queer", but on closer inspection "we realize that they are actually parts of the best explanations of some of our observations and beliefs. At this point they no longer seem ontologically suspicious, although we may find them queer in other ways" (Olson, 2014, p. 87). However, these examples are clearly different from the example of the blind community. To the blind community light is not just *prima facie* queer, it seems genuinely queer, i.e. even on closer inspection its existence would entail a fundamental addition to their ontological framework. The example of the blind community therefore identifies a different, and more problematic way in which queer entities can be ontologically benign.

¹² At first glance, my objection might seem to target **P2**: the premise that says ethical facts are queer. This is not the case. My objection does not establish the inaccuracy of our queerness judgment. It merely shows that the queerness of an entity need not provide a reason to deny its existence.

believing that ethical facts exist. Thus, an accurate and more modest first premise should merely state that it is most reasonable to refrain from believing that queer and explanatorily redundant entities exist.

The difference between the conclusion of the error-theoretical and the modest argument from queerness is the following: the error-theoretical conclusion encourages us to believe that ethical facts do not exist, while the modest conclusion cautions us against believing that ethical facts exist.

This is the modest argument from queerness:

- P1*.** If an entity is queer and explanatorily redundant, it is most reasonable to refrain from believing it exists.
- P2.** Ethical facts are queer.
- P3.** Ethical facts are explanatorily redundant.
- C*.** It is most reasonable to refrain from believing that ethical facts exist.

Thus, the queerness of ethical facts is not necessarily a reason to deny their existence, which undermines the first premise of the argument from queerness. The modest argument from queerness reflects the diminished force of our queerness judgments and has a more modest conclusion than the error-theoretical version: namely, that it is most reasonable to refrain from believing that ethical facts exist. This conclusion entails metaethical scepticism, i.e. suspending judgment on whether ethical facts exist or not, as opposed to an error theory. For it warrants neither the belief that ethical facts exist nor that they do not exist.

This is not to say that metaethical scepticism is a viable metaethical position. Metaethical scepticism might be completely untenable. But this issue is separate from the one discussed here: the issue of what the argument from queerness can and cannot show. As we have seen, the answer

is that the argument from queerness cannot establish the error theorist's ontological claim: ethical facts do not exist. It merely supports metaethical scepticism.

1.4. Objection: Bait-and-Switch

One way in which the error theorist might object to my argument is to claim that it is a bait-and-switch. The error theorist might say the following: you claim to argue that the metaphysical argument from queerness cannot establish the ontological claim. But in your argument you turn to epistemological worries. You then claim we error theorists believe that the epistemological argument from queerness can only get us to agnosticism about ethical facts. Based on this, you conclude that the metaphysical argument from queerness can only get us agnosticism as well. So instead of arguing against the metaphysical argument from queerness, you confuse it with the epistemological argument from queerness and make an uncontroversial claim about that argument.

Although I can understand why one might say this, it is mistaken. My argument is as follows: The metaphysical argument from queerness depends on a metaphysical conception of queerness. This conception of queerness is limited by our epistemic capacities, because what our ontological framework encompasses is limited by our cognitive capacities. That is, our ontological framework encompasses only entities that are accessible to us.¹³ This means that some metaphysical queerness judgments need not entail the ontological suspiciousness of the entity in question. For we might judge an entity to be metaphysically queer, i.e. a fundamental addition to our ontological framework, simply because it is (d) epistemically inaccessible, and (e) the way we could access the entity, were it to exist, is very different from our other cognitive capacities.

Here and only here is where the epistemological argument from queerness comes into play. For the epistemological argument from queerness entails the claim that ethical facts are (d) epistemologically inaccessible. And it entails that (e) access to ethical facts, were they to exist, would

¹³ This does not mean that our ontological framework encompasses everything we could access, because some things that do not exist are in principle accessible to us, e.g. unicorns, hippogriffs, leprechauns, etc. It merely means that an entity must at least be epistemically accessible to be part of our ontological framework.

require “some special faculty of moral perception or intuition, utterly different from our ordinary ways of knowing everything else” (Mackie, 1977, p. 38).

Thus, my argument is that if conditions (d) and (e) *do not* hold for a queer entity we can apply the error theoretical argument from queerness, but if these conditions *do* hold for a queer entity we can only apply the modest argument from queerness. As the epistemological argument from queerness makes clear, (d) and (e) hold for ethical facts. Thus, the relation between the metaphysical and epistemological argument from queerness is that the latter provides the conditions for the failure of the former. This is different from confusing the one with the other. There is therefore no switching of the bait.

1.5. Objection: Accessibility

This does leave us with an obvious way the error theorist might object to my argument: deny that conditions (d) and (e) hold for ethical facts (note that, for ease of exposition, from this point on I will take talk of epistemic accessibility as shorthand for both conditions (d) and (e)). Thus, the error theorist might claim that if ethical facts were to exist, we would be able to access them. But this requires her to explain how we would be able to access them – an explanation they accuse the moral realists of failing to give without relying on “a lame answer” like intuition (Mackie, 1977, p. 39).

Furthermore, their explanation needs to be such that the moral realist cannot simply hijack it. For many non-naturalist moral realists endorse the conceptual claim: they also think that our moral discourse commits us to the existence of irreducibly normative ethical facts. So, if the error theorist gives a brilliant epistemic account of how we could access these facts if they existed, it seems more difficult for them to reject non-naturalism.

1.6. Objection: Doxastic Move

Instead of resisting my criticism, the error theorist might accept it and embrace the modest argument from queerness. They might, however, deny that it leads to metaethical scepticism. They could argue that the modest argument from queerness in fact indirectly shows that ethical facts do not exist.

The error theorist might argue that the modest argument from queerness leaves us in a place not unlike the one potential atheists and agnostics occupy. That is, there are many convincing arguments that show that we have no reason to believe that God exists. Some people believe that these arguments only warrant agnosticism (or scepticism) about the existence of God. Others, however, believe that these arguments warrant atheism, i.e. the wholesale rejection of God. Or, to use a less charged example, since we have no reason to believe that unicorns exist, we can either suspend judgment or deny their existence.

Assume that we are allowed to deny the existence of unicorns.¹⁴ The error theorist might then argue that this also holds in the metaethical case. She can say that, given the conclusion of the modest argument from queerness, we can deny the existence of ethical facts. That is, the error theorist might argue that we can make a doxastic move from it being reasonable to refrain from believing that ethical facts exist to a reasonable belief that ethical facts do not exist. However, for this to work, the error theorist needs to show that there is indeed no reason to believe that ethical facts exist. The modest argument from queerness on its own does not show this.

We might argue that the error theorist cannot show that there is no reason to believe that ethical facts exist, because there in fact are reasons to believe that ethical facts exist. Here are two arguments we might use to show this.

First, we can argue that the Moorean argument for moral realism gives us a reason to believe that ethical facts exist. According to this argument, the statements (f) 'torturing the innocent is wrong', and (g) 'there are no ethical facts' cannot both be true. Since it is initially more plausible that

¹⁴ This might be a contentious assumption, but following the general strategy of this paper, I will simply grant the error theorist all she needs, without granting her the ontological claim.

(f) is true than that (g) is true, the argument concludes that there is reason to believe that ethical facts exist (Huemer, 2005, p. 115-117).

But why does (f) seem plausible? That it appears to us that (f) can simply be the product of evolution or upbringing. But if that is the case it is difficult to see how (f) seeming to be more plausible than (g) gives us any reason to believe that ethical facts exist. For this to provide such a reason it seems that the appearance that (f) should somehow give us insight into the ethical fact that (f): torturing the innocent is wrong.¹⁵ This means the Moorean argument commits us to the accessibility of ethical facts. But that means the error-theoretical argument from queerness is on the table again. After all, it is the inaccessibility of ethical facts that pushes us towards the modest argument from queerness in the first place. The Moorean argument therefore does not help us reject the error theorist's doxastic move.

Second, we might refer to Enoch (2011) who claims that irreducibly normative truths are indispensable to deliberation. That is, according to Enoch, by deliberating we commit ourselves to the existence of irreducible normative truths (2011, p. 70-79). If Enoch's argument is sound, it gives us a reason to believe that ethical facts exist. However, these irreducible normative truths only seem deliberatively indispensable on Enoch's specific conception of deliberation (2011, p. 71-72). Part of this conception is the idea that we discover and do not create irreducible normative truths when we deliberate (Enoch, 2011, p. 73). This means that to deliberate in the way that is significant to Enoch's argument we need to believe that we have epistemic access to irreducible normative truths. How else are we going to discover these truths? If this belief in the accessibility of irreducible normative truths is not justified, Enoch's argument cannot give us a reason to believe that ethical facts exist. And if this belief is justified, the error-theoretical argument from queerness is on the table once again.

¹⁵ It can of course appear that (f) because we hold other moral beliefs that support (f) and therefore make it appear that (f). But this alone does not give us a reason to believe ethical facts exist, because these other moral beliefs can of course themselves be the product of evolution or upbringing.

So, we cannot easily argue that there are reasons to believe that ethical facts exist without committing ourselves to (a belief in) the epistemic accessibility of ethical facts. But such a commitment allows the error theorist to raise her argument from queerness. For my criticism of the argument from queerness relies on the inaccessibility of ethical facts.

What we can argue is that there are important differences between unicorns and ethical facts, and that these differences show why we can make the doxastic move in the unicorn case, but not in the metaethical case. That is, they might show why we can move from it being most reasonable to refrain from believing that unicorns exist, to it being most reasonable to deny their existence, but that we cannot do so with ethical facts. Two of these differences come to mind.

1.6.1. Difference in Epistemic Standards

The first difference concerns the epistemic standards for what one can reasonably believe. This affects the validity of the doxastic move in the following way: the higher the epistemic standards, the stronger our support for what we can reasonably believe must be. This might mean that when epistemic standards are high, refraining from believing does not warrant denial of existence. Thus, if the metaethical case involves higher epistemic standards than the unicorn case, this can explain why we can make the doxastic move in the unicorn case but not in the metaethical case.

To show there is this difference in epistemic standards we can call on epistemic contextualism, according to which the truth conditions of knowledge ascriptions vary based on context (Cohen, 1999, p. 57).¹⁶ This helps us reject the error theorist's doxastic move in two ways: first, the truth of knowledge ascriptions like 'S knows she has hands' varies based on the intellectual context. For example, in most contexts the statement 'S knows she has hands' is true, but when uttered during an academic lecture on scepticism it is false. This is because the context of everyday

¹⁶ Epistemic contextualist considerations normally concern epistemic standards for true knowledge ascriptions. But the same considerations seem to hold for the epistemic standards for reasonable belief as well. They might even more appropriately apply to reasonable belief, because what is reasonable to believe seems, at least *prima facie*, more context variable than knowledge ascriptions.

life seems to involve lower epistemic standards than the context of academia, especially when scepticism is discussed. This is at least partly because in the academic context there is more pressure to not jump to conclusions, consider alternative explanations, etc. in order to get it right. There is a similar difference in intellectual context between the metaethical and the unicorn case. The study of metaethics is purely academic, while the unicorn case does not have this link with academia. One can therefore argue that the epistemic standards in the metaethical case are higher than in the unicorn case.¹⁷

Second, some philosophers (e.g. Lewis, 1996, p. 556; Stanley, 2005) argue that the truth of knowledge ascriptions differ based on the practical stakes. For example, if I have to catch a train to London or risk losing my job, the stakes for being certain about the time of departure are quite high.¹⁸ In this case, having checked the time of departure while sleep-deprived the evening before might not warrant saying that I know the train departs at 1 o'clock. My circumstances require stronger support to attain a true knowledge ascription. The need for strong support in a high stakes case like this shows there are high epistemic standards involved. The stakes in the metaethical case are significantly higher than in the unicorn case. For, according to the error theorist, our moral practice, e.g. making moral judgments and acting accordingly, depends on whether ethical facts exist or not. Given the role these moral practices play in our everyday interactions and decision-making, getting this ontological issue right is very important. In contrast, although interesting, practically there is not much at stake in discovering whether unicorns exist. Thus, the stakes in the metaethical case are significantly higher than the stakes in the unicorn case. This means the same holds for the epistemic standards in these cases.

In short, if contextualism about reasonable belief is correct, then these considerations suggest that the epistemic standards in the metaethical case are significantly higher than those in the

¹⁷ This suggests that in everyday contexts the error theorists should believe that ethical facts do not exist, even though this is the exact opposite of what most of them argue. They argue that we can believe that ethical facts do not exist in an academic context, but that we cannot, or need not, maintain this belief in everyday contexts (Mackie, 1977, p. 16; Olson, 2014, p. 192). However, the intellectual context response does not entail that we can be error theorists in everyday life. For talking metaethics, no matter where or when, will automatically mean there are high epistemic standards. This is what I mean by 'metaethics is purely academic'.

¹⁸ For a similar example, see Cohen (1999, p. 58).

unicorn case. This limits what we can reasonably believe in the metaethical case, which shows that there is a difference in the validity of the doxastic move between these cases. Thus, that it is reasonable to refrain from believing that ethical facts exist does not warrant the belief that ethical facts do not exist. The conclusion of the modest argument from queerness therefore does not indirectly lead to the error theoretical conclusion.

1.6.2. Difference in Accessibility

The second difference between the metaethical and unicorn case concerns their accessibility. My criticism of the argument from queerness only holds for entities that we cannot epistemically access. Error theorists think that ethical facts are inaccessible. Unicorns, however, are epistemically accessible. That is, if unicorns were to exist, we would be able to detect them.

The accessibility of unicorns explains why we can claim that it is most reasonable to believe unicorns do not exist in the following way: if unicorns exist, we can detect them. However, we never see unicorns. The best explanation for this is that unicorns do not exist. This establishes an added presumption against the existence of unicorns.¹⁹

The error theorist cannot make the same inference to the best explanation in the metaethical case. For if an entity is inaccessible, the fact that we never observe it is void of any argumentative force. This is how the difference in accessibility between ethical facts and unicorns explains the difference in the validity of the doxastic move between these cases.

1.6.2.1. Objection: Overgeneralisation

¹⁹ Of course, if we take unicorns to be queer (for example, because their horns supposedly possess magical properties) and explanatorily redundant, their epistemic accessibility allows us to raise an error-theoretical argument from queerness about unicorns. This would do away with the need for a doxastic move in the unicorn case altogether, because that argument from queerness simply allows us to deny the existence of unicorns. But if we do not take unicorns to be queer (or use a different non-queer example, like a Pegasus (Shepski, 2008, p. 375)), the inference to the best explanation does allow us to make the doxastic move.

The error theorist could object that the accessibility response in the previous section overgeneralises. She might say: imagine an invisicorn. Like ethical facts, it is not accessible. It is invisible, immaterial, odourless, causally inert, etc. We have no reason to believe it exists. And because it is inaccessible, we cannot raise an error-theoretical argument from queerness, nor can we make an inference to the best explanation concerning the non-existence of invisicorns. We therefore cannot make the doxastic move in the invisicorn case either. This means that the accessibility response commits us to invisicorn scepticism. But that is not all, for the invisicorn is obviously but one of infinitely many similar examples. Thus, the accessibility response overgeneralises. It commits us to the absurdity of scepticism rather than nihilism about all inaccessible entities.

We can reply in at least three ways: first, it is questionable whether the idea of an invisicorn is coherent. Invisicorns are supposed to be invisible, odourless, causally inert, etc., yet somehow they are also supposed to be animals: living organisms that at the very least can move. If invisicorns are incoherent, they cannot exist. If they cannot exist, we are not committed to invisicorn scepticism.

Second, if invisicorns are coherent, then the arguments in section 6.1 also show that there is a difference in epistemic standards between the metaethical and invisicorn case. The doxastic move is therefore valid in the invisicorn case, which means the accessibility response does not commit us to invisicorn scepticism.

Third, we can embrace invisicorn scepticism. This is not as absurd as the error theorist might think. In general, it seems most reasonable to suspend judgment on whether an inaccessible entity exists. For example, say someone claims that there is a unique and invaluable gem on a specific planet beyond the edge of the observable universe. We have no reason to believe this is true. We also have no reason to believe it is false. And on top of that, we have no way to learn whether it is true or false. It therefore seems just as odd to deny that there is such a gem than to claim that there

is. That is, if we are strict about what we can and cannot reasonably believe. After all, we have no reason to hold either belief. Suspending judgment therefore is the most reasonable thing to do.²⁰

We can frame it like this: someone who claims it is most reasonable to believe that invisicorns exist is obviously wrong. This might tempt us to say that it is most reasonable to believe that invisicorns do not exist. However, there is no reason to believe this either.²¹ Furthermore, we would then in effect be saying that it is most reasonable to believe that an entity we cannot possibly access does not exist, which seems epistemically arrogant. What we should do is take the absurdity of the claim that invisicorns exist to stem from the obvious fact that there is no reason to believe that they exist. This fact is compatible with both denying the existence of invisicorns and suspending judgment about their existence. But, given the absence of reasons to believe that invisicorns do not exist and thus the epistemic arrogance of denying their existence, suspending judgment is the most reasonable thing to do.

In short, both the epistemic standards and accessibility response identify differences between the metaethical and unicorn case. These differences can explain why we can make the doxastic move in the unicorn case, while we cannot do so in the metaethical case. That is, they can explain why we can move from it being reasonable to refrain from believing that unicorns exist to a reasonable belief that unicorns do not exist, while we cannot do so in the metaethical case. The

²⁰ One might think it is more reasonable to deny the existence of the invisicorn than the existence of the gem, and that the gem example therefore does not take away the absurdity of invisicorn scepticism. This might be for at least two reasons: first, there seems to be a difference in the inaccessibility of the gem and invisicorn case. The gem is inaccessible because of the position of the earth in the universe, but the invisicorn is inaccessible because our cognitive capacities cannot detect it. The error theorist can argue that this difference allows for a difference in what is reasonable to believe about these cases. So, she might argue that it is most reasonable to refrain from believing that the gem exists, but also that it is most reasonable to deny the existence of invisicorns. However, this only works if we take our cognitive capacities to be epistemically privileged. That is, if our cognitive capacities as a source of inaccessibility make it most reasonable to deny the existence of the inaccessible entity, one has to claim something that comes down to: if we do not have the cognitive capacities to detect it, it cannot exist. It is not just unclear how one could defend this claim; it also seems arrogant to make it. Second, we might feel it is more reasonable to deny the existence of invisicorns than the existence of the gem because of unease with the idea of invisicorns itself. Remember, the idea of invisicorns might be incoherent. But, as mentioned, this simply shows that the accessibility response does not commit us to invisicorn scepticism.

²¹ The error theorist might argue that the recent fabrication of the invisicorn explains why it is most reasonable to believe invisicorns do not exist. Depending on how much stock we should put into this argument, this might push us away from embracing invisicorn scepticism. However, it would do so by identifying an apparently important difference between invisicorns and ethical facts, for ethical facts are not a recent fabrication. Thus, even if successful in the invisicorn case, this fabrication argument does not translate to the metaethical case.

modest argument from queerness therefore does not indirectly lead to the error-theoretical conclusion that ethical facts do not exist.

1.7. Impact on Our Moral Practice

The modest argument from queerness does not warrant the rejection of ethical facts. It merely supports metaethical scepticism: suspending judgment on whether ethical facts exist or not. Theoretically this is interesting, since the argument from queerness is the most prominent argument in favour of the error theory. But there is also a more practical payoff: metaethical scepticism is less of a threat to our moral practice than the error theory.

It seems clear that the error theory threatens to undermine our moral practice. If the theory is correct, claims like 'torturing people for fun is wrong' are false. And as Wright argues, the error theorist's view "relegates moral discourse to bad faith" (1995, p. 2), because our moral discourse commits us to the existence of ethical facts. This is a bitter pill to swallow.

It is therefore not odd that error theorists often try to argue that, despite the apparent implications of their theory, we can hold on to our moral practice. Mackie, for instance, claims that his error theory is "not merely distinct but completely independent" (1977, p. 16) from normative ethics, but offers no defence for this claim. Olson's moral conservatism supposedly allows us to preserve our moral discourse because of a similar compartmentalization of our beliefs. He claims that "it is a psychologically familiar fact that we sometimes temporarily believe things we, in more reflective and detached contexts, are disposed to disbelieve" (Olson, 2014, p. 192). We can therefore have moral beliefs in an everyday context even if, in a reflective and detached context, we believe there are no ethical facts. And our moral practice can still rely on these moral beliefs. Joyce argues that we should change our moral thought and speech to avoid the consequences of the error theory being true (2001, p. 206-231). His revolutionary fictionalism urges us to make moral claims pretend assertions. That is, to claim that torture is wrong is not to assert the proposition that it is wrong to

torture, it is to pretend to assert it. These pretend assertions still allow us to engage in our moral practice without any ontological commitments.

All these suggestions try to shield our moral practice from the threat of the error theory in some way or other. But there is a general unease with them: they create tension within the error theory. For even though our moral practice commits us to the existence of ethical facts, the fact that ethical facts do not exist apparently does not affect our moral practice.

Error theorists of course try to downplay this tension. But it is interesting to note that a preservationist move on behalf of metaethical scepticism seems less strained. That is not to say there is no tension, just that there is less of it. This is because metaethical scepticism does not imply that our moral convictions are uniformly false, nor does it relegate our moral discourse to bad faith. It only entails that the belief that our moral practice relies on, the belief that ethical facts exist, is unwarranted, not false. Conserving our moral practice therefore seems less problematic if we suspend judgment on whether they exist than if we deny their existence.

To make this even clearer it might be helpful to return to the parallel with the agnostic and the atheist. Imagine that an atheist and an agnostic agree that our religious practices have practical benefits. For instance, suppose they agree that praying is good for one's mental health. The claim then is that if both want to retain these benefits, there is less tension in the view of the praying agnostic than in the view of the praying atheist. This is because the agnostic, unlike the atheist, does not believe the negation of the belief that prayer depends on: the belief that God exists. The agnostic only suspends her judgment on this matter. The same goes for the metaethical sceptic and the error theorist in the case of our moral practice and the belief that ethical facts exist. This is why metaethical scepticism is less of a threat to our moral practice.

1.8. Conclusion

The metaphysical argument from queerness purports to show that it is most reasonable to believe that ethical facts do not exist. I have argued that even if we assume the queerness and explanatory redundancy of ethical facts, the argument cannot show this. In establishing its error-theoretical conclusion, the metaphysical argument from queerness takes the metaphysical queerness of an entity to provide a reason to deny its existence. But, as I have argued, the metaphysical queerness of an entity fails to provide such a reason if two conditions hold: first, the entity is epistemically inaccessible. Second, the way we would be able to access the entity, were it to exist, is utterly different from our other epistemic capacities. The epistemological argument from queerness establishes that both conditions hold for ethical facts. This weakens the conclusion we can draw from the metaphysical argument from queerness. The modest argument from queerness therefore does not establish that ethical facts do not exist, but merely shows that it is most reasonable to refrain from believing that ethical facts exist. This supports metaethical scepticism rather than a moral error theory, which is cause for optimism. For metaethical scepticism is less of a threat to our moral practice than moral error theory.

2. Going Negative on the Moral Fixed Points

2.1. Introduction

Non-naturalism faces major metaphysical and epistemological worries. Cuneo and Shafer-Landau argue that non-naturalists can deal with the metaphysical worries, if they take their non-naturalism to pertain to moral concepts instead of properties (2014, p. 402). For this allows them to remain uncommitted on the metaphysical nature of moral properties. To deal with the epistemological worries, Cuneo and Shafer-Landau argue that some moral propositions, ‘the moral fixed points’, are conceptually true (2014, p. 403). They are therefore knowable through a proper grasp of moral concepts.

Several philosophers (Evers and Streumer, 2016; Ingram, 2015, 2018) have criticized this proposal. In this paper, I argue that this criticism only affects Cuneo and Shafer-Landau’s account because their account focuses on moral fixed points that are positive. Were Cuneo and Shafer-Landau to defend an account that focuses on negative moral fixed points, they could avoid this criticism.

In section 2.2, I explain Cuneo and Shafer-Landau’s account in more detail. In section 2.3, I explain the main criticism of their account. In section 2.4 and 2.5, I discuss how focussing on negative moral fixed points can help deal with this criticism. In section 2.6, I argue that although the negative account makes a weaker claim, this does not create any new explanatory burdens. In section 2.7, I show that the negative account retains the benefits of Cuneo and Shafer-Landau’s positive account.

2.2. Cuneo and Shafer-Landau’s Account

Cuneo and Shafer-Landau argue that some moral propositions are conceptually true. They call these propositions ‘moral fixed points’. Examples of moral fixed points include (Cuneo and Shafer-Landau, 2014, p. 405):

a) It is *pro tanto* wrong to engage in the recreational slaughter of a fellow person

and

b) If acting justly is costless, then, *ceteris paribus*, one should act justly.

Cuneo and Shafer-Landau think a) and b) are candidates for conceptual truth because they have the following features (2014, p. 407-408): 1) if true, they are necessarily true. 2) They have framework status, “fixing the boundaries as to what counts as a type of subject matter” (2014, p. 407); that is, if you engage in moral thought but deny either a) or b), we have good reason to believe you use different moral concepts. 3) Denying either a) or b) bewilders conceptually competent people. 4) a) and b) seem *a priori* knowable.

Cuneo and Shafer-Landau argue that a) and b) are conceptually true because they are true in virtue of the essences of the concepts involved. That is, “a proposition <that x is F> is a conceptual truth if it belongs to the essence of ‘F’ that, necessarily, anything that satisfies ‘x’ also satisfies ‘F’” (Cuneo and Shafer-Landau, 2014, p. 410). Thus, a) is conceptually true because the essence of the concept *PRO TANTO WRONG* is such that anything that satisfies *RECREATIONAL SLAUGHTER OF A FELLOW PERSON* necessarily satisfies *PRO TANTO WRONG*.

This means that, according to Cuneo and Shafer-Landau, there are knowable substantive moral truths: the moral fixed points. And since moral concepts are non-natural, i.e. they are not reducible to natural concepts (Cuneo and Shafer-Landau, 2014, p. 402), these are non-natural moral truths.¹

2.3. Criticism of Cuneo and Shafer-Landau’s account

¹ One might wonder whether an account that only takes moral concepts to be non-natural is truly non-natural, for it seems that most naturalists can accept this as well.

One main criticism of Cuneo and Shafer-Landau's account is that it entails the conceptual deficiency of error theorists (Evers and Streumer, 2016; Ingram, 2015, 2018). Error theorists are cognitivists: they think that moral claims ascribe moral properties. However, error theorists also think that there are no moral properties. If there are no moral properties, no ascription of a moral property can be true. Thus, error theorists conclude, there are no positive moral truths. Since Cuneo and Shafer-Landau argue that there are positive conceptual moral truths, their account entails that error theorists are conceptually deficient (2014, p. 412). For if error theorists fully grasped the content of our moral concepts, they would also grasp the truth of the positive moral fixed points and reject the error theory.

Let us focus on Evers and Streumer's (2016) objection to this conclusion. They argue that the error theorist is not conceptually deficient by drawing a parallel with a theological claim. They argue that, just like Cuneo and Shafer-Landau's moral fixed points, the claim that

c) Benevolence is rewarded by God

bears at least some marks of conceptual truth (2016, p. 3). For instance, c) seems to have framework status. To engage in Christian theological thought, yet to deny c), shows that one uses an unchristian concept GOD.

But c) is not a conceptual truth. According to Cuneo and Shafer-Landau, c) would be conceptually true if the essence of the concept REWARDED BY GOD is such that anything that satisfies BENEVOLENCE necessarily satisfies REWARDED BY GOD. But it is not conceptually true that God exists. God might not exist. This means that benevolence might not be rewarded by God. c) therefore cannot be conceptually true (Evers and Streumer, 2016, p. 3). Evers and Streumer argue that the essences of the concepts in c) can at most show that it is conceptually true that,

c*) *If anything is rewarded by God, benevolence is rewarded by God* (Evers and Streumer 2016, p 4).

They then argue that the same holds for Cuneo and Shafer-Landau's moral fixed points, because it is "not a conceptual truth that wrongness is instantiated" (2016 p. 4). It is therefore not the case that anything that satisfies RECREATIONAL SLAUGHTER OF A FELLOW PERSON necessarily satisfies *PRO TANTO* WRONG. At most the essences of the concepts in a) make it conceptually true that,

a*) If anything is *pro tanto* wrong, it is *pro tanto* wrong to engage in the recreational slaughter of a fellow person.

The error theorist is not conceptually deficient if a*) is all that Cuneo and Shafer-Landau can show to be conceptually true. Since the error theorist denies that anything is in fact wrong, it is fine if our moral concepts entail that, if anything is wrong, it is wrong to recreationally slaughter someone (Evers and Streumer, 2016, p. 4).

This is a problem for Cuneo and Shafer-Landau. Their brand of non-naturalist realism is built up from the moral fixed points. If conditional truths like a*) are the only conceptual truths in the neighbourhood, their account must fail. For these conditional claims obviously fail to establish any form of moral realism.

2.4. Positive and Negative Moral Fixed Points

Note that Cuneo and Shafer-Landau's list of moral fixed points only includes *positive* moral claims (2014, p. 405): claims that, like a) and b), state that φ is morally wrong, or morally admirable, or that there is moral reason to φ , etc. But if Cuneo and Shafer-Landau were to focus on *negative* moral fixed points instead, they can avoid the criticism just mentioned. These negative moral claims state that φ is not morally wrong, or admirable, or that there is no moral reason to φ , etc.

Some negative moral claims seem to bear Cuneo and Shafer-Landau's marks of conceptual truth. For example:

d) It is in itself not morally good to clasp one's hands three times an hour (Foot, 2002, p. 118-119)²

and

e) Justice is not the number 4 (Parfit, 2011, p. 324-325).

Take d): 1) it seems that if d) is true it is necessarily true. And 2), it seems that if someone asserts the negation of d),

d') It *is* in itself morally good to clasp one's hands three times an hour,

they seem to have a strikingly different understanding of the concept MORAL GOODNESS. So, d) seems to have framework status. And, 3) someone who asserts d') seems quite bewildering. Finally, 4) it seems that we can *a priori* know d). For as Foot suggests, we should ask: why would hand clasping three times an hour be morally good in itself (Foot, 2002, p. 119)? Does it have anything to do with virtue, duty, charity, etc.? It does not. There is no conceptual connection between how something can be morally good in itself and hand clasping. The same seems to hold for e).

And we can consider d) and e) conceptually true because the essences of the concepts in d) and e) seem to entail their truth. That is, the essence of the concept MORAL GOODNESS is such that anything that satisfies CLASPING ONE'S HANDS THREE TIMES AN HOUR necessarily *cannot* satisfy MORALLY

² Although Foot does not use the term "in itself" she does reject objections to her claim that introduce a "special background" (2002, p. 118), such as: we can imagine a society where hand clasping is seen as a charitable act. I take this to mean Foot thinks this hand clasping cannot be morally good in itself.

GOOD IN ITSELF. And the essence of the concept JUSTICE is such that anything that satisfies NUMBER 4 necessarily *cannot* satisfy JUSTICE.

But how do these negative moral fixed points help avoid the problem of the conceptually deficient error theorist? Let us first look at a negative theological claim:

f) My cat is not God.

f) seems to be a negative theological conceptual truth. The essence of the concept GOD is such that it cannot apply to my cat. My cat, for instance, is not omnipresent, omnipotent, and sadly not immortal.³

Now note that the theological equivalent of an error theorist, the atheist, can accept that f) is conceptually true. For f) does not posit the existence of God. The conceptual truth f) only excludes the truth of certain positive theological claims: claims that entail that my cat *is* God.

The same holds for negative moral conceptual truths such as d) and e). The error theorist can accept that they are conceptually true, because they do not posit the existence and instantiation of moral properties. They only exclude the truth of certain positive moral claims: claims that entail that it *is* in itself morally good to clasp ones hands three times an hour, and claims that entail that justice *is* the number 4. Claims d) and e) being conceptually true therefore does not entail that the error theorist is conceptually deficient.

Now, one might think that we need to know some positive moral fixed points in order to know the negative ones. If so, the error theorist cannot accept the negative moral fixed points after all. But although we do need some positive knowledge of moral concepts to grasp the negative moral fixed points, this knowledge need not come from positive moral fixed points.

That we need some positive knowledge of moral concepts is evident from the reasons why d) seems conceptually true. We need some positive knowledge of the moral concepts VIRTUE, DUTY,

³ As evidenced by her passing while I worked on this thesis. Rest in peace, Mies.

CHARITY, etc. to see that hand clasping has nothing to do with them. But it is unclear why we would need positive moral fixed points to obtain this knowledge. For example, assume we have a proper understanding of the NUMBER 4. To then know that the number 4 cannot be justice, we only need to know that justice is somehow linked to interests or harm. But it does not seem like we need positive conceptual moral truths to have this very limited knowledge of JUSTICE.⁴

2.5. Shifting Focus

One may represent this in the following way:

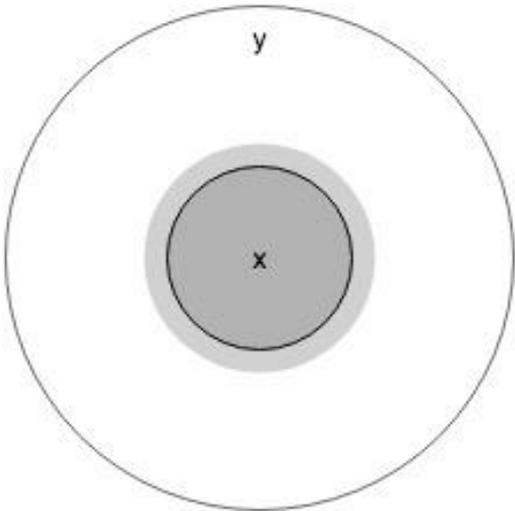


Figure 1.

Figure 1 depicts all possible moral beliefs. Take **x** to represent those moral beliefs whose truth is not settled by our concepts. It will include beliefs such as “lying is wrong” and “civil disobedience can be morally permissible”. Take **y** to represent moral beliefs whose truth is settled by our moral concepts. It will include beliefs such as “a duck is not charity” and “sneezing loudly is in itself not virtuous”. Take the grey space between **x** and **y** to represent moral beliefs of which it is disputed whether our moral concepts settle the matter.⁵ It is less clear what this might include, but the belief that all humans have equal moral worth might be a candidate.

⁴ It is also unlikely that this limited knowledge itself would allow us to grasp any positive moral fixed points.

⁵ Note that figure 1 is not necessarily an accurate depiction of the size of these three domains: claims that are clearly conceptually true (**y**), claims that are clearly not conceptually true (**x**), and claims for which this is disputed (grey space).

Given the seemingly sound criticism by Evers and Streumer (2016), a), b), and all other positive moral fixed points actually fall within x instead of y . For their truth is not settled by our concepts. y can therefore only encompass negative moral fixed points.⁶ These are the only moral truths that our moral concepts do settle.

2.6. Do The Negative Moral Fixed Points Show Enough?

Sections 2.4 and 2.5 suggest that Cuneo and Shafer-Landau should focus on negative moral fixed points. But many might be sceptical about the prospects of this endeavour. For the negative moral fixed points are weaker than the positive ones. That is, they entail less. So, there is this looming question: do these negative moral fixed points show enough to tell a complete non-naturalist tale?

No. If non-naturalists only manage to establish the negative moral fixed points it would be a major blow to the aspirations of non-naturalism. Non-naturalists should therefore supplement their account of the negative moral fixed points with an explanation of how we can come to have true positive moral beliefs.⁷ But this task is not unique to this negative account. Cuneo and Shafer-Landau also need to show how we can come to have true positive moral beliefs that are not conceptually true. So, going negative on the moral fixed points does not come at the cost of any new explanatory burdens.

This, however, might not dispel the worry completely. For Cuneo and Shafer-Landau will have more resources to draw from for this supplementation if they can establish positive moral fixed points instead of just negative ones. That is, because the positive moral fixed points entail more, any account that includes them has more resources to tell a complete non-naturalist tale than any account that includes only negative moral fixed points. Thus, although shifting our focus to negative moral fixed points does not create any new problems, it does limit the resources for answering

Figure 1 is simply a way to visualise these domains. In reality y , x , and the grey space might be a lot bigger or smaller than depicted in figure 1.

⁶ This is not completely true. Conditional moral fixed points, like a^* , can be conceptually true as well, regardless of whether they are negative or positive.

⁷ Although a supplemented negative account does entail that the error theorist is wrong, it does not entail that she is conceptually deficient.

existing problems. Thus, to shift focus is to move down on the scale of explanatory power.

However, here we should return to a theological example. Although the existence of God would be a tremendous explanatory resource, we do not think that agnostics or atheists have moved down on the scale of explanatory power. We view the theist as making an unwarranted move up that scale. This is because this explanatory resource only becomes available when we can establish that God exists, which it seems we cannot. The same holds for Cuneo and Shafer-Landau's account. The positive moral fixed points can explain more than the negative ones, but we are not justified in asserting the positive moral fixed points. So, we should not view going negative as a move down the scale of explanatory power, we should view Cuneo and Shafer-Landau as having made an unwarranted move up that scale.

2.7. Retaining benefits

Besides, it is not as if negative moral fixed points are explanatorily inert. The explanatory power of the negative moral fixed points is comparable to that of the positive ones on certain key issues. That is, Cuneo and Shafer-Landau claim that their positive account can deal with three important objections to non-naturalist realism (2014, p. 422-431). And in this section, I will argue that a negative account can deal with these objections in a similar fashion.

2.7.1 Moral Disagreement

The first objection that Cuneo and Shafer-Landau discuss is the argument from disagreement (2014, p. 422-424). The argument is supposed to undermine our confidence in our moral judgments by pointing out the ubiquity of moral disagreement.⁸ The anti-realist can argue that the best explanation of this widespread moral disagreement is that there are no moral truths. For in a world without

⁸ This might be actual disagreement, or disagreement between more idealized agents.

moral truths we would expect to see widespread moral disagreement, since moral thought would be an exercise of projecting attitudes onto the world rather than accessing moral truths.

Cuneo and Shafer-Landau point out that there is no widespread moral disagreement about the positive moral fixed points. The argument from disagreement therefore does not apply to their account (Cuneo and Shafer-Landau, 2014, p. 423). This means that the argument from disagreement does not undermine our confidence in all our moral judgments, because it does not undermine our confidence in the positive moral fixed points.

We can say the same about negative moral fixed points. There is little disagreement on whether clasping one's hands three times an hour is in itself morally good. Possibly even less than on the positive moral fixed points, since the error theorist can accept the negative ones. The argument from disagreement therefore does not apply to the negative moral fixed points. This means that the argument from disagreement does not undermine our confidence in all our moral judgments, because it does not undermine our confidence in the negative moral fixed points.

2.7.2 Remarkable Coincidence

The second objection that Cuneo and Shafer-Landau discuss is the remarkable coincidence argument (2014, p. 424-428). Cuneo and Shafer-Landau interpret this argument, as for instance put forward by Street (2006, p. 121-122), as follows: there are indefinitely many moral systems, but we do not observe great variance in the moral systems that people endorse. This is because evolutionary forces influence our moral capacities and judgments, and there are only so many moral systems that enhance fitness. But this means that our moral systems track fitness, not truth. It would therefore be a remarkable coincidence if our moral system encompasses a significant number of moral truths.

Cuneo and Shafer-Landau dispute the claim that there are endless possible moral systems (2014, p. 426-428). This is because for any moral system to count as such it must encompass the positive moral fixed points. The set of possible moral systems is therefore limited, meaning the first premise of the remarkable coincidence argument is false.

Negative moral fixed points can do a similar job, because the negative moral fixed points also limit the set of possible moral systems. For any system that encompasses the negation of any negative moral fixed point cannot count as a moral system. So a system that encompasses, for instance,

d') It *is* in itself morally good to clasp one's hands three times an hour

or

e') Justice *is* the number 4

is not a moral system. Find enough negative moral fixed points and the set of possible moral systems starts to dwindle rapidly. And as it happens they are easy to find. The set of possible moral systems is therefore limited, meaning the first premise of the remarkable coincidence argument is false.

Now, one might object that this is not entirely fair. This response to the remarkable coincidence argument targets the first premise of the argument, but this first premise is put formulated too strongly and is therefore an easy target. Instead of there needing there to be an infinite amount of possible moral systems, for the argument to work there would only need to be a large number of possible moral systems. How large a number? Large enough for it to be remarkable for us to be right. So, instead of arguing that the moral fixed points limit the set of possible moral systems to a number smaller than infinity, Cuneo and Shafer-Landau need to show they limit the set of possible moral systems to a number smaller than the number at which it would be remarkable for us to get it right.

This is where the explanatory power of positive and negative moral fixed points might come apart. It is clear that negative moral fixed points entail less than positive ones. This is a virtue when it comes to the matter of the conceptual competence of error theorist, but might be a vice when it

comes to the issue of sufficiently limiting the set of possible moral systems. That being said, it is also not clear whether the positive moral fixed points can rise to the occasion. This depends on how restrictive they are and the number at which we think getting it right would no longer be remarkable, both of which do not have clear-cut answers.

Besides, it is important to focus on the goal of this paper: to show that Cuneo and Shafer-Landau should defend negative instead of positive moral fixed points. Doing so requires arguing that the negative moral fixed points have the same benefits as positive ones. This definitely seems to hold Cuneo and Shafer-Landau's own interpretation of the remarkable coincidence argument. Since that interpretation is the bar that Cuneo and Shafer-Landau set, I will work with it and take the negative moral fixed points to reach it.

2.7.3 The Humean challenge

The third objection that Cuneo and Shafer-Landau discuss is the Humean challenge (2014, p. 428-431).

The Humean challenge relies on three claims:

SUPERVENIENCE: No metaphysically possible world that is identical to a second world in all base respects can be different from the second world in its ethical respects.

BRUTE CONNECTION: The nonnaturalist must take the supervenience of the ethical properties on the base properties to involve a brute necessary connection between discontinuous properties.

And:

MODEST HUMEAN: Commitment to brute necessary connections between discontinuous properties counts significantly against a view.

(Cuneo and Shafer-Landau, 2014, p. 428)

According to the Humean challenge non-naturalism comes at a significant cost. If it is to respect SUPERVENIENCE it has to posit brute (i.e. unexplained) necessary connections between the non-

natural moral properties and natural properties.

Cuneo and Shafer-Landau, however, are uncommitted as to the nature of moral properties (2014, p. 402). Their account is only non-naturalist about moral concepts and is compatible with the claim that moral properties are natural properties. If this is the case, there is no need to posit brute connections, since only theories that take moral properties to be non-natural are committed to such connections (Cuneo and Shafer-Landau, p. 429). A proponent of the negative account can likewise remain uncommitted on the nature of moral properties. Negative moral fixed points do not entail the existence of non-natural moral properties. The Humean challenge therefore does not affect the negative account either.

2.8. Conclusion

Cuneo and Shafer-Landau (2014) argue that there are positive conceptual moral truths: the moral fixed points. This entails the conceptual deficiency of error theorists. Focusing on negative conceptual moral truths avoids this problem. This move does not add any new explanatory burdens and retains the benefits of Cuneo and Shafer-Landau's account. We should therefore go negative on the moral fixed points.

3. Expressivism and Moral Irrelevance

3.1. Introduction

This paper identifies a new challenge for metaethical expressivism. It argues that it is unclear whether expressivism can explain why some irrelevance claims, i.e. claims according to which a certain descriptive feature is not or cannot be morally relevant, seem conceptually true. In section two, I will give several examples of irrelevance claims. In section three, I will argue that some irrelevance claims bear marks of conceptual truth. And I will explain why I focus on moral irrelevance claims in particular. If these claims are indeed conceptually true, metaethical theories have to explain how this irrelevance is grounded in the nature of moral concepts. In section four, I will argue that moral realists can meet this explanatory challenge, but that it is unclear how expressivists can meet it. This chapter does not argue that it is impossible for expressivists to meet this challenge. Rather, it argues that several initially plausible expressivist responses fail, which suggests that meeting this challenge will require a novel response from the expressivist.

3.2. Irrelevance Claims

In this section, I'll discuss three examples of irrelevance claims, by Singer (1975, 1989), Thomson (1971), and Parfit (1984). As will become clear, irrelevance claims are not always made explicitly. Singer's claim, however, is and will therefore receive most attention.

In brief, Singer argues that we should give non-human animals (hereafter "animals") the same moral consideration as marginal humans: humans that lack certain capacities that typical adult humans possess, e.g. severely mentally handicapped people, or (brain-damaged) infants. For the only consistent way to deny this is to take *mere* species membership to be morally relevant. And mere species membership cannot be morally relevant (Singer, 1975, p. 21).

In more detail, Singer notes that we treat humans and animals differently. We regularly subject animals to treatment that we would never inflict on a human. Think, for example, of the

prevalence of the experimentation on and the consumption of animals (Singer, 1975, Ch. 2, 3).¹ This shows that we generally do not care much about the interests of animals. According to Singer, this is down to an “attitude of bias toward the interests of members of one’s own species and against those of members of other species” (Singer, 1975, p. 7). He calls this bias “speciesism”, as akin to racism or sexism.²

However, many believe that this moral divide is justified by the fact that humans have and animals lack certain capacities, such as the capacity of reason. But Singer notes that marginal humans are no more rational than apes, cows, or pigs (1975, p. 20, p. 79; 1989, p. 80). This means that if we truly think that the capacity of reason demarcates the moral community, i.e. the group of beings that deserve moral consideration, marginal humans should fall on the same side as animals.³ This leaves us three options: treat marginal humans like animals, treat animals like marginal humans, or choose another criterion to demarcate the moral community.

Singer swiftly dismisses the third option. Other criteria for demarcation face the same problem. For one can plausibly argue that marginal humans can at most rival animals in any other capacity (autonomy, emotional maturity, etc.) as well (Singer, 1975, p. 22; 1989, p. 83-84). This leaves us the other two options: treating marginal humans like animals, or treating animals like marginal humans. Or, put differently, we either push marginal humans out of the moral community or pull animals into it.

Some want to resist this conclusion. They often point out that marginal humans are still human, and argue that this provides marginal humans a higher moral status.⁴ But, according to Singer, *mere* species membership cannot be morally relevant (1975, p. 21). One can therefore only

¹ Of course humans have been subjected to similar horrific treatment. But that is seen as a moral travesty by most. And if it is not, like in Nazi Germany, that seems to be at least in part due to the dehumanization of those persecuted. Being seen as human is generally taken to protect someone against this kind of treatment.

² This kinship extends beyond mere name. Singer argues that embracing speciesism undermines our ability to denounce racism and sexism, because the reasons that supposedly justify speciesism also justify racism and sexism (1975, p. 9; 2016, p. 31-33).

³ Some might not view the moral community as binary (either you are in or you are out) but as, for instance, a tiered system of significance. Singer’s point should then be taken as saying that animals and marginal humans belong in the same tier.

⁴ Claims about the moral significance of being human need not be this blatant. One might instead use “high-sounding phrases like “the intrinsic dignity of the human individual”” (Singer, 1989, p. 82-83) to try to vindicate the moral divide between humans and animals. But if this is not just a claim about a mere difference in species, it only moves the bump under the rug. For in light of what capacity do only humans possess this dignity (Singer, 1989, p. 83)?

avoid the choice to either push marginal humans out or pull animals into the moral community on pain of moral irrelevance. That is, by grounding one's moral judgment on a morally irrelevant ground. For one can only avoid this choice by being what I shall call a *strict-speciesist*: someone who claims that a *mere* difference in species is morally relevant.⁵

Note that strict-speciesism is an extreme view. Singer likens speciesism more generally to racism (1975, p. 9; 2016, p. 31-33), but the strict-speciesist is unlike any racist we are familiar with. Granted, both racists and strict-speciesists postulate a moral divide between groups. But the average racist tries to justify this with a misguided judgment about racial differences in what we would generally think of as morally relevant features, such as intelligence, autonomy, etc. Thus, morally speaking the racist is not – however odd it might sound – concerned with *mere* race.⁶ In contrast, the strict-speciesist is morally concerned with *mere* species membership.

To illustrate this even further, imagine, as Singer asks us to (2001, p. *xvi-xvii*), that we discover extra-terrestrial life. Also imagine that these aliens are as intelligent, emotionally mature, autonomous, susceptible to suffering and happiness, etc. as we are. They are like us. But because they have followed a different evolutionary path, they belong to a different species. The strict-speciesist thinks this mere difference in species is morally significant. She might therefore believe it is morally permissible to exploit these aliens for our enjoyment, but that exploiting other humans for our enjoyment is obviously reprehensible. Thus, “strict-speciesist” is a technical term that picks out those who, for instance in response to Singer's marginal cases argument, latch onto *mere* species membership as morally relevant.⁷

Singer's marginal cases argument is a consistency argument. It claims that, if we think marginal humans deserve a certain amount of moral consideration, then animals deserve the same.

⁵ Singer uses the term “pure speciesism” for this kind of speciesism (1975, p. 21). I opt for the term “strict” because the term “pure” arguably bears some positive normative connotations that do not appropriately apply here.

⁶ This of course takes the racist at face value. However, even if the racist's concern for intelligence is a result of the rationalization of a concern for *mere* race, the fact that she rationalizes this shows that the stricter position is quite extreme. For it indicates that simply accepting that *mere* race matters is difficult.

⁷ This means one can be a speciesist, without being a strict-speciesist. Just like one can be sexist without believing that a mere difference in sex is morally relevant. For instance, by valuing what stereotypically are considered “masculine” traits more than “feminine” ones.

This argument only works if there is no morally relevant difference between these groups. Singer's irrelevance claim that mere species membership cannot be morally relevant counters the strict-speciesist's claim that it is, and is therefore pivotal to his consistency argument.

Other consistency arguments rely on similar irrelevance claims, although they are often left implicit. Take an applied ethics classic: *A Defense of Abortion* by J.J. Thomson (1971). In this paper, Thomson uses the thought experiment of the ill violinist to show that even if we assume that the fetus is a person from the moment of conception, abortion can still be morally permissible (1971, p. 48-49). Thomson asks us to imagine involuntarily being medically hooked up to an ill violinist for the nine months. And to further imagine that the violinist's life depends on this arrangement. Thomson thinks it is outrageous to suggest that we may not disconnect ourselves from the violinist. This outrage, she argues, should affect our view of abortion (1971, p. 49).

Thomson argues for consistency in our treatment of a fetus and the violinist. This argument only works if there are no morally relevant differences between these two. Although Thomson does not make any explicit irrelevance claims, she should hold that differences between a fetus and the violinist are morally irrelevant. She is therefore committed to, for example, the irrelevance claim that: the location of the hostess (the fetus and violinist) relative to the host (the pregnant person and the person hooked up to the violinist) cannot be morally relevant. That is, whether or not the hostess is inside us or in a bed next to us, cannot matter for whether ending their dependence on us is morally permissible.

Irrelevance claims also figure in other areas of normative inquiry. For example, in *Reasons and Persons* (1984) Parfit argues that it is no good reason prefer a painful surgery on Tuesday to a mildly irritating surgery on Wednesday, just because it falls on Tuesday (Parfit, 1984, p. 124). Parfit's irrelevance claim is that just the day of the week cannot affect the desirability of pains and pleasures.

One might find certain consistency arguments contentious. One might think that Singer overlooks a morally relevant difference between (brain-damaged) infants and pigs. However, this chapter focuses on irrelevance claims, not on the consistency arguments they figure in. It focuses on

claims such as: *mere* species membership cannot be morally irrelevant. This claim seems uncontroversial. In what follows, I will take Singer's claim as the main example.

3.3. Conceptual Irrelevance

There are at least two ways to interpret these irrelevance claims. On a normative reading, irrelevance claims are simply claims about what does not make normative difference.⁸ On a conceptual reading, we instead interpret irrelevance claims as saying that something cannot make a moral difference. This 'cannot be' formulation is indicative of a stronger claim. The natural reading is that it is a matter of necessity that the descriptive feature φ does not make a moral difference. This suggests that our moral concepts do not allow for moral theories that take φ to be morally relevant. That is, it suggests that our moral concepts bar φ from being morally relevant. I will argue that this conceptual reading seems appropriate for some irrelevance claims, because some of them seem conceptually true. I will use Foot's (2002) and Cuneo and Shafer-Landau's (2014) work on conceptual moral truths to argue for this. But before I do, note that I will simply assume that their analyses are correct. This does mean that all that follows is conditional on whether their account of what makes a claim a conceptual moral truth is right.

Foot argues that not just anything can be morally good. The content of the concept MORAL GOODNESS excludes certain things from being morally good. So, it is knowable simply through an adequate understanding of this concept that, for instance, to clasp one's hands three times an hour is not in itself morally good (Foot, 2002, p. 118-119). For if someone claims it *is*, we should ask: What do you mean? Does it display a virtue? Does it alleviate a duty? Can we consider it an act of charity? The answer must be "No" (Foot, 2002, p. 119).⁹ Clearly, just clasping your hands three times an hour

⁸ This normative reading might seem obviously wrong for Singer's Irrelevance claim. For Kant postulates a strict moral divide between humans and animals. However, Kant does not use a *mere* difference in species to demarcate the moral community. He claims that only rational beings fall within the moral community (1785, 4:389; 4:408; 4:411-4:412; 4:425), and he believes that rationality sets humans apart from all other animals (1785, 4:428). Thus, Kant's ethics is compatible with the normative reading of Singer's argument.

⁹ That is, if we do not inject some extra information into the story (Foot, 2002, p. 118). For example, that this concerns a different society with different moral concepts that entail that these actions are virtuous or charitable. But in changing the

has nothing to do with moral goodness. The content of the concept MORAL GOODNESS bars this hand clasping from being morally good. This suggests that it is conceptually true that clasping one's hands three times an hour in itself is not morally good.

Building on Foot's work, Cuneo and Shafer-Landau argue that certain moral claims, the moral fixed points, are conceptually true, e.g. "It is *pro tanto* wrong to engage in the recreational slaughter of a fellow person" (Cuneo and Shafer-Landau, 2014, p. 405).¹⁰ This claim supposedly bears the following marks of conceptual truth (Cuneo and Shafer-Landau, 2014, p. 407-408): One, it is *a priori* knowable. Two, if true, the claim is necessarily true. Three, the claim has "framework status, fixing the boundaries as to what counts as a type of subject matter". This means that if someone denies this claim it is reasonable to doubt whether she uses the same concepts we do. Four, denying the claim bewilders conceptually competent people.

Judging by these features of conceptual truths, Singer's irrelevance claim seems conceptually true. It seems we would answer Foot's questions, "Is it to do with virtue, duty, charity?" with a resounding "No", when asked about the moral relevance of mere species.¹¹ Mere species membership has nothing to do with the criteria for the application of moral concepts. This suggests that the content of our moral concepts bars mere species membership from being morally relevant.

As for Cuneo and Shafer-Landau's marks of conceptual truth: as Foot's work suggests, it seems like we can know, simply through understanding the concepts involved, that mere species membership is morally irrelevant. Given that Singer's claim seems *a priori* knowable, there is also some reason to think it is necessarily true. And Singer's irrelevance claim seems to have some framework status. That is, it seems reasonable to doubt whether the strict-speciesist uses the same moral concepts we do. The strict-speciesist also seems bewildering. Remember that she accepts

content of the concept MORAL GOODNESS this reply can no longer give us insight into the concept MORAL GOODNESS as we know it. There is of course the further question, which I will not address here, of which concept MORAL GOODNESS we should use. For an in-depth discussion of this issue, see Eklund (2017).

¹⁰ They see their project as an elaboration of Foot's work (Cuneo and Shafer-Landau, 2014, p. 406).

¹¹ Note that there is a difference between something not being morally good and it being entirely morally irrelevant. For besides the concept MORAL GOODNESS, other moral concepts like RIGHT, WRONG, PERMISSIBLE, IMPERMISSIBLE, etc. bear on the question of moral relevance. So, if we ask Foot's questions about moral irrelevance, we must imagine asking them not just about MORAL GOODNESS, but about these other moral concepts as well.

Singer's contention that (some) marginal humans are equal to cows and pigs in every respect, except for them being different species. But contrary to Singer, she also thinks that this sole difference in species is morally relevant, such that exploiting cows and pigs is permissible because of it. This seems quite bewildering.

If this does not seem bewildering to you, this might be because a difference in species generally correlates with a difference in morally relevant capacities (intelligence, autonomy, etc.). A way to get around this is to return to Singer's alien example. Remember that these aliens are just like us (equally intelligent, autonomous, etc.), except that they are a different species. Imagine someone who claims that it is morally permissible to forcefully rear, breed, and exploit these aliens for our enjoyment, but that doing so with humans is absolutely forbidden. Imagine further that she claims that the *mere* difference in species vindicates this difference in moral appraisal. This surely is bewildering.

Thus, Singer's claim that mere species membership is morally irrelevant seems to bear several of Foot's and Cuneo and Shafer-Landau's marks of conceptual truth. But one might wonder: why go through all of this effort to establish this about Singer's claim? Why not just use Cuneo and Shafer-Landau's moral fixed points? This is because Singer's claim is *negative*, while Cuneo and Shafer-Landau's moral fixed points are *positive*. Cuneo and Shafer-Landau's moral fixed points are about things that *must be* morally wrong, admirable, just, etc., such as the wrongness of the recreational slaughter of a fellow person. Singer's irrelevance claim is *negative* in that it only says that mere species membership *cannot be* morally relevant. This claim merely excludes the truth of certain moral claims. It does not, like Cuneo and Shafer-Landau's positive moral fixed points, ensure the truth of certain moral claims.¹² This makes Singer's claim less controversial than Cuneo and Shafer-Landau's claims.

Some might, however, think that Singer's claim is more controversial in other ways. That is, almost everyone seems to be against the recreational slaughter of fellow people, but there are a lot

¹² But Foot's claim is negative too, so why not use her claim? The thought behind using Singer's claim is to show that the irrelevance claims actually play a role in applied ethics.

of people that see no problem with speciesism. As mentioned when discussing Singer, the use and consumption of animals is widespread. Here I have to once again press the point that we are not considering the claim whether membership of species is morally relevant, but whether *mere* species membership is morally relevant. And, when pressed to consider this, it seems to me quite uncontroversial to think that it is morally irrelevant.

So, Singer's irrelevance claim seems to bear Foot's and Cuneo and Shafer-Landau's marks of conceptual truth. In section 4.5, I will discuss whether, besides Singer's claim being conceptually true, there are any alternative explanations for why Singer's claim bears these marks. Bracketing that for now, we can take Singer's claim having these marks as making it plausible that this claim is conceptually true. As suggested, a defence of why these marks are suggestive of conceptual truth is needed to affirm this more strongly. I will not do that here, and I defer to Foot (2002) and Cuneo and Shafer-Landau (2014) for this. Assuming that what they say is correct, Singer's irrelevance claim seems conceptually true.

3.4. Metaethics: Realism and Expressivism

The conceptual truth of Singer's claim gives rise to an explanatory burden. Metaethical theories need to explain how it can be conceptually true. That is, they need to explain why the content of our moral concepts is such that it bars mere species membership from grounding moral claims. Or, alternatively, they need to explain why Singer's claim *seems* conceptually true, even though it is not.

I argue that, assuming we have solved all other relevant problems in metaethics, moral realism can meet this explanatory burden while it remains unclear whether expressivism can meet it. I will first give the realist's solution and then I will go through some potential expressivist solutions and argue that they fail.

3.4.1. Realist Solution

Moral realists are committed to at least two theses. The first is cognitivism (Shafer-Landau, 2003, p. 17): moral claims express beliefs. And these moral beliefs involve an ascription of moral properties to an act. The second is a specific metaphysical view: “there are moral facts involving moral entities, relations, and properties that do not consist in what anyone’s attitudes are” (Finlay, 2007, p. 821). That is, realists hold that there are mind-independent moral properties, relations, and entities.

These two theses together suggest that we use moral concepts to pick out moral properties. Thus, when we say, “stealing is morally wrong” we use the moral concept *WRONG* to pick out the moral property of wrongness possessed by the act of stealing. In these specific ways moral concepts and properties resemble descriptive concepts and properties.¹³ To claim, for instance, that something is wooden is to express the belief that it has the property of being made of wood. We take this property to exist mind-independently and we use our concept *WOODEN* to pick it out.

The realist can argue that this metaphysical worldview and our use of concepts suggest a particular mechanism that (in part) determines the content of our concepts. And she can further argue that it follows from this mechanism that our moral concepts bar certain descriptive features from grounding certain claims.

Let us first look at a descriptive example. Just like the strict-speciesist’s claim, the claim “this table is made of wood merely because it has three legs” seems conceptually confused. The content of the concept *WOODEN* bars the mere three-leggedness of a table from grounding its woodenness. To explain this, the realist can argue as follows: tables and wooden things exist. We learn, from encounters with these properties, that the property of being made of wood is unaffected by certain descriptive features. Since we use concepts to pick out properties, this information should affect our understanding of the concept *WOODEN*. This in turn allows us to know that certain descriptive features cannot affect the instantiation of woodenness from the content of the concept *WOODEN* alone. Having three legs is such a condition. For the world is such that the three-leggedness of a table is never the reason why the table is wooden. We observe that three-legged tables can be made from

¹³ In many other respects moral and descriptive concepts and properties might be vastly different.

all kinds of materials. And our concepts have come to reflect this fact. That is why mere three-leggedness is conceptually irrelevant for questions about whether a table is made of wood or not.

The moral realist can say the same about Singer's irrelevance claim: moral properties exist. We encounter them daily. From these encounters we learn that moral properties like goodness, badness, etc. are unaffected by certain descriptive features. Since we use moral concepts to pick out these properties, this information should affect our understanding of the moral concepts like GOOD, BAD, etc. This, of course, need not provide us with infallible knowledge of what moral properties are instantiated under which descriptive features. But it does mean that we can know that certain descriptive features cannot affect the instantiation of moral properties from the content of our moral concepts alone. For instance, it allows us to rule out that the mere shape of the moon grounds the virtuousness of defending one's *polis* in war. It allows us to rule out that the mere colour of a child's eyes grounds a duty to save it from drowning. And it allows us to rule out that mere species membership grounds the moral permissibility of exploiting animals.

Of course, the realist will have to say more about how information about the world affects our understanding of our concepts for this solution to be convincing. Exactly how this works depends on our conception of concepts. There are two dominant views on this: concepts as mental representations and concepts as abstract objects (Margolis and Laurance, 2007, p. 561).

On the first view, concepts are mental representations, which are the constituents of propositional attitudes (Margolis and Laurance, 2007, p. 562). My belief that the earth revolves around the sun is constituted by the more basic mental representations: earth, movement, sun, etc. On this view, these more basic mental representations are concepts.

In the case of, for instance, our concept SUN, the mental representation represents the actual sun. Now imagine our early ancestors' mental representation of the sun. It surely would have been vastly different from ours. For example, (some aspects of) divinity would have been central to their representation of the sun, while this plays no part in ours. It seems natural to think that this change in the mental representation of the sun is, at least partly, driven by our growing understanding of

what the sun actually is.¹⁴ For being sensitive to new information about that which is represented seems essential to being a representation. There is no reason for the realist to think this does not hold for the mental representations of moral properties as well. New information about these properties will therefore affect our mental representation of them. This is how our understanding of our moral concepts can be shaped by information about moral properties.

On the second view, concepts are abstract mind-independent objects that are the constituents of propositions (Margolis and Laurance, 2007, p. 562). It is less clear how information about moral properties can affect our understanding of moral concepts on this view. However, two things are clear: one, the realist can claim that we can grasp abstract objects, and that we can therefore grasp the boundaries of our moral concepts, which allows us to recognize certain moral claims as conceptually true. And two, the expressivist cannot claim this. It would be odd for the expressivist to suggest that our understanding of moral concepts depends on our grasp of abstract objects. On the expressivist account, it seems that the meanings of our moral terms will have to be determined by our mental attitudes.

So, on the mental representation view there is a clear realist answer to the question how information about moral properties affects our understanding of our moral concepts. On the abstract objects view this is less clear, but it is clear that the expressivist cannot hold this view.

However, there seems to be a fundamental problem for both these views: they require an explanation of how we can epistemically access moral properties. This problem is especially pressing for the non-naturalist (Mackie, 1977, p. 38-39). However, this problem is irrelevant for my purposes in this paper. As mentioned, this paper aims to argue that, when all other problems for moral realism and expressivism are solved, the realist can explain why our moral concepts bar mere species membership from grounding moral claims, but that it remains unclear whether the expressivist can.

¹⁴ One might think that mental representations are stable and that it is just the things we associate the mental representations with that change. On this view, we have the same mental representation of the sun as our ancestors. What changed is that we just do not associate it with the divine anymore. This need not undermine the realist's explanation. She can argue that on this view new information about the sun changes the conceptual ties we take the concept SUN to have. It therefore remains the case that information about the world affects our understanding of our concepts.

In this section, I have argued for the first half of this aim. In the following sections, I will argue for the second half.

3.4.2. First Expressivist Answer: Attitudinal Interpretation

Expressivists are non-cognitivists: they believe that moral claims express moral attitudes, and that moral sentences derive their meaning from these attitudes. Thus, “stealing is wrong” expresses and gets its meaning from the attitude of, say, disapproval of stealing (Blackburn, 1984, 1993), or having a plan not to steal (Gibbard, 2003). This means that expressivists do not explain what it is for stealing to be wrong, they explain what it is to think that stealing is wrong.

So, to explain why mere species membership is conceptually morally irrelevant the expressivist might explain what it is to think this. She can argue that to think this is to disapprove of changing one's moral attitude about harming ϕ just because of ϕ 's species, or to plan not to change one's moral attitude about harming ϕ just because of ϕ 's species.

But this alone does not solve the problem at hand. The problem is not just to understand what it is to think that something is morally irrelevant. The problem is to explain why some claims are *conceptually* morally irrelevant or why they seem to be even though they are not. For instance, we need to know why someone who approves of changing one's moral attitude about harming ϕ just because of ϕ 's species bewilders us.

3.4.3. Second Expressivist Answer: Supervenience

The expressivist might argue that the conceptual truth of these irrelevance claims is grounded in the conceptual truth of a different claim about morality: the supervenience claim. This claim states that, if two events are descriptively identical, they must be morally identical as well. And expressivists such as Blackburn have argued that this claim is conceptually true (1993, p. 137).

Now, to be clear, supervenience itself cannot explain why we should treat infants and pigs as equals. There are obvious descriptive differences between the two. But the explanation of our observance of supervenience might explain it. This explanation ensures the consistency of our moral judgment when we consider two descriptively identical situations. And given its ability to ensure that, it might also be able to ensure the consistency of moral judgments that only differ in morally irrelevant ways.

Now, to explain why we observe supervenience the expressivist obviously cannot claim that that is simply how moral properties behave. Blackburn (1984, p. 186; 1993, p. 143-144) and Gibbard (2003, p. 56) argue that the role of the supervenience claim is to secure the action-guiding function of moral judgements. They argue that if we do not assume supervenience two descriptively identical acts could differ in moral status. That is, one and the same act might be both good and bad, which undermines the action-guiding function of moral judgments.

However, it seems that as long as one can clearly distinguish between humans and animals, the strict-speciesist's theory is a clear and reliable guide for action. So it does not seem plausible that we accept the moral irrelevance of mere species membership because we want to avoid undermining the action-guiding function of moral judgements. It is therefore unclear how the conceptual truth of the supervenience claim would help to explain why Singer's irrelevance claim is conceptually true or appears to be.

3.4.4. Third Expressivist Answer: Evolutionary Explanation

Instead, the expressivist could argue that there is an evolutionary explanation of why Singer's irrelevance claim seems conceptually true. Expressivists already argue that we developed a sense of morality because this was (evolutionarily) beneficial to us (Blackburn, 1984, p. 186, 1993, p. 137; Gibbard, 1990, p. 26-30). They might therefore also argue that evolutionary forces affect the content of our moral concepts. And they could further argue that they do so in a way that makes Singer's irrelevance claim conceptually true.

The expressivist might argue as follows: take the case of Foot's hand clasper. It would seem evolutionarily beneficial for us not to participate in idle hand clasping. It does not enhance our fitness and is potentially an enormous time sink. To believe that this hand clasping is morally good therefore seems detrimental to the propagation of our genes. Thus, natural selection will favour those that do not believe that this is morally good.

But this can only be part of the expressivist's answer. She has only shown why we might believe that hand clasping is not morally good. She has yet to show why it is or seems conceptually true that mere hand clasping is not morally good. It is far less clear how that would work. The expressivist can of course rightly claim that natural selection favours those whose moral concepts ensure that hand clasping cannot be morally good for the same reasons that it favours those who merely believe this. But a revision of one's moral concepts seems far more laborious than simply adopting a belief. And since both seem capable of preventing the detrimental behaviour of hand clasping, it is unclear why evolution would have had us take the laborious route.

But even if we assume that it did, it does not seem like the same evolutionary story holds for Singer's irrelevance claim. That is, it seems that natural selection would favour those with strict-speciesist moral concepts. Being able to exploit animals for our own benefit, free from any moral qualms, seems highly beneficial to our survival, and thus the propagation of our genes.

The expressivist can respond in two ways: first, she can argue that Singer's irrelevance claim is an outlier in that it might seem conceptually true, but it in fact is not. Second, she can tell a more sophisticated evolutionary tale that can explain why Singer's irrelevance claim is conceptually true. For the remainder of this section I will focus on the second response. I will discuss the first response in the next section.

The expressivist's more sophisticated evolutionary tale might go as follows: evolution does not just produce capacities that are fitness enhancing in every respect. It produces capacities that are fitness enhancing in general. This means that evolved capacities can have effects that reduce our fitness. This is what happened in the case of Singer's irrelevance claim. Evolution provided us with

many rational capacities that are generally fitness enhancing. But when we use those capacities to reflect on morality, they lead us to believe that mere species membership is conceptually morally irrelevant, which is not fitness enhancing.

The expressivist might even use Singer's own work to fill in the details of this story. In *The Expanding Circle* (2011) Singer argues that an ethical system can emerge through evolutionary forces (p. 91-92), and that our rational capacity allows us to apply that ethical system to others outside of the community it emerges in (Singer, 2011, p. 87-124). That is, reason allows us to expand the circle of moral concern. It does so by having us realize that, if we want to justify our behaviour to others properly, we will have to give reasons for our behaviour that are "equally acceptable to all" (Singer, 2011, p. 93). We will have to adopt a disinterested viewpoint. Once reason reveals that this is the best way to justify one's behaviour, we cannot restrict this disinterested thinking to just the members of our own community. For reason is autonomous (Singer, 2011, p. 113-114). So, when we come across outsiders, we will still feel the force of reason to justify our behaviour in a disinterested way. All of this is possible even though giving as much moral consideration to outsiders as insiders might not be the best evolutionary strategy. Since animals are outsiders, the expressivist might argue that eventually this process leads us to include animals into the moral community (Singer, 2011, p. 120).

But, again, this cannot be all that the expressivist says. Granted, this might explain how animals become part of our moral community. It might therefore explain why Singer's claim that mere species membership is morally irrelevant seems true. But that is a far cry from explaining why it is or seems *conceptually* true.

The expressivist can of course add that this reason for becoming disinterested also affects our moral concepts. And she can suggest that it affects our moral concepts in such a way that they incorporate this disinterest. In that case the irrelevance of mere species membership might simply follow from our moral concepts and that would make mere species membership conceptually morally irrelevant.

The problem is that this story about disinterested moral concepts does not seem to line up with what we observe. Most people seem to believe that giving priority to kin is at the very least morally permissible. Yet, if our moral concepts were generally disinterested, kinship would be conceptually morally irrelevant.¹⁵ And even though those that believe kin priority is morally permissible might be wrong, they do not seem conceptually confused.¹⁶

The expressivist could reply that this objection misconstrues what it is for our moral concepts to be disinterested. It does not mean that we cannot give priority to our kin. It only means that everyone's appeal to kinship counts equally (Singer, 2011, p. 117-118). That is, our moral concepts are disinterested when they entail that my child being mine and your child being yours are equally good moral reasons for both of us to prioritize our own child. The disinterested viewpoint therefore does not diminish the moral relevance of kin priority.

Granted, on this view of disinterest, kinship can remain morally relevant. But this hurts rather than helps expressivists. The evolutionary story is meant to account for the conceptual moral irrelevance of mere species membership. It tries to do so by explaining how we become disinterested in mere species membership when it comes to morality. But on the alternative view of disinterest, moral relevance and disinterest come apart. That is, we can be disinterested with respect to a certain descriptive feature while it remains morally relevant. So, we might be disinterested with respect to mere species membership, yet still take it to be morally relevant. We do this when we believe that another *homo sapiens* being a member of our species and another creature being a member of another species are equally good reasons for each of us to prioritize members of our own species. Thus, on this alternative view of disinterestedness, by showing that we are disinterested with respect

¹⁵ Especially because the disinterested viewpoint comes into play when we start to take others apart from our own kin into moral consideration (Singer, 2011, p. 90-96).

¹⁶ The expressivist might of course suggest that our moral concepts are not generally disinterested, but just disinterested with respect to mere species membership. The issue with this is that the reason given by the sophisticated evolutionary story for why we become disinterested with respect to certain descriptive features does not just pertain to mere species membership. This reason also holds for kinship, membership of a tribe, a nation, a race, a religion, etc. It would therefore be quite remarkable if, when this reason affects the content of our moral concepts, it only does so by making moral concept disinterested regarding mere species membership.

to mere species membership, the expressivist has not yet shown that we take it to be morally irrelevant.

3.4.5. Fourth Expressivist Answer: Illusory Conceptual Truth

As mentioned, instead of trying to show that Singer's irrelevance claim is conceptually true, the expressivist could try to argue that it only *seems* conceptually true, but actually is not conceptually true.

She can argue that the general acceptance of the normative reading of Singer's irrelevance claim makes it seem as if the claim is conceptually true, even though it is not. That is, because we generally believe that mere species membership *does not* make a moral difference it seems as if it *cannot* make a moral difference.

More specifically, the expressivist might argue as follows: there is broad agreement that mere species membership does not make a moral difference. Neither consequentialism, Kantianism, nor virtue ethics suggest that it does. Given this broad agreement, it seems we would answer Foot's questions about virtue, duty, and charity with "No" when it comes to mere species membership. This is because we read Foot's questions as asking about specific theories. We read them as: do we have a Kantian duty to clasp our hands? Is it an act of charity to clasp our hands, according to virtue ethics? etc. It also explains why we think Singer's irrelevance claim seems knowable just through a proper understanding of the concepts, and why it seems to have framework status. It is just commonplace in our moral thought to believe that mere species membership does not make a moral difference. And lastly, the utter heterodoxy of the strict-speciesist's beliefs baffles us. Thus, the widespread belief that mere species membership does not make a moral difference explains why Singer's claim seems to bear the marks of conceptual truth. That is, it explains why Singer's claim seems conceptually true.

The first thing to note is that all of this is compatible with mere species membership being conceptually morally irrelevant. The second thing to note is that it is quite remarkable that there is

this broad agreement. Granted, there is a lot of agreement on certain general issues, such as whether we should treat each other respectfully. But when it comes to more narrow and specific issues we tend to see a lot of disagreement. We might therefore ask: what could explain the broad agreement on this specific topic? One obvious answer available to the realist is that mere species membership is conceptually morally irrelevant. That is, that the vast majority agrees that mere species membership does not make a moral difference could be explained by mere species membership being morally irrelevant as a conceptual matter. But this answer is obviously not available to the expressivist that opts for this strategy of illusory conceptual truth. An obvious answer that is available for the expressivist is an evolutionary one. She can argue that evolutionary forces push us to believe that mere species membership is morally irrelevant, and that this explains the broad agreement on this matter. The problem is that, as we have seen in the previous section, evolutionary forces would most likely push us to take mere species membership to be morally relevant, not morally irrelevant.¹⁷ Thus, the expressivist might explain why Singer's claim only seems conceptually true by citing the broad agreement, but it seems difficult for them to explain why there would be such a broad agreement.

3.4.6. Fifth Expressivist Answer: Moral Attitudes

So, explaining away the seemingly conceptual character of Singer's irrelevance claim runs into some important issues. Alternatively, expressivists might take another stab at what it means for something to be morally relevant. They could say that for a feature to be morally relevant is for it to be able to affect our moral attitudes. The expressivist might then argue that the nature of moral attitudes is such that they are insensitive to certain descriptive features like, for instance, mere species membership.

¹⁷ As an interesting side note, the error theory does not seem to fare any better than the expressivist. The error theorist, just like the expressivist, cannot adopt the realist's solution. For that solution relies on the existence of moral properties. And it cannot rely on this strategy of illusory conceptual truth.

Is there anything in the expressivist account of what constitutes a moral attitude that can support this? According to Blackburn (1998, p. 9-14), moral attitudes are unique because they dispose you to have certain related attitudes towards others. For instance, if you have moral attitude M, you are disposed to approve of people that also have attitude M, and disapprove of those who do not. But it seems that the strict-speciesist could also have these dispositions. She could want others to share her attitudes.

Gibbard argues that normative claims express the acceptance of a system of norms (1990, p. 126-150). If we take these norms to regulate anger and guilt, i.e. if these norms tell us to feel guilty at having done wrong, and angry with those who do wrong, then we can speak of a system of moral norms (Gibbard, 1990, p. 126). The strict-speciesist presumably thinks it is wrong to eat humans for strict-speciesist reasons. And there is no reason to think she would fail to express acceptance of norms to feel guilty at having eaten a human, and to be angry with cannibals.

Björnsson and McPherson's also identify social hostility, i.e. being angry with wrongdoers, as a unique feature of moral wrongness-judgments¹⁸ (2014, p. 16). They also take these judgments to be uniquely grounded in judgments about harms to others, and respecting the boundaries for sustainable social cooperation (Björnsson and McPherson, 2014, p. 15). The strict-speciesist might very well ground her wrongness-judgments in judgments about harms to others. It is just that she takes a specific stance on whose harms we ought to consider. Likewise, nothing in the strict-speciesist's view seems to bar her from grounding her wrongness-judgments in a respect for the boundaries that sustain social cooperation between humans.¹⁹

¹⁸ Instead of discussing moral attitudes more generally, Björnsson and McPherson focus on wrongness-judgments as a paradigmatic example of a moral attitude (2014, p. 4).

¹⁹ Björnsson and McPherson also discuss three features of wrongness-judgments that they share with judgments more generally. "These features are *aversion*, *personal-level acceptance*, and *engagement*" (2014, p. 14). Again, there seems to be no reason to think that the strict-speciesist's judgment would fail to possess these features: the strict-speciesist would be averse to treating marginal humans and animals as moral equals. The strict-speciesist actually believes *mere* species membership is morally relevant and not just feels that it is. That is, she accepts on a personal-level that it is morally relevant. And the strict-speciesist "is disposed to engage with the moral judgments of others" (Björnsson and McPherson, 2014, p. 15). She would try to convince others of her position and would argue against those that reject her judgments.

These accounts give us no reason to think of the strict-speciesist's attitudes as anything but ordinary moral attitudes.²⁰ They therefore fail to explain why moral attitudes are insensitive to features like mere species membership.

3.5. Conclusion

This paper has identified a new challenge for expressivism. It argues that, as it stands, expressivism, in contrast to moral realism, struggles to explain why certain irrelevance claims seem conceptually true. This struggle is not presumed to be insurmountable for the expressivist. This paper therefore does not offer a knockdown argument against expressivism. It does, however, show that some initially plausible expressivist answers to this challenge fail, thereby shifting the burden of answering this challenge squarely onto the expressivist's shoulders.

²⁰ Going hybrid does not seem to provide a solution either. Take Ridge's hybrid expressivist account: a normative judgment consists of two elements: a normative perspective, and a "belief that X would be highly ranked as an end by any admissible ultimate standard of practical reason" (2014, p. 119). It is unclear why an end that is defended on strict-speciesist grounds would not be ranked highly by any admissible standard for practical reason. That a standard of practical reason is admissible simply means the normative perspective does not rule it out (Ridge, 2014, p. 120). And the strict-speciesist has a normative perspective. That is, she has "a set of relatively stable policies against accepting certain kinds of standards of deliberation" (Ridge, 2014, p. 115). She has a policy against accepting standards of deliberation that do not take mere species into account. Ends defended on strict-speciesist grounds could presumably rank highly amongst the remaining strict-speciesist standards for deliberation. Thus, the hybrid nature of Ridge's account does not seem to provide any additional resources to explain why such a strict-speciesist policy, and therefore a strict-speciesist normative perspective, is problematic.

4. Relativism and Conceptual Moral Truth

4.1. Introduction

Cuneo and Shafer-Landau (2014) are the most recent defenders of a particular view on moral truths: they believe that there are substantive conceptual moral truths (p. 403). That is, they believe that some substantive moral claims are true because the contents of our moral concepts make them true. They call these conceptual truths the moral fixed points. The most obvious opponents for this view are moral error theorists: those who believe that there are no moral truths whatsoever.¹ Several philosophers (Evers and Streumer, 2016; Ingram, 2015, 2018) have discussed this clash between Cuneo and Shafer-Landau's account and moral error theory.

However, Cuneo and Shafer-Landau's account does not only pose a problem for those who believe there are no moral truths. Conceptual truths are necessary or absolute truths. They are true for everyone, no matter the context. The moral fixed points therefore pose a problem for anyone who believes that there are no absolute moral truths. Relativists should therefore follow the error theorist's lead and engage with Cuneo and Shafer-Landau's account as well.

In this chapter, I explore some of the ways the relativist might deal with these moral fixed points. I will focus on two general strategies: reconciliation and rejection. Although the former might seem doomed from the start and the latter seems the obvious way to go, I will argue that things are more complicated than that. A reconciliatory approach might actually work, although it would of course come with costs. And the rejection approach might work, but comes with more costs than one might at first expect.

In section 4.2, I give an overview of some of the more prominent relativist theories. In section 4.3, I categorise these relativist theories along three axes. In section 4.4, I describe Cuneo and Shafer-Landau's moral realist account. In section 4.5, I formulate the problem this account poses for

¹ Interestingly, Cowie (2020) uses Cuneo and Shafer-Landau's account to establish a moral error theory. I will not discuss this here. More on this in chapter 5.

relativist accounts. In section 4.6, I introduce two potential relativist strategies to deal with this problem: reconciliation and rejection, which I discuss in section 4.7 and 4.8.

4.2. Kinds of relativism

In this section, I will outline several relativist theories. The proponents of these theories all take relativism to be a hermeneutic and not a revolutionary theory. That is, they take their theories to describe our moral practice and not prescribe a moral practice. Throughout this paper I will stick to this conception of relativism.

4.2.1 Gilbert Harman: Moral Relativism Defended

This is what Harman says about ought claims:

If S says that (morally) A ought to do D, S implies that A has reasons to do D which S endorses. (...) [S]uch reasons would have to have their source in goals, desires, or intentions that S takes A to have and that S approves of A's having because S shares those goals, desires, or intentions. So, if S says that (morally) A ought to do D, there are certain motivational attitudes M which S assumes are shared by S, A, and S's audience.

(Harman, 1975, p. 9)

So, according to Harman, a moral judgment is a judgment that depends on the motivational attitudes of the agent who is judged.² So, for my claim “you ought to let me speak” to be true, it must follow from your motivational attitudes that you ought to do so. These motivational “attitudes M derive from an [(often implicit)] agreement. That is, they are intentions to adhere to a particular agreement on the understanding that others also intend to do so” (1975, p. 11-12). So, my claim “you ought to let me speak” is true, if we and our audience intend to keep to an agreement that involves a

² Or at least, the only fully-fledged moral judgments are such inner moral judgment, according to Harman (1975, p. 11).

commitment to let each other speak. Thus, Harman believes moral judgments are relativized to a shared agreement.

4.2.2 David B. Wong: *Moral Relativity*

According to Wong, a judgment of the form “A ought to X” should be read as:

By not doing X under actual conditions C, A will be breaking a rule of an adequate moral system that applies to him or her.

(Wong, 1984, p. 40)

There are several ways in which this introduces relativism into ought-judgments. The clause “under actual conditions C” makes explicit what is generally left implicit (Wong, 1984, p. 40): moral judgments are relative to a set of conditions C that are determined by the context at the time of utterance (Wong, 1984, p. 42). Here Wong makes a distinction between conditions that make a difference within a moral system to an individual, such as making a promise or not, and conditions that determine which adequate moral system applies to a society, such as the material resources available to a society (Wong, 1984, p. 42-43). The adequate moral system for a desert nomadic people will differ from the adequate moral system of an affluent trading society.

There is, however, another level of relativity: the extension of “adequate moral system” can vary. That is, a moral system can only be an adequate moral system relative to a certain moral ideal (1984, p. 39-40). For example, the moral ideal of autonomy is the aim of Kantian moral theory. But what people take to be moral ideals can vary, so what gets recognized as an adequate moral system will vary accordingly. This, it seems, is the most fundamental level of relativity. For according to Wong, this leaves open the possibility of it both being true that “A ought to X” and “A ought not to X”. For these claims can be made by different people with different moral ideals and thus different ideas about which adequate moral system applies to A under conditions C (1984, p. 45). So, Wong

believes that moral claims are relativized to an adequate moral system that is itself relativized to a moral ideal.

4.2.3 Sharon Street: Constructivism about Reasons

Street takes facts like “you have a reason to X” to mean the following:

[T]he fact that X is a reason to Y for agent A is constituted by the fact that the judgment that X is a reason to Y (for A) withstands scrutiny from the standpoint of A’s other judgments about reasons.

(Street, 2008, p. 223)

So, according to Street (2008, p. 225), X can only be a reason to Y for A, because A’s other normative convictions entail that A has reason X to Y. X cannot just be a reason for A to Y. Whether anyone has reasons to do anything depends on the other normative convictions of the agent whose reasons are assessed.

This, however, does not yet constitute a deep or all-the-way-down kind of relativism. There are constructivists, such as Korsgaard (1996), who would be on board with this analysis, but who take any agent who values anything at all to be committed to certain fundamental normative convictions. These convictions will entail that such an agent must have certain reasons. Street rejects this. She says she is “skeptical that there are any particular substantive judgments about reasons to which every agent is committed simply in virtue of valuing anything at all” (2008, p. 244). Thus, Street takes normative claims to relativize to other normative judgments and believes we are not committed to any particular normative judgments.

4.2.4 Stephen Finlay: *Oughts and Ends*

Finlay argues that the moral ought is just a species of the more ordinary modal ought, e.g. it ought to be dry by now. Finlay analyses this modal ought in the following way:

Modal ought: It ought to be the case that p = It is more likely, given circumstance C , that p than that [any p -alternative] obtains.

(Finlay, 2009, p. 322)

Ought claims are claims about probabilities, according to Finlay. If we say something ought to be some way, we say that under certain circumstances it is most likely that this something will be that way.

Instrumental oughts, according to Finlay, are simply modal oughts that occur under an “in order that” modifier (2009, p. 320). This yields the following analysis of instrumental oughts:

Ought _{e} : In order that e it ought to be the case that p = It is more likely, given circumstances C including its being the case that e , that p than that [any p -alternative] obtains.

(Finlay, 2009, p. 323)

The variable e stands for an end that one can reach by bringing about p . This end can be anything. I can say “In order that you become an impressionist painter it ought to be the case that you visit the Van Gogh museum in Amsterdam”, and this can be true, even if you do not actually have the goal to become an impressionist painter.

The moral ought is an instrumental ought with a moral end e . So, if I claim that we ought to stop eating meat because this reduces suffering, my claim should be understood as saying: it is more likely, given circumstances C including it being the case that we have reduced suffering, that we stopped eating meat than that any p -alternative obtains.³ Thus, Finlay thinks that moral claims are relativized to ends.

³ The alleged categorical nature of moral oughts is a matter pragmatics, according to Finlay. That is, Finlay takes categorical normative oughts to be instrumental oughts with some pragmatic posturing (2009, p. 333). When we seem to be making a

4.2.5 John MacFarlane: *Assessment Sensitivity*

MacFarlane believes that normative claims are only true relative to what he calls the context of assessment. The context of assessment is simply the context in which assessors of a proposition find themselves, which includes normative standards. MacFarlane therefore makes the following claim about how the relative truth of aesthetic claims is determined:

A sentence S is true as used at a context c_1 and assessed from a context c_2 iff S is true at c_1 , (Wc_1, Sc_2) , where Wc_1 is the world of c_1 and Sc_2 is the aesthetic standard of the agent of c_2 .

(MacFarlane, 2014, p. 67)

So, if I say “John’s new haircut looks great” to a friend, and another friend Mary wants to know if that’s true, the following will happen: Mary applies her aesthetic standards (Sc_2) to my claim and determines whether they entail that John’s new haircut in the world of me making that compliment (Wc_1) does look great. Thus, to determine the truth of an aesthetic claim the assessor applies their own aesthetic standards at the context of assessment (Sc_2) to the claim, taking the context of use (Wc_1) into account.

We can extrapolate this analysis to the moral domain. This means that if I claim that John ought to keep his promise and take me to the same barber, and Mary wants to know whether this is true, the following will happen: Mary applies her moral standards (MSc_2) to my claim and determines whether they entail that John should keep his promise at the time of me demanding John keep his promise (Wc_1). Thus, MacFarlane believes that the truth of normative judgments is relative to standards, and in particular, the standards of the person assessing the judgment.

4.3. Axes of Variation

categorically normative ought claim, we are making an instrumental end-relational ought claim, but we do it in a way that demands the interlocutor accept the end.

So, there is a lot of variance in theories that fall under the label of relativism. In this section, I parse out these theories along three axes: relative to who, relative to what, and content or truth relativism.

4.3.1 R_{who} : Agent, Speaker or Assessor

Relativist theories vary in *whom* they relativize to: R_{who} . There seem to be three options: agent-relativism, speaker-relativism, or assessor-relativism. I will illustrate the difference between these options with an example: Imagine a teacher says, “John ought to be kinder to his fellow students” during a parent-teacher meeting with John’s parents. How would these different kinds of relativism analyse the truth of the teacher’s moral judgment “John ought to be kinder to his fellow students”? Assume throughout this section that we relativize to moral standards:

Agent-relativism: What matters are the standards of the agent at whom the judgment is aimed: John’s standards. Thus, “John ought to be kinder to his fellow students” is true *iff* John’s standards entail that he ought to be kinder to his fellow students. And it is false *iff* John’s standards do not entail this.

Speaker-relativism: What matters are the standards of the speaker: the teacher’s standards. Thus, “John ought to be kinder to his fellow students” is true *iff* the teacher’s standards entail that John ought to be kinder to his fellow students. And it is false *iff* the teacher’s standards do not entail this.

Assessor-relativism: What matters are the standards of the assessor of the judgment. There is not just one person that can be the assessor. Anyone can be the assessor. In our example, John’s parents are assessors. But John, the teacher, and anyone else can assess the teacher’s claim. Take John’s parents as the assessors we

are interested in for now. Then “John ought to be kinder to his fellow students” is true *iff* the parents’ standards entail John ought to be kinder to his fellow students. And it is false *iff* the parents’ standards do not entail this.

Street defends a form of agent-relativism. According to Street, a normative judgment is true *iff* it is entailed by the agents’ other normative judgments (2008, p. 223). So, when the teacher judges that John ought to be kinder to his fellow students, the teacher thinks that John’s other normative judgments entail that John ought to be kinder to his fellow students. That is, if the teacher wants to say something true. Imagine that John believes that he ought to be kind to those who are kind to him. Further imagine that John believes his fellow students have indeed been kind to him. When his cat passed away they comforted him. In this situation, according to Street, the teacher’s judgment that John ought to be kind to his fellow students is true, since John’s other normative judgments and the non-normative facts entail that John ought to be kinder to his fellow students. Now imagine a different scenario where John instead believes that kindness toward others is a form of weakness and that a man must never be weak. In this situation the teacher’s judgment would be false, since John’s other normative judgments and the non-normative facts do not entail that John ought to be kinder to his fellow students.

Both Wong and Finlay defend a form of speaker-relativism. Wong believes that the teacher’s judgment is true *iff* it is true that John would be breaking a rule of an adequate moral system that applies to John were John not to be kinder to his fellow students. The speaker, the teacher, determines what would be the adequate moral system that applies to John (1984, p. 45). Thus, if the teacher is correct in believing that the adequate moral system that applies to John entails that he ought to be kinder to his fellow students, the teacher’s judgment is true. If the teacher is not correct in this way, for example because she is mistaken about what the adequate moral system that applies to John requires of him, the teacher’s judgment is false.

Finlay's account might seem neutral on the topic of R_{who} . However, the relevant end e in judgment end-relational judgments of the form "in order that e , X ought to Y ", is typically the speaker's end in cases of moral judgments (Finlay, 2009, p. 333). Finlay is therefore a speaker-relativist. Now, assume that the teacher wants to make a moral claim about what John ought to do. In that case, the teacher's statement is true *iff* John being kinder to his fellow students makes it most likely that John achieves the teacher's end e . And it is false *iff* this is not the best means for John to achieve end e . Thus, if the end e is to become compassionate human, the teacher's claim seems true. If the end e is to become infamous, the teacher's claim is false. The speaker, in this case the teacher, is the one that determines which end e we are talking about.

MacFarlane defends a form of assessor relativism. He believes that the teacher's claim is true *iff* the standards of the assessor entail that John indeed ought to be kinder to his fellow students, given the context in which the teacher made the claim (2014, p. 67). And it is false *iff* the assessors' standards do not entail this. Anyone in the scenario can be an assessor, but let us focus on the parents as assessors. If the parents believe kindness is a virtue and the context of the claim involves the fact that John has not been kind to his fellow students, the teacher's claim will be true. The teacher will most likely agree, since she makes the judgment in the first place. But now take John as the assessor. If John does not believe that kindness is a virtue, the teacher's claim will be false. So, there is not one standard that fixes the truth of the teacher's judgment on MacFarlane's account. The truth-value of the judgment will vary along with who assess the judgment and their standards.

Harman's view is harder to categorise neatly. Harman says that if the teacher judges that John ought to be kinder to his fellow students, the teacher implies that John has reasons to be kinder to his fellow students that the teacher and her audience endorse (1975, p. 9). One such reason might be the moral conviction that kindness is a virtue. So, if the following conditions hold then the teacher's judgment is true: 1) John has not been very kind to his fellow students, 2) John believes that kindness is a virtue, 3) the teacher believes kindness is a virtue, and 4) John's parents (and any other audience member) believe kindness is a virtue. This makes Harman a complicated case. The

combination of 1) and 2) adds an agent-relativist element. The combination of 1) and 3) adds a speaker-relativist element. And the combination of 1) and 4) adds relativity to the motivational attitudes of the audience. Harman's relativism thus seems to be a combination of several forms of relativism.

4.3.2 R_{what} : Standards, Normative Judgments, Ends

Relativist theories also differ in *what* they relativize to: R_{what} . As mentioned, Harman takes inner moral judgments to be true relative to shared agreements (1975, p. 11-12). Street believes that normative judgments are true relative to other normative judgments. That is, whether a normative judgment is true or not depends on whether other normative judgments, held by the agent in question, entail the normative judgment (2008, p. 223). Wong takes moral judgments to be relative to an adequate moral system: a set of rules that properly promotes a certain moral ideal such as autonomy (1984, p. 39-40). Finlay defends an end-relational theory of normativity. This means that he takes normative judgments to be relative to ends. Normative judgments should be understood as hypothetical judgments of the sort "X ought to Y in order that we achieve end *e*" (Finlay, 2009, p. 320). Lastly, MacFarlane takes normative judgments to be relative to normative standards of an assessor of the proposition expressed (2014, p. 67).

4.3.3 Content or Truth

Relativist theories also vary in the role that they let R_{who} and R_{what} play. This is the issue of indexicality: R_{who} and/or R_{what} can show up in the content expressed by normative judgments (content-relativism) or R_{who} and R_{what} only affect the truth-values of what is expressed (truth-relativism) (MacFarlane, 2014, p. 72-74).

According to content-relativism, R_{who} and/or R_{what} are part of the content that is expressed by normative judgments. Imagine R_{who} is the speaker and R_{what} are normative standards. In that case,

the proposition expressed in the example in section 4.2.1 would be: *John ought to be kinder to his fellow students according to the teacher's moral standards*. But if we were to add to the example that John's parents reassert this judgment when they talk to John, they would express the proposition: *John ought to be kinder to his fellow students according to the parents' moral standards*. On content-relativism, the variables of relativity (R_{who} and/or R_{what}) are part of the content, i.e. the proposition, that is expressed.

According to truth-relativism, the proposition expressed by the judgment "John ought to be kinder to his fellow students" is always the same: *John ought to be kinder to his fellow students*. R_{who} and R_{what} do not show up in the content that is expressed. But R_{who} and R_{what} do affect the truth-values of the proposition that we express (MacFarlane, 2014, p. 90). Imagine R_{who} is the assessor and R_{what} are normative standards. On truth-relativism, the sentence "John ought to be kinder to his fellow students" always expresses the same proposition: *John ought to be kinder to his fellow students*. But whether this proposition is true or false depends on who assesses it and their moral standards. Thus, the proposition is true when assessed by the teacher, since the teacher believes that kindness is a virtue. But the proposition is false when assessed by John, if he does not believe that kindness is a virtue. On truth-relativism, the variables of relativity (R_{who} and R_{what}) are not part of the content, i.e. the proposition, that is expressed. A normative sentence always expresses the same proposition. Instead, R_{who} and R_{what} affect the truth-values that apply to the proposition.

All authors discussed in section 1 (Finlay, 2009, p. 328; Harman, 1975, p. 10; Street, 2008, p. 225; Wong, 1984, p. 40), except for MacFarlane, are content-relativists. To give an example, Street claims that

the question "Is X a reason to Y?" is ill-formed in the absence of any (at least implicit) answer to the question "For whom?" In the absence of such specification, one has failed to point to the standard that makes the question make sense;

(Street, 2008, p. 225)

That Street takes questions like “Is X a reason to Y?” to be ill-formed or nonsensical shows that she is a content-relativist. Here is why: Street compares these normative questions to questions like “Is the Empire State Building taller?” (2008, p. 225). It is obvious that this question does not express a complete proposition. The predicate “taller” is a two-place predicate, not a single-place predicate. To fail to fill both places is to fail to express a complete proposition. She believes the same holds for the predicate “reason”. Although it might seem to be a two-place predicate, “X is a reason to Y”, it is in fact a three-place predicate, “X is a reason to Y for agent A”. A well-formed normative judgment that does make sense is therefore of the form “X a reason to Y for agent A”. This judgment expresses the proposition: *X is a reason to Y for A*. Thus, according to Street, normative judgments make sense *iff* the variable of relativity $R_{\text{who}}(A)$ is part of the proposition that is expressed. This makes Street a content-relativist.

MacFarlane is a truth-relativist. He is a truth-relativist, because he believes content-relativism cannot explain certain of our disagreement intuitions. As mentioned, John might claim that “John ought not to be kinder to his fellow students” and in doing so he would seem to be disagreeing with the teacher. But on content-relativism this is not the case, at least not in the sense that they have different views about the truth value of the same proposition. Let me explain.

Assume R_{who} is the speaker and R_{what} are moral standards and assume content-relativism. John’s and the teacher’s judgment are now cotenable, meaning that both John and the teacher can coherently believe that both their and the other’s judgment are true. This is because John and the teacher express different but compatible propositions. John expresses the proposition: *John ought not to be kinder to his fellow students according to John’s moral standards*. While the teacher expresses the proposition: *John ought to be kinder to his fellow students according to the teacher’s moral standards*. But, according to Macfarlane, noncotenability is often taken to be a sign of disagreement (2014, p. 120). That is, we take two people to disagree in case both of them would have to drop their own belief to be able to coherently accept the truth of the other’s belief. According to MacFarlane, content-relativism cannot account for the apparent disagreement between

John and the teacher, at least not in the sense of them ascribing different truth values to the same proposition. That's why MacFarlane is a truth-relativist.

So, how is the judgment "John ought to be kinder to his fellow students" relativized? According to MacFarlane, this judgment simply expresses the proposition: *John ought to be kinder to his fellow students*. But the truth-value of that proposition depends on R_{who} and R_{what} . According to all the other authors discussed in section 1, this judgment expresses the relativized proposition that looks something like: *John ought to be kinder to his fellow students according to R_{who} 's R_{what}* .

4.3.4 Recap

As we have seen, relativist theories can be parsed along at least three lines: 1) the variable R_{who} : who do we relativize to? The options here are: the speaker, the agent, or the assessor. 2) The variable R_{what} : what do we relativize to? The common options are: ends, standards, and other moral judgments. 3) Do R_{who} and R_{what} affect the content or only the truth of propositions? The options here are: content-relativism or truth-relativism.

4.4. Conceptual Moral Truth

Now that we have a better grasp of relativism and the diversity within it, it is time to bring the second element into play: conceptual moral truths.

Cuneo and Shafer-Landau (2014) argue that there are conceptual moral truths. They call these truths the moral fixed points. Examples include (p. 405):

a) It is *pro tanto* wrong to engage in the recreationally slaughter of a fellow person,

and

b) If acting justly is costless, then, *ceteris paribus*, one should act justly.

Both a) and b) seem to bear several marks of conceptual truth (Cuneo and Shafer-Landau, 2014, p. 407-408): 1) If true, they are necessarily true, 2) they have framework status. They in part determine the boundaries of a certain subject area, just like, say, the claim “God is omnipotent” has framework status within Christianity. That claim determines a boundary of Christian theological debate. If you do not accept this claim, yet make claims about a god, you cannot be thought of as referring to the Christian God. 3) Denial of this claim bewilders competent speakers. 4) The truth of the claims seems *a priori* knowable.

Cuneo and Shafer-Landau take a) and b) to have these four marks of conceptual truth. But they believe that a) and b) are conceptual truths, because the following holds for both claims: it belongs to the essence of the moral concept involved that, necessarily, the moral concept applies to any action of the kind in question. In other words: “a proposition \langle that x is F \rangle is a conceptual truth if it belongs to the essence of ‘F’ that, necessarily, anything that satisfies ‘x’ also satisfies ‘F’” (Cuneo and Shafer-Landau, 2014, p. 410). One need not buy into the talk of “essences of concepts” to accept this, as long as one believes particular concepts have particular contents. Thus, we can say a) is conceptually true, because it follows from the content of *PRO TANTO WRONG* that, necessarily, anything that constitutes the recreational slaughter of a fellow person also satisfies *PRO TANTO WRONG*. And the same holds for b), and other moral fixed points.

4.5. The Problem for Relativism

Assume for the sake of argument that Cuneo and Shafer-Landau’s arguments are sound. This means the conceptual moral truths pose a problem for any relativist. The relativist believes that substantive moral judgments can at most be relatively true. But if Cuneo and Shafer-Landau are right, this is not true, because the moral fixed points are absolute truths. That is, they are necessarily true and their

content has no relativizing elements. It does not matter which agreement you intend to keep, what your other moral judgments are, which adequate moral system applies to you, what ends you pursue, which moral standards you endorse, who makes the claim, about who the claim is made, or who assesses the claim, it is *pro tanto* wrong to recreationally slaughter a fellow person. So, it seems that if there are conceptual moral truths, relativism must be false.

4.6. Relativist Solutions to the Problem

The relativist can of course try to offer a solution to this problem. I will discuss two general strategies the relativist might go for: reconciliation or rejection. The reconciliatory strategy tries to integrate the moral fixed points into the relativist account; the rejection strategy tries to save relativism by rejecting conceptual moral truths. I will deal with these strategies in order.

4.7. Reconciliation

Some relativists have made what sound like reconciliatory remarks. They mostly do this from a desire to account for the strong absolutist intuition that acts like torturing people for fun must be wrong. I will first discuss some solutions based on these reconciliatory remarks. After showing why these solutions fail, I will discuss some more general solutions and explore their costs. But before I do so, we must lay down some desiderata for a successful reconciliatory solution.

There are four desiderata:

- 1) *Conceptual truth*: The solution must entail that some moral claims are conceptually true. This also means that these claims are necessarily true, since conceptual truths are necessary truths. After all, the moral fixed points are conceptual truths.

2) *Substantive moral truths*: The conceptual truths that the solution yields must be non-trivial, important moral claims, like Cuneo and Shafer-Landau's moral fixed points. That is not to say that Cuneo and Shafer-Landau's list of moral fixed points must follow from the solution exactly, but it must produce conceptual truths similar to those moral fixed points. This is because Cuneo and Shafer-Landau's moral fixed points are supposedly some of the most fundamental moral truths.

3) *No Counterintuitive Truths*: The solution cannot entail that counterintuitive claims are conceptually true as well. The aim is to reconcile Cuneo and Shafer-Landau's moral fixed points with relativism and allowing for counterintuitive conceptual truths undermines that. So, even if a solution ensures the conceptual truth of a), it must not ensure the conceptual truth of claims like

c) it is *pro tanto* right to recreationally murder kittens.

4) *Fixing Boundaries*: The solution must produce conceptual truths that fix the moral boundary. A truth fixes the moral boundary when a normative system cannot count as a moral system unless it encompasses that truth. Thus, for a solution that establishes a) as conceptually true to be a proper solution, it must also ensure that any normative system that does not encompass a) cannot count as a moral system. Again, the aim is to reconcile relativism with Cuneo and Shafer-Landau's moral fixed points, and they believe that this boundary fixing property is an essential part of these moral fixed points (2014, p. 406).

Thus, we will judge reconciliatory solutions by these four desiderata, although not every proposed solution will be assessed for their ability to satisfy all of them. But even if a solution manages to

satisfy all desiderata there remains one ultimate question: can relativists accept the solution? That is, does it sit well enough in their relativistic view of morality or is it a problematic (*ad hoc*) addition to an otherwise coherent theory?

4.7.1 Specific Solutions

In this section, I will discuss potential solutions based on remarks by specific authors. I will start by discussing Harman's remarks.

4.7.1.1 Harman

As mentioned, Harman believes that moral judgments are relativized to a shared agreement. That is, the claim that "John ought to be kinder to his fellow students" is true *iff* John has reasons to be kinder to his fellow students that are endorsed by the teacher and the parents (and any other audience members). These reasons or motivational attitudes "are intentions to adhere to a particular agreement on the understanding that others also intend to do so" (Harman, 1975, p. 12).

Harman also claims that there are certain things we cannot seem to intend. For example, Harman says that we can rationally intend to not preserve our own life only in very extreme circumstances (1975, p. 14-15). This is because we know we have no control over our behaviour when our life is threatened. In general, we would simply save our own lives. And "you cannot now intend to do something later which you know you would not do" (Harman, 1975, p. 14-15). This also means "you cannot now keep an agreement not to preserve your life if it is threatened" (Harman, 1975, p. 14-15). So, the right to self-preservation has to be part of any agreement from which we get our moral reasons, which means that everyone will always have a reason to preserve their own life. And since Harman seems to think that the intention to save our own life will override any previous intention to the contrary (1975, p. 14), it seems that everyone has an overriding reason to preserve their own life. We can interpret this as Harman saying "X may preserve their own life" is always true.

Harman might therefore claim that his form of relativism allows for substantive moral judgments that are always true, just like the conceptual moral truths.

Now let us examine whether this solution meets the desiderata mentioned above. This solution fails to meet desideratum 1) *conceptual truth* in two ways: first, for the solution to establish conceptual truths, it needs to establish that “X may preserve their own life” is necessarily true. But Harman himself admits, there are extreme circumstances under which “X may preserve their own life” is false (1975, p. 14). For example, imagine an inattentive truck driver whose only choice is between hitting a group of children crossing the road or swerving off the road into a large statue. Hitting the children results in the truck driver, but not the children, surviving, whereas swerving off the road will result in the children, but not the truck driver, surviving. It does not seem unthinkable that the claim “The truck driver ought to swerve off the road” is true. That is, it seems that the speaker (me), the agent (the truck driver), and the audience (you and whoever else), can all endorse reasons that say the truck driver ought to swerve off the road. This shows that “X may preserve their own life” is not necessarily true.

Second, even if we assume that this claim is necessarily true, Harman’s solution does not establish that it is conceptually true. Moral fixed points are true in virtue of the content of the concepts involved. However, under the assumption that “X may preserve their own life” is necessarily true, this will be because we cannot rationally form the intentions that make it false. Thus, “X may preserve their own life” might be necessarily true, but it will be because we cannot intend to do things we know we cannot do, and not because the content of the concepts involved entail its truth.

This is also the reason why the solution fails to meet desideratum 2) *substantive moral truths*. The only necessary truths that this solution can establish are those where we psychologically are unable to intend otherwise. But there are ample cases in which people have intended (and done) things that are contrary to the moral fixed points. People have intended to murder, torture, etc. recreationally. And people have intended to act unjustly even if acting justly came without any extra

costs. The ground of Harman's "necessary" moral truth, the psychological inability to intend otherwise, therefore is too barren to entail substantive moral truths like the moral fixed points.

4.7.1.2 Wong

Maybe the relativist can build on Wong's reconciliatory remarks instead. Wong acknowledges the absolutist's intuitions that "there are substantive moral beliefs that we cannot conceive to be false" (1984, p. 59). He gives two examples (Wong, 1984, p. 59-60):

d) just the colour of someone's skin is irrelevant for how we ought to treat them

and,

e) we ought not to torture people for our own amusement.

Judgments d) and e) might be good candidates for moral fixed points.

Wong claims that we cannot conceive d) and e) to be false for different reasons. For d) he believes that our being part of a moral tradition whose morality can be seen as an expanding circle means we have to believe d) (Wong, 1984, p. 59-60). That is, first we only had concern for others of our family or tribe. But, throughout our history, we gradually came to understand that outsiders should receive moral consideration as well. And Wong (1984, p. 59-60) seems to claim that this understanding is incompatible with just taking someone's skin colour to dictate a certain kind of treatment. Whether this is indeed incompatible is up for discussion. However, what is clear is that, according to Wong, this expansion of the moral circle is not necessary. He writes:

It maybe asked whether this process of gradual inclusion is a necessary consequence of the nature of morality. I see no logical necessity in the fact that kinship groups usually unite into larger groups. And if a larger group remains isolated from the rest of the

world, it may have no occasion to question whether its rules should be extended beyond its boundaries. Its morality will remain a purely tribal morality, and I see no compelling reason based on the concept of morality for it to be otherwise.

(Wong, 1984, p. 56)

This means that d) is not necessarily true, which in turn means it does not meet desideratum 1) *conceptual truth*.

What about e)? This is a bit more complicated. Wong argues that we cannot take e) to be false (1984, p. 60), since taking e) to be false would undermine one of the primary functions of morality: the resolution of inter-personal conflict (1984, p. 38). According to Wong, “[a] system of rules that allowed torture for amusement could not possibly achieve this function” (1984, p. 60). This could be taken to mean that a system of rules that does not encompass e) is not a moral system. e) would fix a boundary of what is and what is not moral, just like Cuneo and Shafer-Landau take the moral fixed points to do. So, this solution might satisfy desideratum 4) *fixing boundaries*. But does the solution deliver on the other desiderata?

Let us start with desideratum 2) *substantive moral truths*: does the solution establish substantive moral truths like the moral fixed points? e) resembles Cuneo and Shafer-Landau’s (2014, p. 405) moral fixed point

a) it is *pro tanto* wrong to engage in the recreational slaughter of a fellow person.

Only two things differ: the horrible act, and whether that act is *pro tanto* wrong or ought not to be done. The second difference seems most significant. But if something ought not to be done, it also *pro tanto* ought not to be done. And if something *pro tanto* ought not to be done, it is likely *pro tanto* wrong as well. So, the resemblance between a) and e) seems close. Thus, if this solution works it could establish some substantive moral truths like the moral fixed points.

Now, to address this “if”: does Wong’s solution meet desideratum 1) *conceptual truth*? There are two reasons to think it does not: first, it is not clear that e) is necessarily morally wrong. Other

than relying on intuition, which Wong presumably hopes everyone shares, he does not provide an argument for why a system of rules that fails to encompass e) cannot help resolve inter-personal conflict. And there seem to be systems of rules that do serve this function, but do not encompass e), which would mean e) is not necessarily true.

Hedonistic utilitarianism seems to be such a system of rules. Imagine a hedonistic utilitarian society. This society's system of rules will not include e). It might not condone all instances of torture for amusement, but it will condone some. It will condone all those instances where the amusement or pleasure derived from the torture outweighs the pain inflicted by the torture. It is not at all clear why hedonistic utilitarianism could not help resolve the inter-personal conflict in this society. To be clear, we might not like the way conflict is resolved on this view, but that does not matter. What matters is whether hedonistic utilitarianism can resolve inter-personal conflict, not whether it can do so satisfactorily according to our standards. And it seems that it can do so, especially in a society where it is endorsed as a system of conflict resolution. This means that hedonistic utilitarianism, according to Wong, counts as a moral system. This in turn means it is possible to deny e) on the basis of a moral system. Thus, e) is not necessarily true and therefore not conceptually true.

Second, even if we assume that e) is necessarily true, it is not clear whether Wong's solution establishes conceptual truths. Conceptual truths are true because of the content of the concepts involved in those claims. Under our assumption, e) is true because of the function of morality more broadly. It is not evident that the function of a moral system is part of the individual moral concepts like RIGHT or WRONG. Thus, even if we disregard hedonistic utilitarianism and assume that e) is necessarily true, it is unclear whether Wong can establish conceptual truths and satisfy desideratum 1) *conceptual truth*.

4.7.2 General Solutions

So, these reconciliatory remarks seem to remain just that, remarks. Building them into a proper reconciliation of the moral fixed points with relativism does not seem possible. But there are of

course other ways to attempt to produce a reconciliatory account. The one I will discuss here tries to secure conceptual moral truths for the relativist by demarcating the moral domain in a specific way.

4.7.2.1 Demarcation: Utilitarian Demarcation

The claim “Y ought to X” is normative. But within the normative domain there are many sub-domains. An instance of “Y ought to X” can be an instance of an aesthetic judgment. This would be the case when I claim that Jeff Bezos ought to read *Das Kapital* by Marx since it will allow him to experience the beauty of the writing. Or it can be an epistemic judgment. For example, when I claim that Jeff Bezos ought to read *Das Kapital* to better understand politics and economics. Or it can be a prudential judgment. For example, when I claim a that Jeff Bezos ought to read *Das Kapital*, since having an understanding of the struggles of the working class might help him strategize to prevent a worker uprising. Or it can be a moral judgment. For example, when I claim that Jeff Bezos ought to read *Das Kapital*, since understanding the exploitation he is responsible for might lead him to treat his workers better and improve their well-being.

What makes one claim an aesthetic one and another a moral one are the reasons we cite to support the judgment, or so the relativist could argue. When I claim that Jeff Bezos ought to read *Das Kapital* because the writing will give him an experience of beauty, I cite the aesthetic value of beauty as my reason for making the judgment. This makes it an aesthetic judgment. When I claim that Jeff Bezos ought to read *Das Kapital* because it might improve the well-being of his employees, I cite the moral value of well-being as my reason for making this judgment. This makes it a moral judgment.⁴

The relativist might then claim that the same holds for concepts. General normative concepts like RIGHT or WRONG can belong to different normative sub-domains. Something might be the right colour to use, another the wrong one. Some act might be the right one, another the wrong one.

⁴ In some way this resembles Wong’s strategy for appeasing absolutist intuitions. He too believes that the way the moral domain is demarcated entails the truth of certain substantive moral claims. What is different, however, is that Wong focuses on the function of morality to demarcate the moral domain, not the values on which claims are grounded (1984, p. 60).

What makes one instance of RIGHT an aesthetic one and another a moral one depends on the reasons cited for why the concept applies. In the aesthetic case we might use the concept RIGHT to describe the use of a certain colour because it looks beautiful. In the moral case we might use the concept RIGHT to describe an act because that act promotes the well-being of the less fortunate.

This suggests that there are some conceptual limits to the reasons for which moral concepts might apply in any given situation. And that, the relativist might argue, entails some conceptual truths. To illustrate how this might work, and to keep things simple, imagine that for any instance of RIGHT or WRONG to be moral, the reasons cited for its application must take well-being to matter.⁵ If they do not, the claim is not a moral one. So, the relativist might argue that for the concept MORALLY WRONG to apply to an act, that act must at the very least negatively affect well-being. And for the concept MORALLY RIGHT to apply to an act, that act must at the very least positively affect well-being.⁶ So, well-being is all that counts for wrongness and rightness. Call this the *utilitarian demarcation* of the moral domain or D_u for short.

This D_u , the relativist could go on to argue, entails the truth of certain substantive moral claims like the following moral fixed point (Cuneo and Shafer-Landau, 2014, p. 405):

- a) It is *pro tanto* wrong to engage in the recreational slaughter of a fellow person.

This is because the content of the concept SLAUGHTER is such that slaughtering a fellow person, on its own, will always affect the well-being of the slaughtered person negatively.⁷ So, to an extent, or in so far as only certain facts are concerned (i.e. *pro tanto*), the fact that recreational slaughter involves slaughter means that under D_u it is always morally wrong to recreationally slaughter a fellow person.

⁵ There is a legitimate worry that this might be keeping things a little too simple. I will address this worry in the next section.

⁶ This, of course, cannot be all that the relativist says. An act can still be right, even if it negatively effects well-being. For instance, if the act has a greater positive effect on well-being, or if all alternative acts have an even greater negative effect on well-being.

⁷ One might wonder whether this is true even if the person that is being killed is better off dead than alive, as for instance, when they are in an endless state of suffering. Here we should distinguish SLAUGHTERING from MERCIFUL KILLING. You can mercifully kill someone who suffers endlessly, by ending their life as painlessly as possible. To slaughter someone who suffers endlessly in no way implies a process that is as painless as possible. This means that even in cases of endless suffering, slaughtering a person negatively affects their well-being compared to mercifully killing them.

The content of the concepts SLAUGHTER, *PRO TANTO*, and MORALLY WRONG (under D_u) just bear this out. So, it seems that the relativist can take a) to be conceptually true.

So, why does this illustration of a demarcation solution not work? It seems to establish a) after all. The solution does not work because it faces a dilemma. For, depending on how you think we should evaluate *pro tanto* claims, this solution fails to satisfy either desideratum 1) *conceptual truth* or 3) *no counter-intuitive truths*. I will explain how it fails to satisfy 3) *no counter-intuitive truths* first.

If the relativist's reasoning for why a) is conceptually true under D_u is sound, then it will also be conceptually true that

a[^]) it is *pro tanto* right to engage in the recreational slaughter of a fellow person

This is because, to a certain extent, or in so far as only certain facts are concerned (i.e. *pro tanto*), the fact that recreational slaughter is done recreationally means it is done with the aim of improving your own well-being. That in turn means that it is always *pro tanto* right. For, under D_u , the moral concept MORALLY RIGHT is applicable to an act *iff* that act promotes well-being. So, the content of the concepts RECREATION, *PRO TANTO*, and MORALLY RIGHT (under D_u), entail that a[^]) is conceptually true.

Now, although believing that both a) and a[^]) are conceptually true is not contradictory (because of the meaning of '*pro tanto*'), it does mean that this solution cannot deliver on desideratum 3) *no counter-intuitive truths*. For, given that a) is a quintessential moral fixed point, it seems that a[^]) is clearly a counter-intuitive conceptual truth.⁸

The relativist might reply that this misconstrues how we should evaluate *pro tanto* claims. Even though '*pro tanto wrong*' does mean *wrong to an extent, or wrong considering only certain facts*, this does not mean we can only consider the recreational part of recreational slaughter. We

⁸ The problem is greater than just a[^]) of course. The formula for deriving a[^]) can be repeated for any horrible act by inserting that it is done recreationally. Thus, genocide, torture, rape, etc. all seem at least *pro tanto* wrong. But if we add that these acts are committed recreationally, it would be conceptually true that they are *pro tanto* right under D_u as well.

have to consider the act of recreational slaughter of a fellow person as a whole. If this is true, this means that a[^]) cannot be conceptually true, because we can only establish that it is conceptually true by just focussing on the recreational part of that act. That is, by not considering all descriptive facts.

Whether the argument for the conceptual truth of a[^]) does indeed misconstrue how we should evaluate *pro tanto* claims is up for discussion. However, when we assume that it does and think that the act of recreational slaughter should be considered in its entirety, the solution based on D_u fails to satisfy 1) *conceptual truth*. This is because, on this ‘as a whole’ reading, the solution fails to establish not only a[^]) but also a) as conceptually true.

We can only establish a) as conceptually true if we just focus on the slaughter part of the act of recreational slaughter. If we have to consider the act of recreational slaughter of a fellow person as a whole and we are using D_u , we will need *a posteriori* information about how much well-being is created and how much well-being is destroyed by this act. This is because the fact that a) concerns recreational slaughter specifically creates the possibility of this act positively affecting well-being (recreation), as well as the obvious possibility of this act affecting well-being negatively (slaughter). Thus, to evaluate whether the moral concept *WRONG* applies to this act *pro tanto*, we would need *a posteriori* information about the balance of these positive and negative contributions to well-being. This means that the relativist cannot claim the moral concept *WRONG* (under D_u) applies to this act just in virtue of the content of the concept alone. That is, it means we cannot take a) to be conceptually true, which means that this solution does not satisfy desideratum 1) *conceptual truth*.⁹

⁹ This is disputable. We reach this conclusion because the act of recreational slaughter of a fellow person contains an element that can affect well-being negatively (slaughter), but also contains an element that can affect well-being positively (recreation). But consider

a’) It is *pro tanto* wrong to engage in the slaughter of a fellow person.

Unlike a), a’) does not contain an element that can affect well-being positively. It only contains an element that negatively affects well-being. Under D_u , it therefore seems that a’) is conceptually true. This is because the same reasons the relativist gave for why a) is true if we solely focus on slaughter also apply to a’), which only concerns slaughter. The relativist might then point out that desideratum 2) *substantive moral truths* explicitly mentions that the solution need not produce Cuneo and Shafer-Landau’s moral fixed points exactly. It should, however, produce conceptual truths sufficiently like those moral fixed points. And a’) seems to be sufficiently like a). So, this solution based on D_u does not fail to satisfy desideratum 1) *conceptual truth*. This means that one horn of the dilemma falls away.

So, this illustrative D_u solution either entails that a[^]) is conceptually true and thus fails to satisfy desideratum 3) *no counterintuitive truths* or it entails that a) is not conceptually true and thus fails to satisfy 1) *conceptual truth*.

There is a way out of this dilemma however. The second horn seems to be produced by the fact that the *utilitarian demarcation* of the moral domain D_u is an aggregative view. For this means that, if we consider an act that can both positively and negatively affect what is morally valuable, we need *a posteriori* information about the extent of positive and negative effect to determine its moral status. So, if the relativist keeps the structure of this solution, but replaces D_u with a demarcation of the moral domain that is non-aggregative, she can avoid the dilemma. This means that she has to figure out which values are plausibly moral ones and show that a demarcation based on those values delivers the moral fixed points, while ensuring that this demarcation is non-aggregative. But this seems like quite a task. Instead, she might piggyback on Cuneo and Shafer-Landau.

4.7.2.2. Demarcation: Piggybacking

Given that the aim of the reconciliatory strategy is to establish conceptual moral truths similar to Cuneo and Shafer-Landau's moral fixed points, the relativist might simply use their demarcation of the moral domain. That is, the relativist might use the solution from the previous section, but swap the utilitarian demarcation for Cuneo and Shafer-Landau's demarcation of the moral domain. Their demarcation surely allows the relativist to establish the moral fixed points. It seems that would be a bit too quick, because the relativist will still have to determine what this demarcation looks like and whether it is non-aggregative. Instead, she might fully piggyback on Cuneo and Shafer-Landau and simply adopt their way of establish the moral fixed point. That is, she might adopt the view that moral concepts have essences and that these essences entail certain conceptual moral truths.

This seems correct. However, it is unclear whether a very thin demarcation like D_u can establish the other moral fixed points, like

b) If acting justly is costless, then, *ceteris paribus*, one should act justly.

Besides, the objections in section 4.7.2.2. will also apply to the relativist solution based on D_u .

This solution obviously satisfies the four desiderata, which means there is one question left to answer: can the relativist accept the solution? This, it turns out, depends on the extent of absoluteness that the relativist is willing to tolerate.

Let us assume that the relativist is at least willing to accept the absoluteness of the moral fixed points, else this reconciliatory effort is doomed from the start. I will argue that the absoluteness of the moral fixed points cannot be contained to just those fixed points. Their absoluteness seems to spread in at least two ways. First, the essences of our moral concepts have to be absolute. If they are not the solution fails to meet desiderata 4) *fixing boundaries*. Let me explain. Suppose the relativist thinks that different people might take the same moral concept to have different essences. And further imagine that there is not one objectively correct view about this. That is, imagine that the matter of what essences moral concepts have is a relative matter. Cuneo and Shafer-Landau, and thus the relativists, believe our moral concepts to have essences E_m , but someone else might think they have essences E_m^* . What, we might wonder, is the status of moral fixed points on this view of essences of moral concepts? It seems that a relativist with this view must say that what Cuneo and Shafer-Landau take to be the moral fixed points, are just conceptual truths for E_m . A different conception of the essences of our moral concepts, like E_m^* , E_m^{**} , E_m^{***} , etc. might entail different conceptual moral truths. That is, if these are sufficiently different from E_m .

But this means that Cuneo and Shafer-Landau's moral fixed points cannot fix the boundaries of the moral domain. They might fix the boundaries of a certain moral systems, those that adhere to E_m . But it is very well possible for a different set of systems, stemming from E_m^* , E_m^{**} , E_m^{***} , etc., not to encompass

a) it is *pro tanto* wrong to engage in the recreational slaughter of a fellow person.

And if not one conception of the essences of our moral concepts is privileged, there is no reason to think that systems that do not encompass a) are not moral systems. So, if the question of what the

essences of our moral concepts are at most has a relative answers, Cuneo and Shafer-Landau's moral fixed point a) can hardly be said to be boundary fixing.

This means that the piggybacking solution can only satisfy desideratum 4) *fixing boundaries*, if we take Cuneo and Shafer-Landau's view of the essences of our moral concepts (E_m) to be absolute. That is, if we take E_m to be the only correct way to understand the essences of our moral concepts. The relativist might not be willing to accept this absolutism.

As mentioned, there is a second way in which absoluteness spreads on this solution. It does so through nearly identical moral claims. Remember that we have assumed that the relativist can establish Cuneo and Shafer-Landau's moral fixed points. This means that

a) it is *pro tanto* wrong to engage in the recreational slaughter of a fellow person,

is not only true, but is conceptually true. This in turn means that a) is absolutely true. It is true no matter whom it concerns, or who utters or assesses it and their moral standards, other normative judgments, ends, etc.

But what about a claim like

f) it is *pro tanto* wrong to engage in the recreational slaughter of a twelve-year-old?

If true, in what way is it true? It seems that personhood is gradual. A human is more of a person than a dolphin, but a dolphin is more of a person than a robin. One might think that, on average, a human twelve-year-old is right on the cusp of full personhood (if not, there surely is an age for which this is true on average). This means that f) (or a claim similar to f)) is nearly identical to a). But does the truth of f) fully depend on the variables of relativity (R_{who} and R_{what})? That is, in contrast to a), can f) only be relatively true?

And what about other claims that are nearly identical to a), like

g) it is *pro tanto* wrong to engage in the recreational slaughter of an eleven-year-old

and,

h) it is *pro tanto* wrong to engage in the recreational slaughter of a ten-year-old?

Or what about

i) it is *pro tanto* wrong to engage in the recreational slaughter of a fully-abled adult orangutan?

The relativist can react in basically two ways: dig in and maintain that only the moral fixed points are absolutely true or accept that there are other moral claims that are absolutely true.

Digging in requires accepting the odd situation where, although f) is nearly identical to a), the truth of a) in no way depends on the variables of relativity (R_{who} and R_{what}), while the truth of f) fully depends on them. Or, put differently, digging in requires you to fit a set of seemingly gradually different moral claims (claims a), f), g), h), i), etc.) within a very rigid, black and white, view of which moral claims are absolutely and which are relatively true. This might make us doubt the hermeneutic nature of such a relativist theory. That is, it seems like the relativist would be prescribing a certain division between absolute and relative moral truth, instead of describing that division. This of course is not necessarily problematic, but it is problematic for the prominent relativist authors I have discussed, since they all take their theories to be descriptive or hermeneutic.

However, accepting that there are other moral claims, outside of the moral fixed points, that are absolutely true surely is not something the relativist will be too happy about either. The cost of reacting this way of course depends on how far the matter of near-identity can go. It depends on

how large this set of near-identical absolute truths is. Determining this precisely will be hard. But it being at least somewhat plausible that claims g) through i) might be absolutely true and remembering that something similar goes for all other moral fixed points suggests that this set will be several times larger than the set of moral fixed points.

Now, note that the issue is not that the relativist cannot draw a precise line when moral claims are different enough from the moral fixed points so that there is no longer pressure to consider them absolutely true. Nor that this leads to the conclusion that all moral claims must be absolutely true, given that they are nearly identical to claims, that are nearly identical to claims, that are ... nearly identical to claims that are absolutely true. The issue is that the relativist can accept that the moral fixed points are absolutely true, but that if she does, she cannot just take these moral fixed points to be absolutely true. There are moral claims that are nearly identical to the moral fixed points that therefore also seem absolutely true. And the larger this set of moral claims, the less likely the relativist is willing to accept the piggybacking solution.

So, somewhat surprisingly, it seems relativism and the moral fixed points might be reconciled. Unsurprisingly, however, this reconciliation requires a lot of explanation: it has to explicate the relativists' demarcation of the moral domain or it requires fully adopting Cuneo and Shafer-Landau's account about the essences of moral concepts. This reconciliation would also come with some potentially heavy costs: an objectivist view of the essences of moral concepts, and either adopting a seemingly prescriptivist view about which moral claims are and are not absolutely true or accepting that absolute truth spreads beyond the moral fixed points. It is up to the relativist to give this explanation and decide whether she finds these costs acceptable.

4.8. Rejecting the Moral Fixed Points

Relativists who are not willing to accept the conciliatory strategy are likely to turn to a rejection strategy. In pursuing this strategy, the relativist need not reinvent the wheel. She can borrow the

arguments from error theorists that defend themselves against Cuneo and Shafer-Landau's account by rejecting the moral fixed points.

Error theorists deny that there are any substantive moral truths. This means that they think the moral fixed points are not true. But Cuneo and Shafer-Landau take the moral fixed points to be conceptually true. They therefore believe that error theorists suffer from some kind of conceptual deficiency (Cuneo and Shafer-Landau, 2014, p. 412-414). Error theorists, they believe, do not fully grasp the content of moral concepts. After all, if they did, they would not reject the moral fixed points. Some philosophers (Evers and Streumer, 2016; Ingram, 2015, 2018) have tried to defend error theorists from this accusation of conceptual deficiency. They do so by rejecting the moral fixed points.

Let us look at Evers and Streumer's argument. Evers and Streumer argue that the reasons Cuneo and Shafer-Landau cite for why a) is conceptually true also entail the conceptual truth of

j) benevolence is rewarded by God.

Evers and Streumer, for instance, argue that j) fixes a boundary of Christian theology (2016, p. 3). That is, if you participate in a discussion on the value of benevolence within Christianity and you do not think that j) is true, it seems you are talking about a different God than the Christian one.

But, Evers and Streumer argue, j) cannot be conceptually true (2016, p. 3). For if j) were conceptually true and we sometimes are benevolent, our benevolence would be rewarded by God, which means that God would exist. But surely the existence of God cannot follow from just the truth of j) and our benevolent acts. So, j) cannot be conceptually true. Nonetheless, j) might seem conceptually true, because there seems to be a conceptual truth in j)'s neighbourhood (Evers and Streumer, 2016, p. 4):

j*) If anything is rewarded by God, benevolence is rewarded by God.

Given the content of BENEVOLENCE and the Christian concept GOD, j*) seems conceptually true (at least of a Christian God). This conceptual truth, however, has no metaphysical implications, because it does not entail the existence of God even if we act benevolently.

Evers and Streumer argue that the same holds for the moral fixed points such as a) (2016, p. 4). It is not necessarily true that wrongness is instantiated. This means that it is not the case that any act that satisfies the concept RECREATIONAL SLAUGHTER OF A FELLOW PERSON necessarily satisfies the concept *PRO TANTO* WRONG. So, a) is not conceptually true. However, perhaps a) seems conceptually true because there is a conceptual truth in a)'s neighbourhood:

a*) If anything is *pro tanto* wrong, it is *pro tanto* wrong to engage in the recreational slaughter of a fellow person.

So, the best thing Cuneo and Shafer-Landau can hope for is to establish *conditional* moral fixed points like a*).

This means that the error theorist is not conceptually deficient. She can accept a*) as conceptually true. After all, just like j*), a*) being conceptually true has no metaphysical implications. It does not entail that wrongness is ever instantiated. So, the error theorist can reject a) as conceptually true and accept a*) as conceptually true.

Although the relativist can copy the error theorists' rejection of a), she cannot as easily accept a*) as conceptually true. The metaphysical implications of the moral fixed points are not necessarily the main issue the relativist has with them. The truth of the moral fixed points not depending on the variables of relativity (R_{who} and R_{what}) is their main contention with them. And this issue holds for all moral fixed points, whether conditional or not. For even though it is a conditional claim, a*) would still be an absolute conceptual moral truth, meaning it would be true regardless of

whom it concerns, who utters or evaluates it and their ends, or moral standards, or other moral judgments, etc.¹⁰

Of course the relativist can try to reject the conditional conceptual moral truths as well. But rejecting the conditional moral fixed points leaves an explanatory gap, for they explain why moral fixed points like a) appear to be conceptually true. And, setting aside that concern, the relativist obviously cannot reject a*) for the same reasons she rejects a). So, how might the relativist reject the conditional moral fixed points?

She could follow Street who seems to claim that there are only formal, but no substantive conceptual constraints on what can count as a normative reason (2008, p. 228-229, 243-245). Street, for example, does believe that one “cannot take oneself to have conclusive reason to Y without taking oneself to have reason to take the means to Y” (208, p. 228). But she also seems to believe the

¹⁰ Now, someone might suggest a fully relativist version of the error theoretical argument would not concern a*), but would concern

a**) if anything is absolutely *pro tanto* wrong, it is absolutely *pro tanto* wrong to engage in the recreational slaughter of a fellow person.

And a**), they might argue, is something that the relativist can take to be conceptually true. It would also allow the relativist, just like the error theorist, to explain why people like Cuneo and Shafer-Landau think that a) is conceptually true. They simply mistake it for a**).

The relativist’s argument could go as follows: in our world φ moral claims are at most relatively true. But in the closest possible world where there are absolute moral truths, world φ^{**} , a) would be true. This is because in this world φ , I, the relativist, believe that a) is true, just not conceptually so. So, the possible world where there are absolute moral truths and where a) is absolutely true, world φ^{**} , is closer than any possible world where there are absolute moral truths and a) would not be absolutely true, world φ^{\wedge} . This is because we can get to φ^{**} just by changing which metaethical theory is correct, while, to get to world φ^{\wedge} , we would need to change that and also change which first-order moral claim is true (the negation of a)).

The problem with this is that someone in this world φ who does not believe that a) is true, but who does believe that the negation of a) is true (but not conceptually so) can assert that φ^{\wedge} , a world where the negation of a) is absolutely true, is closer than world φ^{**} for the exact same reasons. In their view, we only need to change which metaethical theory is true to get from this world φ , to get to φ^{\wedge} . While, to get to φ^{**} , we would need to change that and we would also need to change which first-order moral claim is true. So, according to her, the following claim would be conceptually true in our world φ :

a \wedge) if anything is absolutely *pro tanto* permissible, it is absolutely *pro tanto* permissible to engage in the recreational slaughter of a fellow person.

And in our world φ , a world with at most relatively true moral claims, a) is not more true than its negation. So, neither sides’ first-order moral claim, which are used to claim that a**) or a \wedge) is conceptually true, is more privileged than the other.

This means that, if the relativist’s argument works, both a**) and a \wedge) are conceptually true in our world φ . But that leaves us puzzled as to what it means for a claim to be conceptually true. It cannot mean, contrary to what we generally seem to think, that they are also absolutely true. For whether you take a**) or a \wedge) to be true, depends on which other moral convictions you have (a) or the negation of a)), and those in turn depend on your moral standards, or ends, or other normative beliefs, etc. And if the argument does not work, there is of course no ground to think that a**) or a \wedge) is conceptually true in our world φ .

concept REASON does not constrain the content of what can count as a reason outside of this formal constraint and others like it. The relativist might think the same holds for the moral concept WRONG. This would mean she can reject that a*) is conceptually true, since you will need at least some constraints on the content of the concept WRONG to claim that a*) is conceptually true.

Now, even though Street might be right about the concept REASON, you cannot simply say the same holds for the moral concept WRONG. You need to argue for this. Considerations weighing against it are that the moral concept WRONG seems like a less thin concept than the concept of REASON. But, setting aside this worry, rejecting a*) this way also comes at the cost of not being able to account for an even less controversial kind of moral fixed point: *negative* moral fixed points.

Cuneo and Shafer-Landau's moral fixed points concern claims about what is morally wrong, permissible, just, etc. (2014, p. 405). But others, for instance Foot (2002), claim that there are conceptual moral truths about what *cannot be* good, bad, etc. Perhaps Foot's most famous example is the following (2002, p. 118-119):

k) clasping one's hands three times an hour cannot be morally good in itself.

This negative moral fixed point seems far less controversial than the conditional conceptual moral truths, since it does not even entail that some things are potentially morally right, wrong, good, bad, etc. It only rules out the truth of moral claims that do entail that clasping one's hands three times an hour is morally good. But it seems that there need to be some constraints on the content of the concept MORAL GOODNESS for k) to be conceptually true.

Another worry for relativists following Street is that it also is not clear why we clearly recognize the differences between the adjectives 'morally' and 'aesthetically' when talking about why something is good. When we take there to be no substantive conceptual constraints to the concept MORAL GOODNESS and AESTHETIC GOODNESS, it seems that anything can be morally good for any reason, and anything can be aesthetically good for any reason. I might say "X is good, because it is

beautiful” and insist that I am talking about moral goodness. This does not fit well into our normative practices. So, following Street might not be the relativist’s best course of action.

This means that the rejection strategy, contrary to what we might have initially thought, is not the obviously successful strategy. There are costs to pursuing this strategy, such as not being able to take negative moral fixed points to be conceptually true and losing the ability to explain why we understand the difference between (and find it useful to use) adjectives like ‘morally’ and ‘aesthetically’ when using normative concepts.

4.9. Conclusion

So, how can the relativist respond to Cuneo and Shafer-Landau’s assertion that there are conceptual moral truths? I have discussed two general strategies the relativist might employ: reconciliation and rejection.

Given the *prima facie* incompatibility of relativism and conceptual moral truths, the reconciliatory approach might have seemed doomed from the start. But it turns out that demarcating the moral domain in a specific way might allow the relativist to adopt the moral fixed points. However, the major hurdle for this strategy is the cost of spreading absoluteness.

In contrast, the rejection strategy seemed like the most obvious relativist response, but things are more complicated. The relativist might reject the moral fixed points, but has a harder time rejecting the conditional versions of these moral fixed points. And rejecting those seems to imply one has to reject the even more uncontroversial negative moral fixed points as well.

All in all, there does not seem to be a clear-cut way to go for the relativist. What seemed like the most obvious approach, the rejection strategy, is more complicated than it appears at first. And the approach that seemed doomed from the start, the reconciliatory strategy, might in fact work. It all depends on what costs the relativist is willing to incur.

5. The First-Order Argument for Moral Error Theory: The Return of the Impasse

5.1. Introduction

Moral error theory is generally taken to consist of two claims: the conceptual claim states that in making moral judgments we ascribe irreducible and categorically normative properties to acts. The ontological claim states that no such property is ever instantiated. Thus, we are in error when we make moral judgments. The most influential arguments in favour of this theory unsurprisingly concern the alleged metaphysical peculiarities of moral properties (Mackie, 1977; Joyce, 2001; Olson, 2014; Streumer, 2017). But, as Cowie mentions, this is often believed to lead to an impasse between the error theorist and the non-naturalist realist (2020, p. 2). For they simply take a different stance on fundamental metaphysical matters (Olson, 2014, p. 136).

To break this impasse, Cowie (2020) puts forward a new argument for moral error theory. His argument relies on an inconsistency between a set of seemingly true first-order moral judgments in population ethics, and does not concern metaphysics. This, Cowie believes, allows his argument to avoid the impasse (2020, p. 2). Because of this, and other arguably less important reasons, Cowie takes the first-order argument to have a dialectical advantage over the more traditional metaphysical arguments for moral error theory (Cowie, 2020, p. 2-3).

I argue that Cowie's first-order argument for moral error theory trades one impasse for another, and therefore does not have the dialectical edge he claims it has. The argument is straightforward: the first-order argument relies on two assumptions that, when rejected, allow us to deny the first premise of his argument. And your realist¹ or error theorist leanings will affect whether you find these assumptions plausible or not, resulting in an impasse. As a separate argument, I will also argue that a moral error theory supported by first-order argument does worse than other error theorists when it comes to answering the "Now what?" problem for moral error theory. I conclude

¹ By realist I mean any cognitivist that is also a success theorist.

that the first-order argument for moral error theory does not have a dialectical edge over the more traditional metaphysical arguments.

In section 5.2, I will lay out Cowie's first-order argument. In section 5.3, I will discuss the two assumptions in this first-order argument that, when rejected, allow us to deny the first premise of this argument. In section 5.4 and 5.5, I explain how we can reject these assumptions and I will assess those explanations. In section 5.6, I will end by assessing how a moral error theory based on a first-order argument fares when it comes to the "Now what?" problem for moral error theory.

5.2. The First-Order Argument for Moral Error Theory

Cowie argues that an inconsistency between highly intuitive first-order moral claims discredits the entire domain of first-order morality, and thus leads to an error theory. To establish that there is such a first-order inconsistency, Cowie refers to Arrhenius's (2009) impossibility proof. In this work, Arrhenius argues that five axiological claims in population ethics that individually seem intuitively true are in fact inconsistent. The exact details of the inconsistency are fairly unimportant for the purpose of this chapter and the first-order argument. What is important is that there seems to be a genuine first-order ethical inconsistency.

With that being said, this is Cowie's first-order argument (2020, p. 6):

- P1:** If the propositions that constitute the impossibility proof aren't all true, then no moral propositions are true.
- P2:** They aren't all true (they're inconsistent).
- C1:** (P1, P2) No moral propositions are true.
- P3:** When we make moral judgments we're trying to state true moral propositions.
- C2:** (C1, P3) When we make moral judgments, we're thereby in error.

The best way to defend **P1**, according to Cowie, is to argue that the axiological propositions of the impossibility proof are all conceptually true (2020, p. 8-9). He argues as follows (2020, p. 8-9):

a concept that grounds a contradiction cannot be used to assert a true proposition. For instance, we cannot assert any true propositions with the concept ROUND SQUARE.² But we do need moral concepts to assert moral propositions. So, if moral concepts ground a contradiction then no moral propositions are true. Now assume that the axiological propositions of the impossibility proof are inconsistent. This means that, if these axiological propositions are all conceptually true, the moral concept GOODNESS grounds a contradiction.³ This would in turn mean that we cannot assert any true proposition using this concept, which entails that all the propositions of the impossibility proof and all other goodness-judgments are false. So, **P1** is true if all the axiological propositions of the impossibility proof are conceptually true.⁴

To address this “if”, Cowie points to the work of Cuneo and Shafer-Landau (2014). They defend a particular version of non-naturalist realism by arguing that certain substantive moral propositions, the moral fixed points, are conceptually true. In their paper, Cuneo and Shafer-Landau identify several marks of conceptual truth (2014, p. 407-408). Cowie claims that the axiological propositions seem to bear these marks, making them good candidates for being conceptually true (2020, p. 12-13).

So, with **P1** supported and **P2** assumed to be true based on Arrhenius’s work, **C1** follows. If we then also assume the truth of cognitivism, and thus the truth of **P3**, **C2**, i.e. the error theory, follows (Cowie, p. 6).

As mentioned, Cowie takes the first-order argument for moral error theory to have some advantages over more traditional, metaphysical arguments for moral error theory. These advantages are mostly dialectical (Cowie, 2020, p. 2-3). One, and arguably the most important, of these

² To be a bit more precise, we cannot assert any positive, substantive, and true propositions with this concept. We can say true negative things with the concept ROUND SQUARE, such as “There are no round squares”. And we can of course also say trivially true things, such as “A round square is a round square”.

³ This is true, but it is not the only way to get this result. Not all axiological propositions need to be conceptually true for the moral concept GOODNESS to ground a contradiction. It also grounds a contradiction when at least one of them is conceptually true, while the others are simply true. I will discuss this further in section 5.5.1.

⁴ Strictly speaking this is false. All that the axiological propositions being conceptually true shows is that the moral concept GOODNESS grounds an inconsistency. This means we cannot make any true moral goodness judgments, but that need not affect, for instance, deontic moral judgments. Cowie acknowledges this, but also notes that he thinks that we can raise a similar argument for deontic moral concepts (2020, p. 5-6).

advantages is that it does not lead to an impasse on fundamental metaphysical positions. Cowie, for instance, mentions Olson's (2014) metaphysical argument for moral error theory, which takes the claim that there are no irreducible normative relations as "metaphysical bedrock" (p. 136). This then leads to an impasse with non-naturalist realists who do believe there are such relations. And that is exactly what Cowie believes the first-order argument can avoid by not appealing to anything other than inconsistent first-order ethical propositions.

In the following sections, I will argue that this first-order argument does not avoid an impasse and therefore does not have a dialectical edge.

5.3. Conceptual Competence

As mentioned in the introduction, this first-order argument relies on two assumptions that, when either of them is rejected, enable us to argue that **P1** goes unsupported. This is because these assumptions figure in the conceptual argument for **P1**. These assumptions are needed to make the claim that the moral concept **GOODNESS** is incoherent, because they are needed to establish the conceptual truth of the axiological propositions. And if **P1** is not supported, the first-order argument does not get off the ground.

The two assumptions are the following:

1) that we are fully conceptually competent with the moral concept goodness

and,

2) that our conceptual competence entails that our beliefs about the implications of this concept are true.⁵

Imagine that 1) is an unwarranted assumption and we reject it. That is, imagine that we are not fully conceptually competent with the moral concept GOODNESS. And further imagine that we are unaware of this. Then we might mistakenly take one (or more) of the axiological propositions to be conceptually true, being led astray by our imperfect competence. This would mean that the moral concept GOODNESS is not incoherent, which means that the conceptual argument for **P1** fails.

But even if 1) is warranted and we accept it, that is we take ourselves to be conceptually competent, the first-order argument does not necessarily follow. Because imagine that 2) is an unwarranted assumption and we reject it. Based on work by Eklund (2002), we can then argue that it is precisely our conceptual competence that leads us astray. It disposes us to take all axiological propositions to be true, even though one or more are false. This would mean that the moral concept GOODNESS is not incoherent, which once again means that the conceptual argument for **P1** fails.

I will discuss this second option first.

5.4. Rejecting the Second Assumption

Eklund (2002) argues that our conceptual competence with the concept TRUE leads us to take a false assumption in the liar paradox to be true.

The liar paradox goes as follows. Take:

a) "a)" is false.

⁵ One might frown at these assumptions being pulled apart. Is correctly judging what is implied by a concept not part of what it is to be fully conceptually competent? So, if you take 1) to be true, you must also take 2) to be true. In section 4.1, I will explain why this is not necessarily the case.

If we assume that a) is true, a) turns out to be false. After all, a) says that a) is false. But if we assume that a) is false, a) turns out true. After all, a) says that a) is false. So, either assumption leads to a contradiction.

Eklund claims that it is widely agreed that this paradox is brought about by the reliance of some steps in the paradox “on the assumption that all instances of the disquotation schema are true” (2002, p. 259). That is, some steps of the liar paradox presuppose that ““p” is true if and only if p” is true in all instances. So, to resolve the paradox, we must deny that the disquotation schema is true in all instances. But, as Eklund mentions, this seems odd, because “[p]redicating “true” of a sentence is, or is intended to be, to say that what the sentence says is the case actually is the case” (2002, p. 259). The disquotation schema has considerable intuitive appeal. How can that be, if it is false?

Eklund (2002, p. 259) argues that our competence with the concept TRUE might dispose us, barring contrary evidence, to take the disquotation schema to be true even though it is false. That would explain why we do think the liar paradox appears to be a genuine paradox, even if it is not. The idea is that “the meaning of “true” can require that [p] and “[p] is true” come out equivalent without [p] and “[p] is true” actually being equivalent” (Eklund, 2002, p. 259). Thus, our competence with the concept TRUE leads us astray about the falsity of the claim that the disquotation schema is true in all instances. And that makes us mistakenly think that the liar paradox is a genuine paradox.

Might we say something similar about the impossibility proof? That is, might our conceptual competence with the moral concept GOODNESS dispose us to take all axiological propositions to be true, barring contrary evidence, even though at least one of them is in fact false? Well, if it did, the true axiological propositions would not be inconsistent, and the moral concept GOODNESS would not be incoherent. This would mean that the conceptual argument for **P1** fails.

5.4.1. Assessment

So, what should we think about this Eklundian argument? A lot of things are still in need of explaining for this to be a viable argument.

The first thing that might come to mind is that it seems impossible to be both conceptually competent with a concept and also make mistakes in judging what follows from that concept. Is knowing what follows from a concept not what it means to be conceptually competent? Eklund discusses this worry (2002, p. 260-266). He agrees that all popular accounts of conceptual competence do think that this is, at least in part, what constitutes conceptual competence (Eklund, 2002, p. 260). But he also claims that,

at least some of these accounts can be slightly modified so as to be compatible [with the possibility of our conceptual competence disposing us to accept false propositions] (...); and that nothing in the general rationale underlying these accounts demands that they not be so modified.

(Eklund, 2002, p. 260)

A recounting and assessment of his argumentation are beyond the scope of this paper. So, let us say that the Eklundian argument is conditional on the soundness of Eklund's argumentation here.

Now, there are at least two other worries to deal with: First, Eklund's argument relies on a very particular reason why the liar paradox is not a genuine paradox: some steps of the paradox rely on the disquotation schema being true in all instances, even though it is not. This argument does not translate to the conceptual argument for **P1** directly. You will therefore need a specific reason why one of the axiological propositions is false. Once you find that, you can, like Eklund, claim that our conceptual competence disposes us to take this false axiological proposition as true.

Second, Eklund does not think we are led astray about all steps of the liar paradox. We are only led astray about those steps that rely on the disquotation schema being true in all instances. For an Eklundian argument to be at least plausible in the case of the impossibility proof, it seems the same should hold. But here someone like Cowie might argue that pinpointing the specific proposition

about which we are led astray is quite difficult in the axiological case. Especially since he argues that it is not clear which proposition would be easiest to deny (Cowie, 2020, p. 14).

To deal with the second worry first, as I will discuss in section 5.5.4., there is a striking difference in the complexity between the axiological propositions. Some are quite easy to grasp, others even academic philosophers struggle to grasp.⁶ It only seems natural to take one of those highly complex axiological propositions to be the one we are misled about.

As for the first worry, this one is potentially far more difficult to set aside. It is quite difficult to give a specific reason why one of the axiological propositions is false. But it does not seem in principle impossible. So, this argument against the incoherence of the moral concept GOODNESS is also conditional on there being such a specific reason.

Where does this leave us? Well, one motivation for Eklund to think that our conceptual competence with the concept TRUE leads us astray about a premise of the liar paradox, is that it is the logical conclusion to draw (2002, p. 253-254). He says,

Upon noting that your semantic intuitions lead you to accept an inconsistent corpus of statements, you can refuse to take these intuitions at face value, rejecting one or more of them as non-veridical, without thereby manifesting a lack of [conceptual] (...) competence. (Rather the opposite: you positively manifest logical skills.)

(Eklund, 2000, p. 253-254)

Now, a first-order error theorist might say that this is fine when it comes to the liar paradox. After all, if the liar paradox is a genuine paradox, our concept TRUE is incoherent, because a) would be true and false at the same time. And the incoherence of our concept TRUE is surely something we would want to avoid. It is a very useful concept, for instance, when it comes to persuading others. But it of course also helps in articulating our sense that some things are actually the case. So, we reject a part of the liar paradox. But this first-order error theorist might argue that a parallel conclusion is of no real concern in the case of the moral concept GOODNESS. That is, the potential

⁶ Yes, this is a confession.

incoherence of the moral concept GOODNESS should not motivate us to believe that our conceptual competence disposes us to take a false axiological proposition to be true.

But such an argument by the first-order error theorist only seems to hold if one is already inclined to accept that we cannot make any true goodness judgments. A moral realist, who typically does think we can make true goodness-judgments, and who does think that the moral concept GOODNESS is useful, would take the potential incoherence of the moral concept GOODNESS as proper motivation to embrace an Eklundian argument about the axiological propositions.

This means that we arrive at an impasse. One's prior inclinations on moral realism and error theory will affect whether one finds an Eklundian argument plausible or not. This means it affects whether one takes **P1** to be supported or not. This of course in turns means that it affects whether one thinks the first-order argument for moral error theory succeeds or not.

5.5. Rejecting the First Assumption

As mentioned, we can reach a similar impasse if we reject 1): the assumption that we are fully conceptually competent with the moral concept GOODNESS.

Let me start by clearing up two things first: one, not being fully conceptually competent, or being imperfectly competent, is different from being completely conceptually incompetent. Competency with concepts comes in degrees. Two, there are at least two ways in which one can be imperfectly competent. You can fail to grasp part of the content of the concept. Or, even if you do fully grasp the content of the concept, you can fail to grasp all that is entailed by the concept.⁷ Unless discussed explicitly, this second point does not matter for what follows.

With that out of the way, we can look at the implications of rejecting 1). The conceptual argument claims that the axiological propositions of the impossibility proof are conceptually true. For

⁷ Here I do take assumption 2) to be part of assumption 1). That is, I take it to be part of full conceptual competence that you also know what is entailed by the concept. This might seem odd, but it is perfectly fine. I only pulled 1) and 2) apart, because that is what the Eklundian argument requires. We are now no longer taking this argument into consideration, so we can collapse 1) and 2).

us to recognize them as such, we would need to be competent (enough) with the moral concept GOODNESS. But what if we are not? If we are not fully competent with the moral concept GOODNESS, we might mistakenly take the axiological propositions of the impossibility proof to be conceptually true, even if they are not all conceptually true.

This can be the case when our imperfect competence ensures that we do not have access to information that would entail the falsity of at least one of the axiological propositions. This might, for instance, happen through our failing to grasp an exclusion-clause within the moral concept GOODNESS. That is, the moral concept GOODNESS might have a clause that rules out the truth of one of the axiological propositions. To illustrate how this might work, consider the following example: we might take the moral concept GOODNESS to entail that we must maximize welfare. And we might therefore think that our moral concept GOODNESS entails that pushing a large person off a bridge, thereby killing them, to stop a trolley from killing five people is good. That is, if we also stipulate that each person's death in this scenario would destroy an equal amount of welfare.⁸ But if we are only imperfectly competent with the moral concept GOODNESS, we might not grasp, for instance, an autonomy-clause that is part of the moral concept GOODNESS. This clause might say that we have to take a person's autonomy into consideration to judge whether an act is good.⁹ This clause then rules out the possibility that it is conceptually true that pushing the large person is good. Because to determine whether this act is good we need *a posteriori* information about how much autonomy we are destroying and saving, were we to push the large person off the bridge.

Something similar might hold in the axiological case. If you are not fully conceptually competent with the moral concept GOODNESS, you might fail to grasp an exclusion-clause that excludes at least one axiological proposition from being conceptually true. We might therefore falsely believe that all axiological propositions are in fact conceptually true, even though at least one of them is not.

⁸ And if we stipulate that the only actions available to us are to push or not to push the large person off the bridge.

⁹ One might argue that respect for autonomy is already part of what constitutes welfare. That might be so. But let us assume, for purposes of illustration, that it is not.

If the axiological propositions are not all conceptually true, the moral concept GOODNESS does not ground an inconsistency. This concept is therefore not incoherent, which means that the conceptual argument for **P1** fails.

5.5.1. Assessment: Inconsistency

The first thing that first-order error theorists like Cowie might note is that this imperfect competence argument only gets rid of a *conceptual* inconsistency (2020, p. 14). But the axiological propositions will still be *inconsistent* if the non-conceptually true one is still true. To avoid this inconsistency we would have to completely deny the truth of the non-conceptually true axiological proposition. But, Cowie argues (2020, p. 14), although one might be able to argue that one of the axiological propositions is not conceptually true, it seems a lot more difficult to deny that this proposition is not true at all, given our strong intuitive sense that all axiological propositions are individually true. So, a non-conceptual inconsistency remains.

Now, it is not immediately obvious what first-order error theorists can get from this non-conceptual inconsistency. As mentioned, **P1** is supported with the claim that the moral concept GOODNESS grounds inconsistency and thus is incoherent. If so, this means that we cannot assert true propositions with it. The incoherence of the moral concept GOODNESS seems integral to this conceptual argument. As Cowie says, “the conceptual argument [to support **P1**] works by demonstrating an inconsistency in the concept of goodness itself” (2020, p. 9). And,

if the essence of a concept grounds a contradiction, then anyone who makes a judgment with that concept has failed to say something true (compare someone who makes judgments with the concept of a round square). This gives us the first premise of the [first-order argument for moral error theory].

(Cowie, 2020, p. 9)

So, it seems that only a conceptual inconsistency can entail the incoherence of the moral concept GOODNESS. So, why does Cowie (2020, p. 14) think that the remaining non-conceptual inconsistency still leads to a moral error theory?

Cowie does not give a clear answer to this question. What he and other first-order error theorists should say, however, is that the non-conceptual inconsistency still entails the incoherence of the moral concept GOODNESS. In that case, first-order error theorists can still use the conceptual argument for **P1**. Let me explain how they might argue for this.

Take A, B, and C to be propositions about concept φ . Also assume that A, B, and C are inconsistent. If A and B are true, C must be false. If A and C are true, B must be false. And if B and C are true, A must be false.¹⁰ Now imagine that A and B are conceptually true, while C is just true. What does this say about concept φ ? Well, on the one hand A and B being conceptually true excludes, in virtue of the content of concept φ , C from being true. On the other hand, C being just true means that concept φ must at least allow for the truth of C. Thus, concept φ both does and does not allow for C to be true. Concept φ is therefore an incoherent concept.

The first-order error theorist can say the exact same thing about the axiological propositions of the impossibility proof. As long as they are all true, and at least one of them is conceptually true, the moral concept GOODNESS must be incoherent. Its incoherence lies in which propositions can be true. Granted, this incoherence is different from the incoherence of the moral concept GOODNESS if all axiological propositions are conceptually true. That incoherence lies in the truths that the moral concept GOODNESS entails. But it would be incoherent nonetheless. Thus, the non-conceptual inconsistency can still entail the incoherence of the moral concept GOODNESS, which the first-order error theorist can then use to argue for **P1**.

This means we must refine the imperfect competence argument. To avoid an error theory we must not only argue that one axiological proposition is not conceptually true, but we must argue that it is not true at all. Cowie might think this spells doom for this argument, because

¹⁰ And if A is true, B and C cannot both be true, if B is true, A and C cannot both be true, and if C is true, A and B cannot both be true.

[t]he propositions that constitute the axiological impossibility theorem are extremely difficult to deny. They are sufficiently hard to deny that, abandoning any one of them may, for a rational, sincere and *conceptually competent person* [emphasis added], simply not be an option.

(Cowie, 2020, p. 14)

But we are not considering someone who is fully conceptually competent. Does that make completely denying one of the axiological propositions easier?

The answer seems to be yes. There are at least two ways in which it does: first, we can argue that the exclusion-clause in the moral concept GOODNESS, which we fail to grasp because of our imperfect competence, does not only prevent at least one axiological proposition from being conceptually true. It prevents at least one proposition from being true altogether.

Recall the trolley example. If, in that example, the autonomy-clause of the moral concept GOODNESS were stricter, it might entail the falsity of the claim that pushing the large person in front of the trolley is good. That is, if this autonomy-clause says that maximizing welfare may not come at the cost of violating a person's autonomy, instead of only saying that we must take autonomy into consideration, pushing the person cannot be good, as a conceptual matter. We can argue that the exclusion-clause in the axiological case exhibits a similar level of strictness. This would mean that one of the axiological propositions is false, that the moral concept GOODNESS is therefore not incoherent, and thus that the conceptual argument for **P1** fails.

The second way to argue for this is to point out that the strong intuitive sense that the axiological propositions are true cannot, in every instance, come from them being conceptually true. The imperfect competence argument claims that one of them is not conceptually true after all. So, what remains? Is it just a moral intuition that they are true? Is that enough? It does not seem to be. We can imagine a Moorean argument that asks: what is more likely, that our moral intuition is off about one particular (highly complex) claim about moral goodness, or that we cannot make any true moral judgments, including judgments like, "slavery is wrong", "eradicating polio was good", etc.?

This response can be bolstered if we consider that not being fully conceptually competent might affect our intuitive sense of truth, especially when we are not aware of this imperfect competence. Imagine for instance, someone who does not grasp some maximizing element in the concept of goodness. Might they not lack the intuitive sense that the following axiological proposition is true?

The General Non-Extreme Priority Condition: For any welfare level A and any population X, there is a number n of lives such that a population consisting of the X-lives, n lives with very high welfare, and one life with welfare A, is at least as good as a population consisting of the X-lives, n lives with very low positive welfare, and one life with welfare slightly above A, other things being equal.

(Cowie, 2020, p. 4)

The answer seems to be yes. After all, if you do not think that more welfare overall is not worse than less welfare overall, then you might just think that the population that includes the person with the highest level of welfare is best. You would then lack the intuitive sense that *the general non-extreme priority condition* is true. This intuitive sense that all the axiological propositions are true, might therefore not be all that Cowie makes it out to be. Especially when it comes to the one (or more) axiological propositions that is (are) true, but not conceptually true.

So, we can argue that one of the axiological propositions is not true altogether. This undoes all inconsistency, not just the conceptual one. This means that the moral concept GOODNESS is not incoherent, defeating the conceptual argument for **P1**.

5.5.2. Assessment: How Can We Be Led Astray?

This does seem to leave us with an oddity. How can it be that the false axiological propositions seemed true in the first place? Our imperfect conceptual competence can help explain this.

Let me first illustrate this with an example. Take the subject of mathematics. Imagine a child who fully understands what the numbers 1 through 10 mean. Further imagine that she is fully

competent with the concepts ADDITION and SUBTRACTION. This child therefore correctly thinks that the following claims are true:

$$\text{b) } 3 + 3 = 6$$

and

$$\text{c) } 6 - 3 = 3$$

But now imagine that she is not fully competent with the concept MULTIPLICATION, but does not realise this. She might correctly think that

$$\text{d) } 3 \times 3 = 9$$

is true. But she might also falsely think that

$$\text{e) } -3 \times -3 = -9$$

is true. This does not seem unthinkable. The fact that multiplying two negative numbers results in a positive number is not immediately obvious.

And her sense that e) is true might be very strong. Just imagine that she has a classmate that tells her that e) is actually false, and that

$$\text{f) } -3 \times -3 = 9$$

is true. This might very well baffle her. She might think that this classmate uses different mathematical concepts. In other words, she might even think that e) bears some of the marks of conceptual truth mentioned by Cuneo and Shafer-Landau (2014, p. 407-408). So, the child in

question believes that something is true, even though it is actually false, because of her (unacknowledged) imperfect competence with the mathematical concept MULTIPLICATION.

So, how might this work in the axiological case? Well, as mentioned, because of our imperfect conceptual competence we might just be failing to grasp part of the moral concept GOODNESS, such as, an exclusion-clause that entails that one of the axiological propositions is not conceptually true. But if we do not notice this, if we believe that we are conceptually competent, we might take all axiological propositions to be conceptually true even though we mistakenly take one of them to be conceptually true. So, we might very well take a false axiological proposition to be conceptually true.

5.5.3. Assessment: How Can We Recognize Other Conceptual Truths?

First-order error theorists might think that that is all well and good, but might also argue that there surely are downsides to considering ourselves to be imperfectly conceptually competent. Does it not also cast doubt on Cuneo and Shafer-Landau's moral fixed points? That is, does it not entail that we may be mistaken to regard Cuneo and Shafer-Landau's highly intuitive claim (2014, p. 405),

g) it is *pro tanto* wrong to engage in the recreational slaughter of a fellow person

as conceptually true?

If the answer is yes, it will depend on your beliefs about Cuneo and Shafer-Landau's moral fixed points whether this bothers you or not. If you do not feel their intuitive pull, or do feel it but do not think it is explained by their conceptual truth, you need not be bothered by no longer being able to regard claims like g) as conceptually true. You would, of course, need to come up with some reason why you remain unconvinced by the arguments for these supposed conceptual moral truths. But I will not be doing that here partly because I am somewhat sympathetic to Cuneo and Shafer-Landau's work, and partly because it does not fit the aims of this paper.

So, what if you do feel the intuitive pull of g), and think this is explained by it being conceptually true? Does this mean that the imperfect competency argument is off the table? First thing to note is that if you want to maintain that g) is conceptually true you must reject the moral error theory, and must therefore reject the first-order argument for this theory. After all, any moral error theory entails that there are no moral truths, let alone conceptual moral truths. But secondly, not being fully conceptually competent does not have to mean that you cannot recognize anything that is entailed by that concept.

Return to the mathematical example. The child's imperfect competence with the concept MULTIPLICATION makes her accept

$$e) -3 \times -3 = -9$$

as true. Which is of course false. But this does not necessarily entail that she cannot recognize the truth of

$$d) 3 \times 3 = 9.$$

Her imperfect grasp of the mathematical concept MULTIPLICATION does not mean she can only make false claims with it. She can make some true claims, but it also leads her to make some false claims such as e). Likewise, the fact that we are not fully conceptually competent with moral concepts need not mean we cannot grasp any conceptual moral truths. It is possible to recognize some truths entailed by concepts even though you are only imperfectly competent with a concept.

5.5.4. Assessment: What About the Axiological Propositions Then?

But now first-order error theorists might turn this around and ask: if we can recognize the conceptual truth of g) even though we are not fully conceptually competent, why would we fail to correctly

recognize the conceptual truth of the axiological propositions of the impossibility proof?

One factor that is likely to be of influence is the complexity of the proposition in question. It seems reasonable to think that the more complex the proposition, the better one's grasp of the concept must be to judge whether it is conceptually true. For instance, the first mathematical truths we grasp will be from easy tables, not the equations that figure in Einstein's theory of general relativity.

It seems that Cuneo and Shafer-Landau's list of moral fixed points, including g), are all significantly less complex propositions than some of the axiological propositions. Granted, the complexity of the first axiological proposition

The Egalitarian Dominance Condition: If population A is a perfectly equal population of the same size as population B, and every person in A has higher welfare than every person in B, then A is better than B, other things being equal

(Cowie, 2020, p. 4)

seems on a par with some of Cuneo and Shafer-Landau's moral fixed points. But the complexity of some of the other axiological propositions is obviously far greater. For example,

The Non-Elitism Condition: For any triplet of welfare levels A, B, and C, A slightly higher than B, and B higher than C, and for any one-life population A with welfare A, there is a population C with welfare C, and a population B of the same size as AUC and with welfare B, such that for any population X consisting of lives with welfare ranging from C to A, BUX is at least as good as AUCUX, other things being equal

(Cowie, 2020, p. 5)

is surely far more complex than the moral fixed point

g) it is *pro tanto* wrong to engage in the recreational slaughter of a fellow person.

There might therefore be good reason to think that one can recognize Cuneo and Shafer-Landau's moral fixed points as conceptually true, but that the same does not hold for every one of the

axiological propositions. Some of them are far more complex than the moral fixed points. They therefore require a better grasp of the moral concept GOODNESS than we have to judge whether they are conceptually true or not.

The first-order error theorist might respond in two ways: first she might say that this is mistaken. She might agree that we do need a better grasp of a concept to determine whether more complex propositions follow from it. But the imperfect competency argument does not argue that we mistakenly take one of the axiological propositions to be conceptually true because it is too complex. According to this argument, we make this mistake because we fail to obtain information that entails the falsity of one axiological proposition, for example, by not grasping an exclusion-clause. And we can just as easily fail to grasp an exclusion-clause for a simple proposition as for a complex one.

This is true. But the argument about complexity does not have to rely on this conception of imperfect conceptual competence. Recall the two ways in which we can understand not being fully conceptually competent. The most obvious way for someone to fail to grasp an exclusion-clause is to just not know that this clause is part of the concept. If we talk about this kind of failure, the complexity of the proposition we are mistaken about does not matter. However, we can also fail to be fully conceptually competent by not grasping all that is entailed by the concept. In this case, we might fail recognise that a clause we do grasp is an exclusion-clause for a particular proposition. This second way to fail to grasp an exclusion-clause surely is affected by the complexity of that proposition. After all, it seems that the more complex the proposition, the more difficult it is to understand what can keep it from being conceptually true. So, for highly complex propositions it would be more difficult to know whether a specific clause excludes its truth. Thus, the difference in complexity between g) and *the non-elitism condition* can entail that we need a better grasp of our moral concepts to correctly judge whether the latter is conceptually true or not, while this does not hold for the former. We just have to understand conceptual competence in this second way.

The second response that first-order error theorists like Cowie will surely give is that this

difference in complexity is “not deep” and merely the “result of a difference in parsing” (Cowie, 2020, p. 10). Cowie argues that Cuneo and Shafer-Landau’s moral fixed points, like g), can be rewritten to match the complexity of some of the axiological propositions:

For any person A, and any group of people B1 to Bn, and any action φ that consists in the deliberate killing of B1 to Bn by A, there is a strength of desire, L, such that it is pro tanto wrong for A to φ in order to satisfy a desire with strength less than or equal to L.

(Cowie, 2020, p. 11)

But this misses the mark. Granted, we can write g) in a complex way. This should not come as a surprise. Philosophers are often in the business of overcomplicating simple claims. But the issue is not whether we can write simple propositions complexly. The issue is whether we can write complex propositions simply. After all, the claim is that the difference in complexity entails that someone with imperfect competence with moral concepts can recognize g) as conceptually true, but cannot correctly determine whether *the non-elitism condition* is too. Being able to write g) as complexly as *the non-elitism condition* does not undermine this claim. Being able to write *the non-elitism condition* as plainly as g) does.

So, can propositions like *the non-elitism condition* be written in a less complex way without losing any of its meaning? This seems a lot more difficult than writing g) in a more complex way.¹¹ I am not claiming it is impossible to write a simple version of *the non-elitism condition*, but it is a significant challenge, and I do not see a way of doing it. This suggests that the complexity of *the non-elitism condition* is inherent to the proposition. That is, the information the proposition is trying to capture is highly complex and cannot be properly communicated without some level of complexity that far exceeds the complexity of g).

This means that the difference in complexity between the moral fixed points and some of the axiological propositions is not merely the “result of a difference in parsing”. It is a robust difference. This in turn means that you can use this difference in complexity to defend the position that being

¹¹ Again, this should not be surprising given that philosophers often fail to write simply about complex matters.

able to correctly recognize g) as conceptually true does not mean you can correctly recognize the truth of the axiological propositions.

5.5.5. Assessment: What Is Your Evidence?

At this point the first-order error theorist might say: that is an interesting story, but you have not given us any evidence for it. Why should we believe that we are not fully conceptually competent with the moral concept GOODNESS such that we mistakenly take all axiological propositions to be conceptually true?

The answer is: the exact same reason for why we should believe that the moral concept involved in these propositions grounds inconsistency: the apparent inconsistency of the axiological propositions. That is, first-order error theorists take the seeming inconsistency of the axiological proposition to be evidence for the idea that the moral concept GOODNESS is incoherent. Which is not surprising for someone who favours moral error theory. But you can just as well think that the moral concept GOODNESS is coherent, and therefore take this inconsistency to be evidence of our failure to fully grasp this concept. A person who leans more towards moral realism is likely to favour this second reading of the evidence. Thus, we, once again, end up at an impasse between error theorists and moral realists.

5.6. Now What?

So, the first-order argument from moral error theory fails to deliver on what is potentially its most important promise: a way around an impasse between realists and anti-realists. But there is another reason to think that the first-order argument does not have a dialectical edge over the traditional metaphysical argument for moral error theory. A first-order moral error theory does worse when it comes to answer the “Now what?” problem.

The “Now what” problem is a problem particular to the error theory. The problem assumes

that moral error theory is true and that all our moral judgments are erroneous. The question then is: now what? Do we just stop making moral judgments and give up on our moral practice and its benefits? That is, should we be abolitionists about morality? Or are there other options?

To answer this last question: yes, there are. Many error theorists are not comfortable with abolishing our moral practice and have articulated ways of getting around it. The alternatives in the literature are: conservatism, fictionalism, and substitutionism (Lutz, 2013).

In this section, I will not discuss whether these are plausible alternatives on their own merit.¹² I will only argue that these alternatives to abolitionism do not seem available to first-order error theorists. Let us start with conservatism.

5.6.1. Conservatism

The conservationist claims that a metaethical theory like moral error theory need not affect our first-order moral views. Mackie, for example, claims that moral error theory is “not merely distinct but completely independent” (1977, p. 16) from first-order ethics. Olson (2014) argues for a similar compartmentalization of our metaethical beliefs. He says that “it is a psychologically familiar fact that we sometimes temporarily believe things we, in more reflective and detached contexts, are disposed to disbelieve” (Olson, 2014, p. 192). So, even though we generally think that recreational slaughter is morally wrong, when we start to think about metaethics we enter a “more reflective and detached context”, which leads us to the conclusion that it is not morally wrong, because nothing is morally wrong.

The main theme to all these conservationist answers is that in some way or other first-order ethics is isolated from the implication of our metaethical theorising. They are imagined as being two separate realms. And because of that we can be moral error theorists who believe that there are no true moral judgments, while maintaining our first-order ethical practice of attempting to make true moral judgments.

¹² See (Lutz, 2013) for this discussion.

This is exactly why this conservationist approach is not available to first-order error theorists. Precisely the thing that makes their position unique, that it is supported by an argument based on first-order ethical claims, means that compartmentalizing metaethics and first-order ethics does nothing for the first-order error theorist. It does not isolate the source of moral error theory from its first-order consequences, like it does for the traditional metaphysical arguments for the error theory.

Now, first-order error theorists can of course suggest that we can isolate the source of her error theory from its consequences for first-order ethics. All we have to do is compartmentalize at least one axiological proposition of the impossibility proof from our first-order ethics. That would rid our first-order ethics of inconsistency that leads to an error theory.

But there seem to be some obstacles for this approach. To name two: one, on what grounds can we compartmentalize one of the axiological propositions? The first-order error theorist cannot rely on the arguments for the standard form that conservationism takes. Remember that Olson argues that we generally take there be a significant difference in how we come to metaethical judgments and first-order ethical judgments (2014, p. 192). This difference, he then argues, allows for a compartmentalization of these areas of thought. This argument obviously does not work for the first-order error theorist, since there is no such difference in judgment formation between the different axiological propositions. Two, which one of the axiological propositions should we compartmentalize? Note that we assume that the first-order error theory is true, and that, according to Cowie (2020, p. 4), all axiological propositions seem intuitively plausible.

These seem like considerable obstacles. The conservationist approach to the “Now what?” problem therefore does not seem promising for first-order error theorists.

5.6.2. Fictionalism

So, can first-order error theorists become fictionalists instead? Joyce (2001, p. 206-231), a fictionalist, argues that we can avoid the first-order consequences of moral error theory if we change what we do when we make a moral judgment. Error theorists, Joyce included, believe that when we

make a moral judgment we assert a moral proposition and thereby ascribe a moral property to an act or person. But what we should do instead, according to Joyce, is only pretend to assert a moral proposition when we make a moral judgment. If we do that, we do not ascribe properties that do not exist, but we can, according to Joyce, maintain our moral practice and with it the practical benefits it offers.

Now, this might work if the problem with our moral practice is that we ascribe non-existent properties. Only pretending to ascribe these properties will mean that you avoid this problematic act. But, according to first-order error theorists, the ascription of moral properties is not where we go wrong when we make moral judgments. We go wrong in using an incoherent moral concept when we make moral judgments. So, only pretending to ascribe moral properties will not help us avoid the problematic part of our moral judgments.

First-order error theorists can of course reply that they can tailor the fictionalism to their account. That is, what if we do not only pretend that we ascribe moral properties, but we also pretend to use moral concepts. The problem with this is that it then becomes unclear whether what we are asserting is moral at all. According to Cowie (2020, p. 9), we need moral concepts to make moral judgments. If we only pretend to use moral concepts it seems we cannot make moral judgments, not even pretend ones. That is, if I say “stealing is morally wrong”, but I only pretend to use the moral concept *WRONG*, what am I actually saying? This is unclear. This also makes it unclear whether judgments that only pretend to use moral concepts actually have the benefits that we take normal moral judgments to have. That is, how can a pretended use of a moral concept help guide our actions, motivate us, etc.? By only pretending to use moral concepts, we are not making judgments that are moral, after all. Fictionalism therefore does not seem to fit the first-order error theory well.

5.6.3. Substitutionism

Might first-order error theorists then opt for substitutionism instead? The answer seems to be ‘No’ again. Remember that error theorists make a conceptual claim: in making moral judgments we

ascribe irreducible and categorically normative properties to acts. Substitutionism proposes that we find a fully-committed attitude that is as close to this conception of moral judgments without it being problematic (Lutz, 2013, p. 361-366). That is, according to the error theorist, the problem with our moral judgments is that they do not refer. There are no irreducible and categorically normative properties. The substitutionist wants to find a non-moral way of asserting “murder is wrong” that does refer. So, instead of having the belief that murder is morally wrong, the substitutionist might say she has the belief that she disapproves of murder (Lutz, 2013, p. 365-366). This substituted belief can refer and should then, at least partly, be able to secure the practical benefits we take moral judgments to have.

If the first-order argument for the error theory is to be a standalone argument, it does not make sense for first-order error theorists to pursue this kind of substitutionism. For the first-order argument does not concern itself with the error theorist’s conceptual claim and its implications. Moreover, substituting moral claims with, for instance, claims about approval or disapproval does nothing to resolve the inconsistency of the axiological propositions. If we did make this substitution, one might wonder why this would not lead to an error theory about approval and disapproval as well.

But a different kind of substitutionism is available to first-order error theorists. According to them, it is the incoherence of the moral concept GOODNESS that leads to an error theory. So, what if we substitute this moral concept GOODNESS with a different concept GOODNESS*, which, unlike its asteriskless counterpart, does not ground an inconsistency? When we adopt this concept GOODNESS*, we can make true moral* judgments. And, the first-order error theorist might argue that, just like our moral practice, our moral* practice has important practical benefits.

Interestingly though, Cowie discusses this kind of substitutionism and rejects it. The reasons why he rejects it are not entirely clear (2020, p. 15-16). And he acknowledges this (Cowie, 2020, p. 16): he claims that to clearly articulate why this solution is unsatisfactory is difficult. But he might be understood as thinking it leads to arbitrariness. It seems that the only standard we have for

GOODNESS* is that it must not entail that all axiological propositions of the impossibility proof are true. But there are presumably many ways of filling out the concept GOODNESS* that clear this bar, and we have no way to determine which one we should use. After all, GOODNESS* might be good* to use, but not good** to use. GOODNESS** might be good** to use, but not good* or good*** to use, etc. So, which of these concepts do we use to substitute the moral concept GOODNESS? It seems that substitutionists should at the very least answer this question. But there seems to be no non-arbitrary way for first-order error theorists to do so.

One perhaps obvious suggestion is that we should choose the new concept that causes the least disruption to our first-order ethical beliefs. That is, out of all possible concepts we can substitute for moral goodness, we should pick the one that entails and allows for moral truths that most closely resemble our beliefs about moral goodness. The problem with this is that there seems to be very little reason to make our current moral concept GOODNESS this important. After all, according to first-order error theorists, it is an incoherent concept.

Cowie might of course argue that regardless of its incoherence, our moral concept GOODNESS and the beliefs that follow from it still deliver the practical benefits we take moral beliefs to have. That might then be the reason why we should use it as a touchstone for picking our new concept. But that just seems like a roundabout way of claiming that our new concept GOODNESS* should also entail that our moral beliefs deliver this practical benefit. And there surely are many ways to conceive of GOODNESS* that would do so, leading us to the worry about arbitrariness once again.

So, substitutionism does not seem to work for the first-order error theorist either.

5.6.4. Abolitionism

Each alternative to abolitionism seems to come with its own set of problems for first-order error theorists. They therefore seem forced to accept abolitionism. This is a bitter pill to swallow. There is a reason why error theorists have proposed alternatives to abolitionism: our moral practices are practically valuable. In abolishing our moral practice, the abolitionist gives up these practical

benefits, such as morality's action-guiding and motivational functions. This is arguably why the abolitionist response is the least desirable answer to the "Now what?" problem. In contrast to the more traditional metaphysical arguments for moral error theory, the first-order error theorists seem forced to accept this undesirable answer.

5.7. Conclusion

I have argued that the first-order argument for moral error theory does not have the dialectical edge Cowie believes it has. This new argument is supposedly preferable to the traditional metaphysical argument for the error theory, since it does not lead to an impasse between realists and error theorists. Cowie is mistaken about this. This first-order argument does lead to such an impasse. This is because this argument relies on two assumptions, which, when either is rejected, entail that the first premise of the first-order argument goes unsupported. And whether one finds it plausible to reject these assumptions or not depends on one's realist or error theorist leanings. Thus, the first-order argument only trades one impasse for another.

And that is not all. The first-order error theory does worse than the traditional metaphysical error theories when it comes to a different issue: the "Now what?" problem. This is the problem of what the error theorist thinks we should do with our first-order ethics, once it turns out their metaethical theory is true. Out of all the answers to this question, only the least desirable one, abolitionism, is available to first-order error theorists.

In short, although the first-order argument for moral error theory is novel, dialectically it is not a step forward.

6. Conclusion

This thesis has put forward several new problems for anti-realism. The majority of these problems were developed by looking at the point of contact between anti-realist theories and the idea of conceptual moral truths. I will now briefly recap the conclusions from each chapter, so that I can then draw some broader conclusion from them.

Chapter 1 addressed the soundness of the most prominent argument in favour of moral error theory: the (metaphysical) argument from queerness. The chapter showed that even if we grant the error theorist the crucial premises of this argument, it still does not establish that there are no ethical facts. This is because there is an alternative to the error theorist's explanation for why we judge ethical facts to be queer. And this alternative explanation does not entail that ethical facts are ontologically suspect. The chapter concludes that the argument from queerness can at most establish metaethical scepticism, which states that it is most reasonable to suspend our judgment on the existence of ethical facts.

In chapter 2, I argued that realists who claim that there are conceptual moral truths should focus on negative conceptual moral truths instead of positive ones. This is because a realist account that focuses on the former can avoid a strong objection to accounts that focus on the latter. One potential drawback is that negative conceptual moral truths have less explanatory power than positive ones. But, I concluded that an account that focuses on negative conceptual moral truths does not entail any new explanatory burdens and can (to some extent) retain the benefits of the positive account.

Chapter 3 looked at a particular kind of negative conceptual moral truth: conceptual moral irrelevance claims. It argued that moral realists can give a satisfactory account of why our concepts bar certain descriptive features from being morally relevant, but that expressivist would struggle to do so. That is, it concluded that the arsenal of replies that expressivists have developed to other

objections does not seem to explain why certain descriptive features are conceptually morally irrelevant.

In chapter 4, I discussed the point of contact between conceptual moral truths and metaethical relativism. I focussed on two general approaches the relativist might adopt: a reconciliatory approach and a rejection approach. Although the reconciliatory approach might have seemed doomed from the start, it turned out to be potentially successful. Its success does depend on how much objectivity the relativist is willing to commit to. In contrast, at first glance the rejection approach might have seemed a natural fit for the relativist's theory. But this chapter showed that it requires the relativist to commit to an extreme form of relativism that simply disregards the strong intuition that there are negative conceptual moral truths.

Lastly, chapter 5 dealt with a recent proposal that suggests that positive conceptual moral truths might actually help establish a moral error theory. This first-order argument for moral error theory can also allegedly avoid an impasse with moral realists, an issue that the argument from queerness does seem to suffer from. But this chapter showed that this is false. The first-order argument relies on two crucial assumptions that. But whether these two assumptions seem plausible or not depends on one's realist or anti-realist leanings. Thus, the chapter concluded that this first-order argument does not avoid an impasse and therefore fails to be a step forward for the error theorist.

So, what can we conclude from these chapters? As we have just seen, this differs for each anti-realist theory. But there is some commonality in what these chapters entail for expressivism and relativism. Rejecting the idea of conceptual moral truths seems quite costly for both of them. Granted, rejecting positive conceptual moral truths might not pose too much of a problem. But rejecting negative (and/or conditional) conceptual moral truths turned out to be quite difficult. Chapter 3 and 4 therefore mostly focus on a strategy of reconciliation or accommodation. The idea being that if it is difficult to reject these conceptual moral truths, these anti-realist theories must make room for them.

The expressivists' reconciliatory efforts differ from those of the relativists, which is unsurprising given that both will try to draw on the unique strengths of their own theory. But in both cases we can see that the explanations that moral realists might give of why there are negative conceptual moral truths cannot be easily co-opted by these anti-realists. In the case of the relativist such copying of the realist's account would entail the spread of absoluteness beyond just the conceptual moral truths. This is a cost that is especially burdensome to relativists. In the expressivist's case, however, the realist's explanation for why certain specific descriptive features are conceptually morally irrelevant simply cannot be copied without giving up on anti-realism, since the realist's explanation relies on the existence of moral properties. So, the reconciliatory approach poses a unique challenge to both of these anti-realist theories. But, given the problems with the rejection approach, these theories will need to accommodate at least negative conceptual moral truths.

The moral error theorist, in contrast, is in a far better position to deal with negative conceptual moral truths. In chapter 2, we saw that a combination of the rejection and reconciliatory approach is successful for the error theorist. That is, we saw that, like other anti-realists, the error theorist can reject positive conceptual moral truths. But we also saw that, unlike the other anti-realists, she can accept negative (and conditional) conceptual moral truths. As chapter 1 makes clear, the problems for the error theorist lie elsewhere. Chapter 1 concluded that the most prominent argument for moral error theory does not support a moral error theory, but a form of metaethical scepticism. The first-order argument, discussed in chapter 5, might then be seen as a way to actually support a moral error theory. But that argument runs into an impasse with moral realism.

These three chapters thereby undermine the motivation for the error theory. Error theorists' most prominent argument, the argument from queerness, does not support a moral error theory. And the new and promising first-order argument for moral error theory ends in a stalemate with moral realism. Now, of course this does not exhaust the list of arguments that are made in favour of the error theory. But the argument from queerness is the main reason that turns people towards a

moral error theory. This means that many moral error theorists will have to find a new way to support their theory.

This leads to an obvious suggestion for further research: is there an alternative way to motivate the moral error theory? As was just mentioned, there have been such proposals. The most recent and prominent way to do so is by arguing that all other non-revolutionary metaethical theories fail to capture some essential element of our moral practice (Streumer, 2017). This then leaves error theory as the only viable option. However, we should note that this allegedly leads to an error theory about all normative judgments. That is something that hardly any error theorist wants to commit to. So, a more refined research question would be: is there an alternative way to motivate *just* the moral error theory?

What about future research on expressivism and relativism? One issue that might be worth exploring is whether the relativist's reconciliatory approach, in particular the demarcation solution, might be co-opted by expressivists. As mentioned, the big drawback for relativists is that this solution runs into a problem of spreading absoluteness. But given that expressivism is not a relativist theory, this should bother expressivists a lot less. Further research might look into the potential costs of this demarcation solution for expressivism.

Another issue that this thesis touches on every now and then but mostly sets to the side is the issue of the status of the moral concepts we use, as discussed by Eklund (2017). Eklund, and others, have observed that there is a lot of discussion about what is morally right or wrong, or good or bad, but that there are very few people that question whether these are the 'right' moral concepts to use. That is, instead of using the concept MORAL WRONGNESS we could be using the concept MORAL WRONGNESS*, or MORAL WRONGNESS**, or MORAL WRONGNESS***, etc. So, which of these concepts ought we to be using? Eklund explains that even asking this question properly is difficult, given that we might use OUGHT in asking it, but we could be using OUGHT*, or OUGHT**, or OUGHT***, etc.

Why does this matter? Well, to question the status of our moral concepts is to question the status of the conceptual moral truths that follow from them. That is, if we forget Eklund's worries for

a moment, then the conceptual truths that follow from our moral concepts seem robust. They seem to have some level of authority. They are the moral truths that we can be absolutely certain of. That is, we can be sure that, for instance, clasping one's hands three times an hour in itself cannot be morally good. That we can be so certain of this truth is one of the reasons why metaethical theories need to have something to say about conceptual moral truths. But this robustness is tied up with the robustness of our moral concepts. If it turned out that instead of the concept MORAL GOODNESS, we OUGHT* to be using the concept MORAL GOODNESS* and that concept entailed that it is in itself morally good* to clasp one's hands three times an hour, then it is unclear what sway the conceptual truth that this cannot be morally good still holds over us. Thus, the power or authority of the conceptual moral truths that follow from our moral concepts stands and falls with the 'rightness' of our moral concepts. So, the issue of 'choosing normative concepts' is hugely important to the matters discussed in this thesis. Further research should do more to address this.

Pending this further research, we can conclude that the anti-realist theories of expressivism and relativism struggle to deal with this idea of (negative) conceptual moral truth and its underlying intuition. And although moral error theory is better equipped to deal with it, it is not motivated as well as we might previously have thought.

7. Nederlandse Samenvatting

Denk eens terug aan de laatste keer dat je een moreel besluit moest nemen. Misschien moest je beslissen of je geld zou geven aan de collectant, die aan de deur kwam voor een goed doel. Misschien moest je twee keer nadenken over de aanschaf van een nieuw paar schoenen, aangezien de schoenen waarschijnlijk in een sweatshop gemaakt zijn. Misschien wilde je eten bestellen en moest je beslissen of je een gerecht met of zonder vlees zou nemen. Of misschien was het iets heel anders. Heb je je in zulke situaties weleens afgevraagd of er objectieve feiten bestaan over welke van je keuzes goed of slecht zijn? Het debat tussen meta-ethisch realisme en anti-realisme gaat over deze vraag.

Realisten geloven dat objectieve of geestonafhankelijke ethische feiten bestaan. Anti-realisten geloven dit niet. Het debat tussen hen is oud en er is voor beide theorieën iets te zeggen. Anti-realisme zou bijvoorbeeld een goede verklaring zijn van de observatie dat morele meningsverschillen wijdverspreid lijken. Immers, als er geen objectieve ethische feiten zijn, dan is er niets buiten ons wat een moreel meningsverschil kan beslechten. Anti-realisme geeft ook een metafysisch spaarzamer beeld van de wereld dan sommige realistische theorieën. Ethische feiten worden regelmatig als metafysisch problematisch gezien, omdat ze fundamenteel anders zouden zijn dan feiten uit andere onderzoeksgebieden, zoals de biologie, de scheikunde en de natuurkunde. En gegeven het feit dat we ethische feiten niet nodig lijken te hebben om ons morele gedrag te verklaren, omdat puur psychologische en evolutionaire verklaringen volstaan, lijkt er geen goede reden te zijn om hun bestaan te postuleren.

Natuurlijk zijn er ook bezwaren geformuleerd tegen anti-realisme. Anti-realisten hebben op hun beurt die bezwaren proberen te weerleggen en realisten hebben die weerleggingen weer proberen te weerleggen. Het doel van dit proefschrift is niet om een weerlegging van een weerlegging van een weerlegging te geven. Het doel is om nieuwe problemen voor het anti-realisme te formuleren, problemen die geen simpele oplossing hebben. Deze nieuwe problemen

beslechten niet het debat tussen realisten en anti-realisten. De strekking van dit proefschrift is bescheidener. De strekking van dit proefschrift is om anti-realisten onder druk te zetten en ze te dwingen iets nieuws te zeggen door middel van het formuleren van nieuwe problemen voor hun theorie.

Er zit samenhang tussen de meeste nieuwe problemen die in dit proefschrift worden uitgewerkt. Het overgrote deel van deze problemen ontstaat namelijk vanuit aandacht voor een specifieke intuïtie: de intuïtie dat sommige morele beweringen wel waar *moeten* zijn. Het gaat dan om beweringen als: ‘drie keer per uur je handen in elkaar slaan, kan op zichzelf niet moreel goed zijn’ (Foot, 2002). Of, ‘het is *pro tanto* verkeerd om een andere persoon recreatief te slachten’ (Cuneo and Shafer-Landau, 2014). Sommige filosofen denken dat deze intuïtie verklaard wordt door het feit dat deze morele beweringen conceptueel waar zijn. Dat wil zeggen, hun waarheid volgt uit de betekenis van de concepten die een rol spelen in de bewering. Deze beweringen zijn wat dat betreft complexere en morele varianten van de bewering “een vrijgezel heeft geen romantische partner”. Hoofdstukken 2 tot en met 5 behandelen het raakvlak tussen deze mogelijke conceptuele morele waarheden en verschillende anti-realistische theorieën.

Hoofdstuk 1 behandelt de onderbouwing van de morele foutenleer. Deze anti-realistische theorie omvat twee basisstellingen. Ten eerste, wanneer wij morele oordelen vellen schrijven wij een morele eigenschap toe aan een handeling. Ten tweede, deze morele eigenschappen komen nooit tot stand in de wereld. Dit betekent dat al onze morele oordelen onwaar zijn. Het meest prominente argument voor de morele foutenleer is het argument van vreemdheid. Dit argument stelt dat ethische feiten metafysisch vreemd zijn en dat ze voor het verklaren van ons moreel gedrag overbodig zijn. Kritiek op dit argument concentreert zich vaak op de kwestie of ethische feiten daadwerkelijk metafysisch vreemd zijn of niet.

Dit hoofdstuk laat echter zien dat zelfs als we accepteren dat ethische feiten metafysisch vreemd en overbodig zijn voor verklaringen van ons gedrag, de conclusie van de foutenleer niet volgt. De conclusie die we wel kunnen trekken uit het argument van vreemdheid is meer bescheiden,

namelijk dat we ons oordeel over of ethische feiten bestaan of niet moeten opschorten. Dit betekent dat dit argument niet een morele foutenleer, maar een vorm van meta-ethisch scepticisme of agnosticisme onderbouwt. Dit is een reden voor optimisme, aangezien meta-ethisch scepticisme een minder grote bedreiging vormt voor onze morele praktijk dan de morele foutenleer. Volgens de morele foutenleer zijn al onze morele oordelen immers onjuist, terwijl volgens meta-ethisch scepticisme ze alleen maar voorbarig zijn.

Hoofdstuk 2 snijdt het sub-thema van conceptuele morele waarheden aan. Dit hoofdstuk bespreekt de meest recente verdediging van conceptuele morele waarheden door Cuneo en Shafer-Landau's (2014) en kritieken op hun theorie. Die kritieken ontstaan voornamelijk vanuit onvrede met een conclusie van hun theorie: de conclusie dat aanhangers van de morele foutenleer conceptueel verward zijn (Evers en Streumer, 2016). Deze conclusie volgt, omdat als er conceptuele morele waarheden zijn, dan moet iemand die denkt dat er helemaal geen morele waarheden zijn onze morele concepten niet volledig snappen. Want had die persoon dat wel gedaan, dan had die geweten dat er wel degelijk morele waarheden zijn.

De kritiek op de theorie van Cuneo en Shafer-Landau lijkt correct. Echter, stelt dit hoofdstuk, filosofen zoals Cuneo en Shafer-Landau kunnen deze kritiek vermijden door een onderscheid te maken tussen *positieve* en *negatieve* conceptuele morele waarheden. Positieve morele beweringen stellen dat iets goed of slecht is. Negatieve morele beweringen stellen dat iets niet goed of slecht is. Aanhangers van de foutenleer kunnen niet accepteren dat iets goed of slecht is, maar ze kunnen wel accepteren dat iets niet goed of slecht is. Volgens hen is immers niets goed of slecht. Dus wanneer filosofen zoals Cuneo en Shafer-Landau zich richten op negatieve conceptuele morele waarheden, dan volgt daaruit niet dat aanhangers van de foutenleer conceptueel verward zijn. Bovendien beargumenteer ik dat het verschuiven van de aandacht op deze manier niet leidt tot nieuwe bewijslasten of het verlies van de voordelen van de originele theorie van Cuneo en Shafer-Landau.

Hoofdstuk 3 richt zich op één specifieke vorm van negatieve conceptuele morele waarheden: conceptuele morele irrelevantie beweringen. Morele irrelevantie beweringen stellen dat bepaalde

descriptieve kenmerken er moreel niet toe doen. Bijvoorbeeld, iemand kan in een gesprek met een racist beweren dat louter de kleur van iemands huid moreel irrelevant is. Sommige van deze morele irrelevantie beweringen lijken conceptueel waar te zijn. Dat wil zeggen, ze hebben kenmerken van conceptuele waarheid: ze zijn *a priori* kenbaar, hun ontkenning verwart conceptueel competente mensen, etc. Dit betekent dat er conceptuele grenzen lijken te bestaan voor wat moreel relevant is. Meta-ethische theorieën moeten een verklaring geven voor waarom dat zo is.

Hoofdstuk 3 laat zien dat dit makkelijker is voor moreel realisme dan voor expressivisme, een vorm van anti-realisme. Expressivisme stelt dat wij non-cognitieve attitudes uiten wanneer we een moreel oordeel vellen. Dus als ik oordeel dat geld doneren aan een goed doel goed is, dan uit ik bijvoorbeeld een attitude van goedkeuring over geld doneren aan een goed doel (Blackburn, 1984, 1993). Deze analyse van morele beweringen belemmert de expressivist echter. De realist kan zich namelijk beroepen op het bestaan van morele eigenschappen om een verklaring te geven van de conceptuele grenzen aan wat moreel relevant is. De expressivist, als anti-realist, kan dat niet. Expressivisten moeten daarom hun eigen verklaring hiervan geven. Dit hoofdstuk laat zien dat veel aanvankelijk plausibele expressivistische verklaringen niet voldoen.

Hoofdstuk 4 bespreekt moreel relativisme. Deze theorie stelt dat er geen objectieve morele waarheden zijn en is daarom een anti-realistische theorie. De relativist stelt dat morele oordelen hooguit waar kunnen zijn relatief aan iets. Wat dat iets dan is, daar is veel discussie over. Belangrijke kandidaten zijn morele standaarden (persoonlijke of culturele), andere normatieve oordelen of doelen. Echter, als er conceptuele morele waarheden zijn, dan zijn er morele waarheden die waar zijn los van wat onze morele standaarden, andere normatieve oordelen of doelen zijn.

Ik bespreek twee manieren waarop de relativist met dit probleem kan omgaan: verzoening of verwerping. Een verzoenende strategie probeert de conceptuele morele waarheden binnen het relativisme te vangen. Het idee is om de conceptuele morele waarheden als objectief te beschouwen, maar alle andere morele oordelen als hoogstens relatief waar te zien. Een verwerpende strategie tracht de conceptuele morele waarheden te ontkrachten. De

verwerpingsstrategie lijkt het meest voor de hand liggend voor de relativist, maar er blijken toch flink wat nadelen aan deze strategie te zitten. Verassend genoeg lijkt de verzoenende strategie, ook al heeft het een aantal nadelen, een redelijk alternatief voor de relativist.

Hoofdstuk 2 liet onder andere zien waarom het bestaan van positieve conceptuele morele waarheden zou betekenen dat aanhangers van de morele foutenleer conceptueel verward zijn. Hoofdstuk 5 richt zich op een recent voorstel dat positieve conceptuele morele waarheden juist kunnen helpen bij het onderbouwen van een morele foutenleer. De morele foutenleer die hieruit volgt zou een dialectische voordeel hebben op de morele foutenleer die volgt uit de meer traditionele metafysische argumenten voor deze theorie. Dat komt omdat de metafysische argumenten volgens velen doodlopen in een metafysische impasse met moreel realisten. Het nieuwe eerste-orde argument voor morele foutenleer zou deze impasse kunnen vermijden. Dit hoofdstuk laat zien dat dit onwaar is. Het eerste-orde argument voor morele foutenleer draait ook uit op een impasse met de moreel realist. Positieve conceptuele morele waarheden helpen de dialectische positie van de morele foutenleer dus niet vooruit. Hoofdstuk 5 concludeert dat een morele foutenleer onderbouwd door dit nieuwe eerste-orde argument dialectisch gezien zelfs slechter af is dan een morele foutenleer die onderbouwd is met traditionele metafysische argumenten.

Uit hoofdstukken 3 en 4 kunnen we concluderen dat het verwerpen van (negatieve) conceptuele waarheden lastig is voor het expressivisme en het relativisme. Verzoening met de conceptuele morele waarheden is echter ook ingewikkeld. Beide theorieën kunnen niet de verklaringen overnemen die de realisten tot hun beschikking hebben. En verklaringen vanuit de sterke punten van hun eigen theorie lijken ook niet te slagen. De morele foutenleer doet het beter op dit punt. In hoofdstuk 2 zien we dat aanhangers van de morele foutenleer de positieve conceptuele morele waarheden kunnen verwerpen en zich verzoenen met de negatieve conceptuele morele waarheden. Hoofdstukken 1 en 5 laten echter zien dat er een ander probleem is voor de morele foutenleer: het ontbreekt aan onderbouwing. Het meest prominente argument voor de foutenleer, het argument van vreemdheid, ondersteunt niet de morele foutenleer, maar meta-

ethisch scepticisme of agnosticisme. En een nieuw voorstel voor het onderbouwen van de foutenleer leidt, net als de metafysische argumenten voor morele foutenleer, tot een impasse met de moreel realist. De morele foutenleer is dus minder goed onderbouwd dan dat wij wellicht aanvankelijk dachten.

8. References

- Arrhenius, Gustaf (2009). One More Axiological Impossibility Theorem. In Lars-Göran Johansson, Jan Österberg, and Ryszard Sliwinski (eds.), *Logic, Ethics and All That Jazz. Essays in Honour of Jordan Howard Sobel*. Uppsala: Uppsala Philosophical Studies, 23-37.
- Bedke, M. (2010). Might All Normativity be Queer?. *Australasian Journal Of Philosophy*, 88(1): 41-58.
- Björnsson, G. and McPherson, T. (2014). Moral Attitudes for Non-Cognitivists: Solving the Specification Problem. *Mind* 123(489): 1-38.
- Blackburn, S. (1984). *Spreading the Word*. Oxford: Clarendon Press.
- Blackburn, S. (1993). *Essays in Quasi-Realism*. Oxford University Press.
- Blackburn, S. (1998). *Ruling Passions: A Theory of Practical Reasoning*. Oxford: Oxford University Press.
- Cohen, S. (1999). Contextualism, Skepticism, and the Structure of Reasons. *Nous*, 33(s13): 57-89.
- Cowie, C. (2020). A New Argument for Moral Error Theory. *Noûs*, (early view), 1-19.
- Cuneo, T. (2007). *The Normative Web*. Oxford: Oxford University Press.
- Cuneo, T. and Shafer-Landau, R. (2014). The Moral Fixed Points: New Directions for Moral Nonnaturalism. *Philosophical Studies* 171(3): 399-443.
- Eklund, M. (2002). Inconsistent Languages. *Philosophy and Phenomenological Research* 64(2): 251-275.
- Eklund, M. (2017). *Choosing Normative Concepts*. Oxford: Oxford University Press.
- Enoch, D. (2011). *Taking Morality Seriously: A Defense of Robust Realism*. Oxford: Oxford University Press.
- Evers, D. and Streumer, B. (2016). Are the Moral Fixed Points Conceptual Truths? *Journal of Ethics and Social Philosophy* (1): 1-9.
- Finlay, S. (2007). Four Faces of Moral Realism. *Philosophy Compass* 2(6): 820-849.

- Finlay, S. (2009). Oughts and Ends. *Philosophical Studies* 143(3): 315-340.
- Foot, P. (2002). *Virtues and Vices*. Oxford: Oxford University Press.
- Garner, R. (1990). On the Genuine Queerness of Moral Properties and Facts. *Australasian Journal of Philosophy* 68(2): 137-146.
- Geach, P. T. (1965). Assertion. *Philosophical Review* 74(4): 449-465.
- Gibbard, A. (1990). *Wise Choices, Apt Feelings: A Theory of Normative Judgment*. Cambridge: Harvard University Press.
- Gibbard, A. (2003). *Thinking How to Live*. Cambridge: Harvard University Press.
- Harman, G. (1975). Moral Relativism Defended. *Philosophical Review* 84(1): 3-22.
- Harman, G. (1977). *The Nature of Morality*. New York: Oxford University Press.
- Huemer, M. (2005). *Ethical Intuitionism*. London: Palgrave Macmillan.
- Ingram, S. (2015). The Moral Fixed Points: Reply to Cuneo and Shafer-Landau. *Journal of Ethics and Social Philosophy* 9(1): 1-5.
- Ingram, S.(2018). Are Moral Error Theorists Intellectually Vicious? *Journal of Ethics and Social Philosophy* 13(1): 80-89.
- Joyce, R. (2001). *The Myth of Morality*. Cambridge: Cambridge University Press.
- Kant, I. (1785). *Groundwork for the Metaphysics of Morals*. Oxford: Oxford University Press.
- Korsgaard, Christine M. (1996). *The Sources of Normativity*. Cambridge: Cambridge University Press.
- Lewis, D. (1996). Elusive Knowledge. *Australasian Journal of Philosophy*, 74(4): 549-567.
- Lutz, M. (2014). The 'Now What' Problem for Error Theory. *Philosophical Studies* 171(2): 351-371.
- MacFarlane, J. (2014). *Assessment Sensitivity: Relative Truth and its Applications*. Oxford: Oxford University Press.

- Mackie, J. (1977). *Ethics: Inventing Right and Wrong*. Harmondsworth: Penguin.
- Margolis, E. and Laurence, S. (2007). The Ontology of Concepts: Abstract Objects or Mental Representations? *Noûs* 41(4): 561-593.
- Olson, J. (2014). *Moral Error Theory: History, Critique, Defence*. Oxford: Oxford University Press.
- Parfit, D. (1984). *Reasons and Persons*. Oxford: Oxford University Press.
- Parfit, D. (2011). *On What Matters, Volume 2*. Oxford: Oxford University Press.
- Ridge, M. (2014). *Impassioned Belief*. Oxford: Oxford University Press.
- Shafer-Landau, R. (2003). *Moral Realism: A Defence*. Oxford: Oxford University Press
- Shepski, L. (2008). The Vanishing Argument From Queerness. *Australasian Journal of Philosophy*, 86(3): 371-387.
- Schroeder, M. (2008). *Being For: Evaluating the Semantic Program of Expressivism*. Oxford: Oxford University Press.
- Singer, P. (1975). *Animal Liberation: A New Ethics for Our Treatment of Animals*. New York: New York Review.
- Singer, P. (1989). All Animals Are Equal. In Regan, T. and Singer, P (eds.), *Animal Rights and Human Obligations*. Oxford University Press. 148-162.
- Singer, P. (2001) *Writings on an Ethical Life*. New York: HarperCollins Publishers Inc.
- Singer, P. (2011) *The Expanding Circle: Ethics, Evolution, and Moral Progress*. Princeton University Press.
- Singer, P. (2016). Why Speciesism is Wrong: A Response to Kagan. *Journal of Applied Philosophy* 33(1): 31-35
- Stanley, J. (2005). *Knowledge and Practical Interests*. Oxford: Clarendon Press.
- Street, S. (2006). A Darwinian Dilemma for Realist Theories of Value. *Philosophical Studies* 127(1): 109-166.

Street, Sharon (2008). Constructivism About Reasons. *Oxford Studies in Metaethics*. Vol. 3 Ed. R. Shafer-Landau. Oxford: Oxford University Press. 207-45.

Streumer, B. (2017). *Unbelievable Errors: An Error Theory About All Normative Judgments*. Oxford: Oxford University Press.

Thomson, J.J. (1971). A Defense of Abortion. *Philosophy and Public Affairs* 1(1): 47-66.

Wong, David B. (1984). *Moral Relativity*. Oakland: University of California Press.

Wright, C. (1995). Truth in Ethics. *Ratio*, 8(3): 209-226.

Curriculum Vitae

Erik Kassenberg

Rembrandt v. Rijnsstraat 49,
Groningen, 9718 PJ
erikkassenberg@gmail.com
+316 4999 4999

Education:

M.A. in Philosophy *Cum Laude*, University of Groningen - 2016

B.A. in Philosophy, University of Groningen – 2014

Publications:

Kassenberg, E. (Forthcoming). Debunking the Argument from Queerness. *Ratio*